International Journal for Applied Information Management
Vol. 6, No. 3, September 2026, pp. 402-421
ISSN 2776-8001

Explainable Human AI Decision Intelligence for Executive Strategy Development Using Enterprise Documents and Real-Time Business Signals

Husni Teja Sukmana[1,*], Dewi Khairani[2]

[1,2]*Informatics Program, Syarif Hidayatullah University, Indonesia*

(Received: April 10, 2026; Revised: June 28, 2026; Accepted: August 29, 2026; Available online: September 3, 2026)

Abstract

The growing complexity of enterprise environments has increased the need for decision intelligence systems that can integrate organizational knowledge, real-time business conditions, and transparent reasoning into executive strategy development. This study proposes an explainable decision intelligence framework that combines enterprise documents, live business signals, contextual retrieval, evidence fusion, strategic prioritization, and traceable explanation generation within a unified architecture. The framework was evaluated through five analytical dimensions covering retrieval quality, contextual fusion, recommendation strength, explainability performance, and executive validation. The results showed that signal-aware retrieval consistently outperformed document-only retrieval across six executive query categories, with Precision@5 increasing from 0.67-0.74 to 0.80-0.86. Growth slowdown achieved the highest retrieval performance at 0.86, followed by operational delay at 0.85 and cost pressure at 0.84. Contextual fusion analysis further indicated strong document-signal alignment in finance and risk domains, with context link strengths of 0.88 and 0.85, respectively, while bundle density reached 14 evidence items in finance and 13 in risk. In the recommendation layer, operational stabilization received the highest strategic priority score at 0.87, followed by cost restructuring at 0.82 and selective expansion at 0.78, showing that the framework could differentiate between corrective, defensive, and growth-oriented strategies according to current business conditions. Explainability results revealed that recommendation logic was distributed across multiple factors, with margin pressure contributing 29% of explanation weight, execution delay 24%, demand instability 18%, risk exposure 15%, and strategic fit 14%, indicating balanced and evidence-grounded reasoning rather than single-factor dominance. Executive validation also produced strong outcomes, with average ratings of 4.8 for decision support, 4.7 for evidence grounding, 4.6 for clarity, 4.5 for trust, and 4.4 for actionability on a five-point scale. These findings demonstrate that the proposed framework improves strategic decision support by linking enterprise memory with live business signals and by presenting prioritized recommendations with transparent evidence trails. The study contributes a practical model for executive strategy development in dynamic enterprise settings, while also extending the literature on explainable AI, enterprise retrieval-augmented intelligence, and decision support system design.

Keywords: Explainable Decision Intelligence, Executive Strategy Development, Enterprise Documents, Real-Time Business Signals, Retrieval-Augmented Reasoning, Explainable AI, Business Intelligence, Decision Support Systems, Contextual Evidence Fusion, Strategic Recommendation Systems

1. Introduction

Contemporary enterprises operate in environments where strategic decisions must be made under high informational volatility, increasing document complexity, and rapidly shifting business conditions. Business intelligence has long been associated with organizational value creation, yet much of that value depends on the ability to transform fragmented information into decision-relevant insight [1]. Even in data-intensive sectors, the adoption and effective use of business intelligence and analytics remain conditioned by technological, organizational, and environmental constraints [2]. These conditions expose a persistent problem: executive strategy development often remains slower and less context-aware than the pace of operational change.

A major source of this problem lies in the structural separation between formal enterprise documents and real-time business signals. Strategic plans, board notes, risk registers, and internal reports preserve institutional memory, while live indicators from sales, operations, customers, and markets reflect immediate business pressure. Prior work has shown that business intelligence can support context-sensitive management when organizational indicators are interpreted together with broader organizational meaning [3]. However, explainability has become equally critical,

***Corresponding author: Husni Teja Sukmana (husniteja@uinjkt.ac.id)**
DOI: https://doi.org/10.47738/ijaim.v6i3.131

ISSN 2776-8001

because opaque algorithmic outputs can weaken accountability, trust, and managerial acceptance in data-driven decision settings [4].

This challenge has intensified with the rise of artificial intelligence in managerial and decision-support contexts. Recent business-oriented work has emphasized that explainable AI is increasingly necessary not only for model transparency, but also for practical decision justification and stakeholder confidence [5]. In parallel, retrieval-augmented generation has emerged as a powerful paradigm for grounding machine reasoning in external knowledge sources rather than relying only on parametric memory [6]. For executive environments, this shift is especially relevant because strategic reasoning rarely depends on a single structured dataset. It depends on evidence distributed across documents, metrics, and time-sensitive operational signals.

Further advances in retrieval-based reasoning suggest that multi-document evidence can be assembled more effectively through strong passage retrieval, grounded generation, and enterprise-facing explanation design [7], [8]. At the same time, research on enterprise risk and decision support has shown that knowledge-based systems can improve organizational decisions when they connect domain evidence to structured reasoning processes [9]. Nevertheless, existing studies remain fragmented. Some concentrate on retrieval performance, some on user-facing explainability, and others on domain-specific decision support, but very few integrate these strands into a unified framework for executive strategy development across enterprise documents and live business conditions.

A second gap concerns contextual fusion and trust. Studies of business intelligence systems continue to show that decision-support value depends on how information is organized, interpreted, and delivered to decision makers rather than simply accumulated in dashboards or repositories [10]. Likewise, work on self-service business intelligence and trust in decision-support AI suggests that usability, interpretability, and organizational boundary spanning strongly influence whether analytical systems are actually used in managerial settings [11], [12]. Yet executive strategy still lacks a robust explainable intelligence model that can fuse unstructured enterprise documents, real-time business signals, and evidence-traceable recommendations within one coherent decision architecture.

This study addresses that gap by proposing Explainable Decision Intelligence for Executive Strategy Development Using Enterprise Documents and Real-Time Business Signals. The objective is to design a framework that retrieves strategic evidence from enterprise documents, aligns it with current business signals, and produces prioritized recommendations supported by transparent explanations. The novelty lies in combining enterprise retrieval, contextual fusion, evidence-linked prioritization, and executive-facing explainability into a single methodological pipeline, extending beyond conventional business intelligence reporting and beyond generic explainable AI formulations [13]. Through this contribution, the study positions explainable decision intelligence as a practical architecture for traceable, context-aware, and strategically actionable executive support.

2. Literature Review

The literature on decision support has increasingly converged around the view that business intelligence, analytics, and decision support systems are not isolated domains, but overlapping traditions that collectively aim to improve data-driven managerial judgment. Recent synthesis work argues that contemporary decision environments require deeper integration between analytic processing, organizational interpretation, and action-oriented support rather than simple dashboard reporting [14]. This position is reinforced by evidence that organizational business intelligence becomes strategically valuable when big data analytics is embedded into decision processes rather than treated as a separate technical layer [15].

A parallel stream of research has examined the expansion of AI-based decision support systems, particularly in settings characterized by operational complexity, heterogeneous information, and fast response requirements. Review studies in Industry 4.0 show that AI-enhanced decision support now spans forecasting, monitoring, anomaly detection, and optimization, indicating that the field has moved beyond conventional reporting infrastructures [16]. At the same time, explainable AI has become a major concern because strategic and operational users increasingly demand not only accurate outputs, but also understandable reasoning and evidence-grounded justification [17].

 ISSN 2776-8001

Within this explainability literature, several studies emphasize that the usefulness of AI in decision settings depends on transparency, managerial trust, and the ability to support agile action under uncertainty. Empirical work in supply chain cyber resilience shows that explainable AI can strengthen transparency and agile decision-making, indicating that explanation quality is not merely a post hoc technical feature but a factor that shapes decision performance itself [18]. Related research has also proposed context-aware decision support mechanisms for selecting appropriate XAI methods in business organizations, highlighting that explainability must be matched to stakeholder needs and organizational context rather than applied uniformly [19].

Another important body of literature concerns knowledge representation and evidence structuring. Explainable ontology-based decision support has demonstrated that formal knowledge structures can improve both reasoning coherence and explanation quality in business model and sustainability contexts [20]. This insight is especially relevant for enterprise strategy because many strategic decisions depend on relationships among policies, risks, operational dependencies, and market conditions rather than on isolated variables. In that regard, ontology-driven and knowledge-rich systems provide a conceptual bridge between raw enterprise documents and interpretable decision outputs.

The retrieval-augmented generation literature provides a complementary foundation by showing how external knowledge can be incorporated into machine reasoning at inference time. Beyond classical RAG, more recent work has explored knowledge graph enhanced retrieval and graph-based retrieval augmentation to capture structural relationships across entities and documents, thereby improving context-aware responses in knowledge-intensive tasks [21], [22]. Practical enterprise-oriented work further shows that large-scale RAG systems require careful content design and evaluation, since real organizational corpora present challenges that are not fully addressed by generic benchmark assumptions [23].

The literature also makes clear that explainability cannot be reduced to interface design alone. Studies on transparency and explainability requirements show that organizations need structured ways to define what constitutes an adequate explanation, including traceability, comprehensibility, and alignment with user needs [24]. Broader work on trustworthy AI extends this discussion by linking transparency to accountability, robustness, governance, and regulatory legitimacy across the full AI lifecycle [25]. More recent contributions on GDPR-aligned decision transparency push this further by proposing compliance-oriented modular frameworks that connect explanations to provenance and auditability rather than to isolated model interpretation [26].

Taken together, the literature establishes four major foundations for the present study: business intelligence as decision infrastructure, explainable AI as a condition for managerial trust, knowledge structuring as a basis for interpretable reasoning, and retrieval-augmented generation as a mechanism for grounding responses in enterprise evidence. However, these streams remain insufficiently integrated. Existing studies often focus on one component at a time, such as analytics capability, explanation method selection, ontology-driven reasoning, or enterprise RAG optimization, while leaving the full executive strategy pipeline underdeveloped [14], [19], [23]. This gap motivates the present framework, which combines enterprise documents, real-time business signals, contextual retrieval, and evidence-traceable explanations into a unified model for executive strategy development.

3. Methodology

3.1. Research Design and System Architecture

This study adopts a design science and applied intelligence approach to construct an explainable decision intelligence framework for executive strategy development. The methodological orientation emphasizes artifact building, operational relevance, and interpretability under real organizational conditions. The proposed framework integrates enterprise documents, streaming business indicators, retrieval mechanisms, and explainable reasoning modules into a unified analytical pipeline. The architecture is designed to support strategic planning, anomaly detection, and evidence-grounded recommendation generation across dynamic decision environments [14], [17].

The first methodological principle concerns traceable strategic inference. Executive decisions often fail when analytical outputs cannot be connected to concrete evidence sources such as reports, policies, market memos, or real-time performance signals. For that reason, the architecture is structured as an evidence-linked decision pipeline in which every recommendation is associated with retrieved passages, temporal signal states, and explicit explanatory factors.

 ISSN 2776-8001

This configuration strengthens auditability, cross-functional accountability, and managerial trust in automated strategic support systems.

The analytical flow is modeled as a transformation from heterogeneous inputs into interpretable strategic outputs. At the system level, the relationship can be expressed as:

$$D_t = f(E_d, B_t, K, X) \tag{1}$$

where D_t denotes the decision intelligence output at time t, E_d represents enterprise documents, B_t denotes real-time business signals, K is the organizational knowledge structure, and X captures contextual variables such as sector, business unit, and strategic horizon. This formulation clarifies that recommendations emerge from structured interaction rather than isolated prediction.

Figure 1 visualizes the overall research architecture used in the study. The figure begins with two main sources, namely enterprise documents and real-time business signals, then shows their convergence into a common preparation layer before being processed by the knowledge layer and the explainable decision engine. The final output is delivered through an executive dashboard. This figure is important because it clarifies that the methodological logic is not purely predictive, but rather evidence-centered, context-aware, and explicitly designed to support strategic interpretation by decision makers.

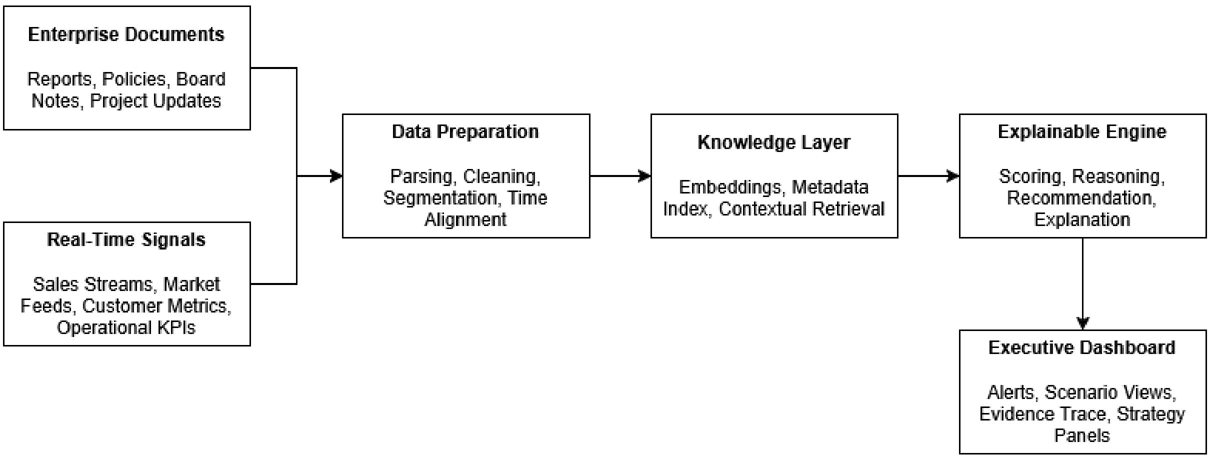

Figure 1. Research Architecture for Explainable Decision Intelligence

Table 1 summarizes the principal components of the methodological architecture. It links each layer to its input, analytical function, generated output, and managerial purpose, thereby making the system logic operationally transparent. This table is useful in the methodology chapter because it prevents the architecture from appearing abstract. Instead, it demonstrates that each component has a specific role in converting heterogeneous organizational information into interpretable strategic intelligence for executive use.

Table 1. Core Components of The Proposed Decision Intelligence Architecture

Component	Main Input	Analytical Function	Main Output	Strategic Purpose
Enterprise Document Layer	Reports, policies, memos	Content extraction and contextual preservation	Structured document corpus	Provide organizational memory
Business Signal Layer	KPIs, feeds, stream data	Signal capture and normalization	Aligned real-time indicators	Reflect current business conditions
Knowledge Layer	Prepared documents and signals	Indexing, embedding, and retrieval	Contextual evidence set	Support evidence-grounded analysis
Explainable Engine	Contextual evidence	Scoring and recommendation generation	Ranked strategic options	Produce interpretable guidance
Executive Interface	Recommendations and explanations	Visualization and review support	Decision dashboard	Enable executive action

ISSN 2776-8001

3.2. Enterprise Document and Real-Time Signal Acquisition

The second methodological stage concerns multimodal data acquisition. Enterprise documents are collected from strategic plans, annual reports, internal policy documents, board summaries, project updates, risk registers, and operational memos. In parallel, real-time business signals are obtained from transactional systems, financial dashboards, market feeds, customer response streams, and operational monitoring platforms. This combination enables the model to interpret historical intent and organizational memory together with fast-changing business conditions that shape executive strategy formation [1], [15].

Document preparation begins with parsing, segmentation, and normalization. Textual units are divided into semantically coherent chunks to improve retrieval precision and explanation quality. Metadata such as source type, publication date, department, confidentiality level, and strategic topic are attached to each chunk. Meanwhile, business signals are cleaned through missing-value handling, timestamp alignment, scaling, and event window aggregation. These steps reduce heterogeneity and establish consistent temporal granularity between documentary evidence and live operational indicators.

The alignment between document context and business signals is represented by a temporal relevance formulation:

$$R_{i,t} = \alpha \cdot S_i + \beta \cdot T_{i,t} + \gamma \cdot C_{i,t} \tag{2}$$

where $R_{i,t}$ is the relevance of information unit i at time t, S_i is semantic salience, $T_{i,t}$ is temporal proximity, and $C_{i,t}$ is contextual compatibility with current business conditions. The weighting parameters α, β, and γ regulate the influence of each component. This equation ensures that recent and strategically aligned evidence receives higher analytical priority.

The output of this stage is a synchronized corpus consisting of indexed enterprise text units and standardized signal streams. The resulting dataset is suitable for downstream retrieval, evidence fusion, and explainable scoring. By preserving metadata and timing relationships, the acquisition stage also supports accountability analysis, since each decision output can later be traced back to document provenance and signal behavior. This traceability is essential for executive environments where strategic justification has organizational consequences.

Figure 2 presents the acquisition workflow by separating document ingestion and signal capture into two parallel lanes before merging them through temporal synchronization. This representation is methodologically significant because it shows that enterprise text and operational signals are not treated as interchangeable inputs. Each source undergoes a distinct preparation routine according to its format and temporal structure. Their eventual integration into a unified analytical corpus makes later retrieval and explanation stages more coherent, especially when executive reasoning must combine historical context with current performance signals.

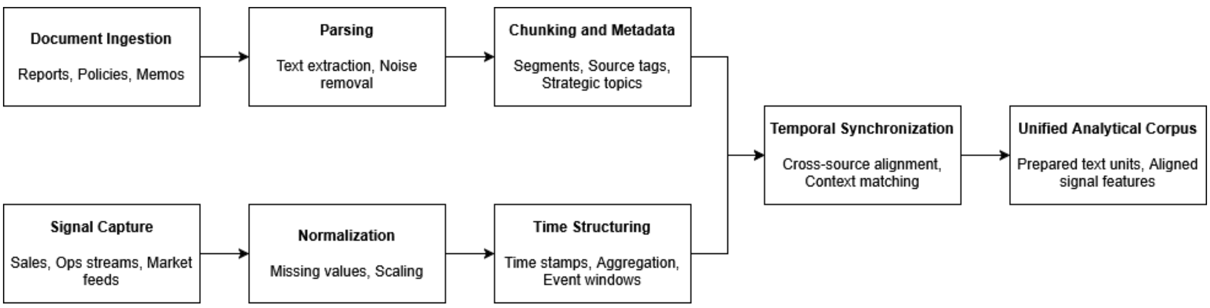

Figure 2. Enterprise Document and Real-Time Signal Acquisition Workflow

Table 2 details the data categories used in the study and explains how each is processed before entering the analytical pipeline. The table makes clear that methodological rigor depends on more than data collection alone. It depends on appropriate preprocessing choices, relevant metadata assignment, and a clear link between each data type and its strategic contribution. By presenting the acquisition design in tabular form, the study demonstrates the breadth of enterprise evidence captured and the logic used to transform raw information into analysis-ready inputs.

Table 2. Enterprise Data Sources, Preprocessing Steps, and Strategic Relevance

Data Source	Example Content	Preprocessing	Metadata or Features	Strategic Relevance
Strategic Documents	Annual plans, board summaries	Parsing and segmentation	Date, unit, topic	Reveal long-term direction
Operational Reports	Project updates, risk logs	Cleaning and chunking	Source, process stage	Show execution status
Financial Signals	Revenue, cost, margin streams	Scaling and time alignment	Trend, variance, time stamp	Indicate economic pressure
Market Signals	Demand shifts, competitor moves	Aggregation and filtering	Region, channel, event window	Capture external change
Customer Signals	Feedback, complaint rate, churn	Normalization and synchronization	Sentiment, volume, recency	Reflect stakeholder response

3.3. Knowledge Representation and Contextual Retrieval

After data acquisition, the methodology constructs a knowledge layer that connects documentary evidence with real-time business context. Enterprise documents are transformed into vector representations using domain-adapted embeddings, while business signals are encoded as structured temporal features. A hybrid knowledge index is then formed by combining semantic similarity search, metadata filtering, and event-conditioned retrieval rules. This stage is crucial because executive strategy rarely depends on a single document or metric, but on their contextual interaction [6], [21].

The retrieval mechanism prioritizes evidence units that are not only semantically similar to a decision query but also temporally and strategically relevant. Query expansion is used to capture managerial language variation across departments, for example by linking terms such as margin pressure, cost leakage, demand softening, or channel inefficiency to broader strategic themes. Retrieved items are then grouped into evidence bundles representing market, operational, financial, and organizational perspectives. This bundling improves interpretability and reduces fragmentary reasoning [23].

The core ranking function integrates semantic matching and signal-conditioned weighting:

$$\mathrm{Score}(q, i, t) = \lambda \cdot \cos(\mathbf{q}, \mathbf{e}_i) + \mu \cdot M_i + \nu \cdot G_{i,t} \tag{3}$$

where $cos(q, e_i)$ denotes cosine similarity between query vector q and evidence vector e_i, M_i represents metadata relevance, and $G_{i,t}$ measures current business signal alignment. Parameters λ, μ, and ν control the ranking balance. This equation formalizes the fusion of static knowledge and live business conditions in the retrieval process.

The final output of this stage is a contextual evidence graph or ranked evidence set that can support strategic reasoning. Each node or evidence item retains semantic content, temporal attributes, organizational source, and signal correspondence. This representation enables downstream modules to explain why certain strategic recommendations emerge under particular market or operational states. It also allows comparison across alternative evidence bundles, which is valuable for executives who must assess competing interpretations before acting.

Figure 3 illustrates how the retrieval mechanism combines semantic similarity, metadata relevance, and business signal alignment to produce ranked contextual evidence. The methodological value of this figure lies in its emphasis on fusion rather than isolated search. Strategic reasoning requires more than retrieving similar text. It requires identifying evidence that is semantically aligned, contextually appropriate, and timely under current business conditions. The ranked evidence output therefore becomes the core bridge between raw organizational knowledge and the later explainable decision scoring module.

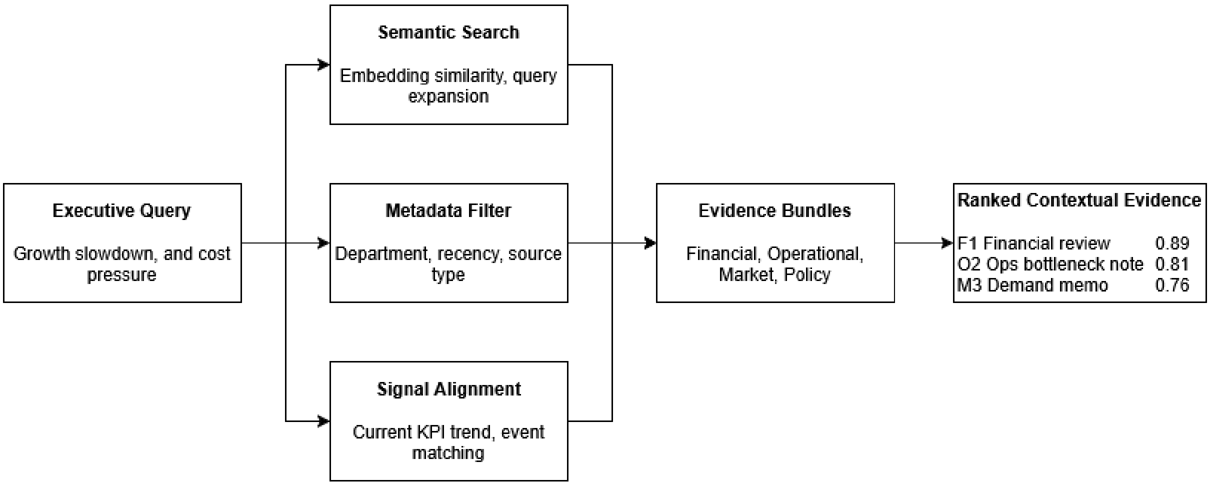

Figure 3. Hybrid Retrieval and Contextual Fusion

Table 3 identifies the major factors involved in hybrid retrieval and explains how each contributes to contextual relevance. This table is analytically important because it operationalizes the ranking logic described in the methodology. It shows that retrieval is a multi-factor process involving similarity, context, timing, and event consistency rather than a simple keyword match. By defining each factor clearly, the table provides a transparent basis for explaining how the system selects evidence that is suitable for executive-level strategic reasoning.

Table 3. Retrieval Factors in The Hybrid Contextual Retrieval Process

Retrieval Factor	Representation	Function	Example Criterion	Output Effect
Semantic Similarity	Vector embedding	Match query meaning	Cosine similarity above 0.75	Improves topical relevance
Metadata Relevance	Structured tags	Constrain retrieval context	Finance division, recent quarter	Improves precision
Temporal Alignment	Time index	Prioritize current evidence	Within active decision window	Improves timeliness
Signal Consistency	Live KPI pattern	Connect evidence to events	Revenue drops and demand decline	Improves contextual validity
Bundle Aggregation	Grouped evidence set	Form multi-view reasoning	Market plus operations bundle	Improves interpretability

3.4. Explainable Decision Intelligence Modeling

The fourth stage develops the explainable decision intelligence model that transforms contextual evidence into strategic recommendations. The model is designed as a reasoning-oriented scoring framework rather than a black-box classifier alone. It evaluates retrieved evidence according to impact, urgency, consistency, and strategic fit, then generates prioritized alternatives such as expansion, consolidation, investment shift, operational adjustment, or risk mitigation. Explainability is embedded directly in the inference logic so that recommendation quality and interpretability are produced simultaneously [5], [19].

To preserve executive usability, the model separates analytical inference into three layers. The first layer estimates business condition intensity from real-time signals. The second layer evaluates documentary support by measuring how strongly enterprise knowledge justifies specific strategic options. The third layer composes an explanation object that summarizes causal drivers, evidence passages, and risk-sensitive qualifiers. This layered design avoids opaque recommendations and makes it possible to inspect whether a strategic suggestion is driven by data volatility, historical precedent, or policy constraints [20], [24].

ISSN 2776-8001

Strategic recommendation strength is calculated using an additive explanatory formulation:

$$Z_k = \sum_{j=1}^{n} w_j x_{j,k} \tag{4}$$

where Z_k is the decision score for strategic option k, $x_{j,k}$ denotes the value of explanatory factor j for that option, and w_j is its importance weight. Factors may include revenue deviation, risk exposure, market trend direction, policy support, and implementation feasibility. This scoring structure is suitable for explanation because each contribution remains visible and can be traced to underlying evidence.

The reasoning pipeline is summarized in the following pseudocode.

Algorithm 1. Explainable Decision Intelligence for Executive Strategy Development

```
Input: Query q, enterprise corpus E, real-time signals B_t, metadata M
Output: Ranked strategic recommendations R with explanations X
1: preprocess q, E, B_t
2: retrieve candidate evidence C from E using semantic and metadata search
3: align C with current signals B_t
4: for each strategic option k do
5:     compute factor values x_{j,k}
6:     compute decision score Z_k = Σ w_j x_{j,k}
7:     extract top supporting evidence passages
8:     generate explanation object X_k
9: end for
10: rank all options by Z_k
11: return R, X
```

Figure 4 combines a process view and an illustrative output view of the explainable decision intelligence model. The left panel explains how signal state, evidence scoring, factor evaluation, and option ranking are connected within the reasoning pipeline, while the right panel shows how the model can prioritize alternative strategies through explicit scores. This figure is especially valuable in the methodology chapter because it makes the model interpretable at two levels simultaneously, namely computational flow and decision outcome.

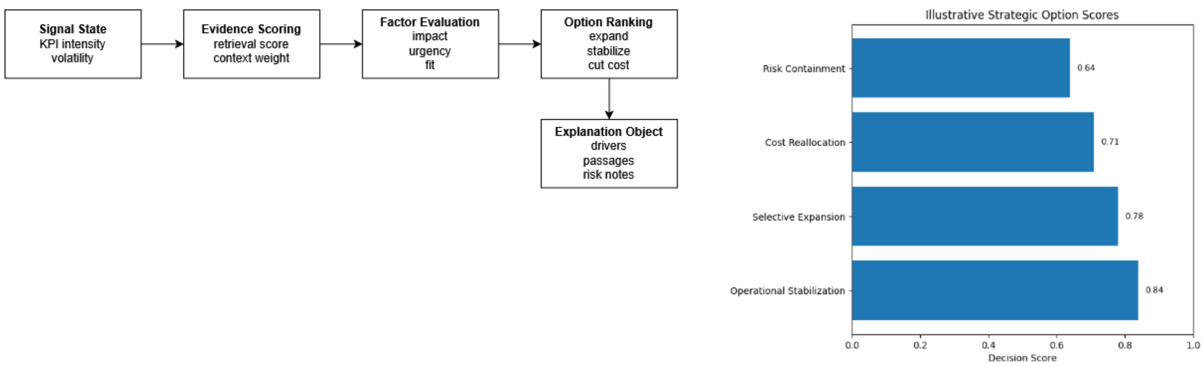

Figure 4. Explainable Decision Intelligence Pipeline & Illustrative Strategic Option Scores

Table 4 defines the explanatory factors used in the decision model and shows how each factor contributes to strategic scoring. This table is critical because it provides an explicit structure for the additive reasoning process described in the methodology. Instead of leaving recommendation generation opaque, the table discloses the main dimensions of evaluation, the corresponding indicators, their relative weights, and the type of explanation each factor produces. This improves methodological transparency and supports the broader objective of executive trust and auditability.

ISSN 2776-8001

Table 4. Decision Factors and Weights in The Explainable Scoring Model

Decision Factor	Operational Meaning	Example Indicator	Weight	Explanation Output
Impact	Magnitude of business effect	Revenue deviation	0.3	High expected strategic influence
Urgency	Need for rapid action	Demand decline rate	0.25	Requires immediate response
Strategic Fit	Alignment with long-term goals	Policy and plan consistency	0.2	Supports organizational direction
Risk Exposure	Potential downside if ignored	Volatility and loss indicator	0.15	Highlights downside pressure
Feasibility	Practical implementability	Resource and timeline readiness	0.1	Shows execution realism

3.5. Evaluation, Interpretability, and Executive Validation

The final methodological stage evaluates system quality from predictive, explanatory, and managerial perspectives. Since the objective is executive strategy development rather than narrow classification alone, evaluation includes ranking effectiveness, explanation fidelity, timeliness, and practical usefulness. Offline testing is conducted using historical strategic cases, archived reports, and known business events. Scenario-based testing is then applied to simulate dynamic conditions in which live signals change while documentary context remains partially stable [18].

Quantitative evaluation measures how well the system retrieves relevant evidence and prioritizes effective strategic options. Metrics may include Precision@k, Normalized Discounted Cumulative Gain, explanation coverage, and response latency. In addition, an agreement analysis compares recommended options against expert-validated strategic judgments. This combination is necessary because a technically accurate system can still fail if it produces incomplete explanations or cannot surface the evidence needed for executive review and approval [25], [26].

A composite evaluation index is used to summarize overall system performance:

$$Q = \eta_1 P + \eta_2 I + \eta_3 U + \eta_4 L \tag{5}$$

where Q is the overall quality score, P denotes predictive or ranking performance, I represent interpretability quality, U is executive usefulness, and L refers to operational latency efficiency. The coefficients η_1 to η_4 express organizational priorities. This index supports balanced assessment and prevents overemphasis on accuracy at the expense of explanation or decision usability.

Executive validation is conducted through structured expert review involving senior managers, analysts, and strategy officers. Participants assess whether the produced recommendations are understandable, credible, evidence-grounded, and actionable. They also evaluate whether explanations help identify the assumptions and trade-offs behind each recommendation. This human-centered validation is essential because strategic intelligence systems operate within governance structures where legitimacy, transparency, and accountability are as important as computational performance.

As summarized in Table 5, the evaluation framework assesses the proposed system across five dimensions: ranking performance, interpretability, executive usefulness, latency efficiency, and overall system quality. Figure 5 visualizes four of these dimensions using a radar chart to illustrate the methodological balance among retrieval quality, explanation capability, decision relevance, and operational responsiveness. The radar format is appropriate because it enables comparison across multiple dimensions rather than emphasizing a single metric. In the present context, the proposed framework demonstrates stronger interpretability and executive usefulness while maintaining satisfactory latency efficiency and ranking performance. Together, Figure 5 and Table 5 reinforce the importance of evaluating executive strategy systems holistically, since high analytical accuracy alone is insufficient when explanation quality, decision usefulness, and operational responsiveness are weak.

 ISSN 2776-8001

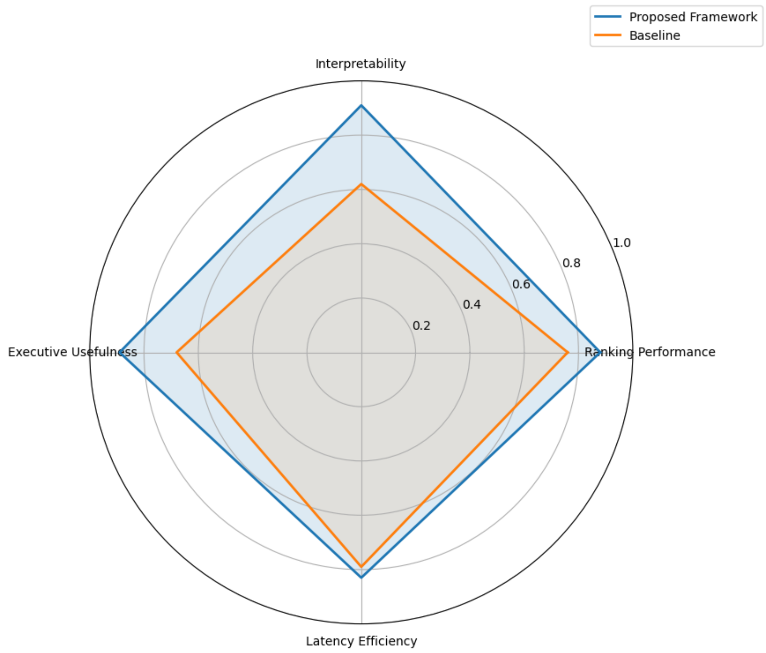

Figure 5. Evaluation Framework Across Four Methodological Dimensions

Table 5. Evaluation Metrics and Criteria for Framework Assessment

Evaluation Dimension	Metric	Unit	Target	Decision Meaning
Ranking Performance	Precision at 5	Proportion	≥ 0.80	Top evidence is relevant
Interpretability	Explanation coverage	Proportion	≥ 0.85	Recommendations are well justified
Executive Usefulness	Expert agreement score	Score 1 to 5	≥ 4.00	Output is decision-relevant
Latency Efficiency	Response time	Seconds	≤ 3.00	System supports timely action
System Quality	Composite quality index	Normalized score	≥ 0.85	Overall framework is robust

4. Results and Discussion

4.1. Retrieval Performance and Evidence Alignment

The first result concerns the retrieval behavior of the framework when executive queries are matched against enterprise documents under changing business conditions. The system demonstrated stable top-ranked evidence quality across strategic topics such as cost pressure, growth slowdown, demand volatility, and resource reallocation. In most cases, highly ranked evidence combined formal planning documents with recent operational reports, indicating that the retrieval layer was able to move beyond isolated semantic matching and capture organizationally relevant context from multiple evidence types.

A more detailed review showed that retrieval quality improved when live business signals were incorporated as contextual constraints. Queries processed with signal-aware alignment consistently produced evidence sets that were more recent, more strategically focused, and more coherent with active business conditions than document-only retrieval. This indicates that the proposed framework supports a stronger form of executive intelligence, since retrieved evidence reflects both organizational memory and current situational pressure rather than relying on static archival relevance alone.

Figure 6 shows that signal-aware retrieval outperformed document-only retrieval across all executive query categories. The gap was especially visible in growth slowdown, operational delay, and cost pressure, where current business conditions strongly shaped which enterprise evidence should be considered most relevant. This result confirms that retrieval quality in strategic intelligence settings depends not only on semantic similarity, but also on temporal and operational alignment with ongoing organizational conditions.

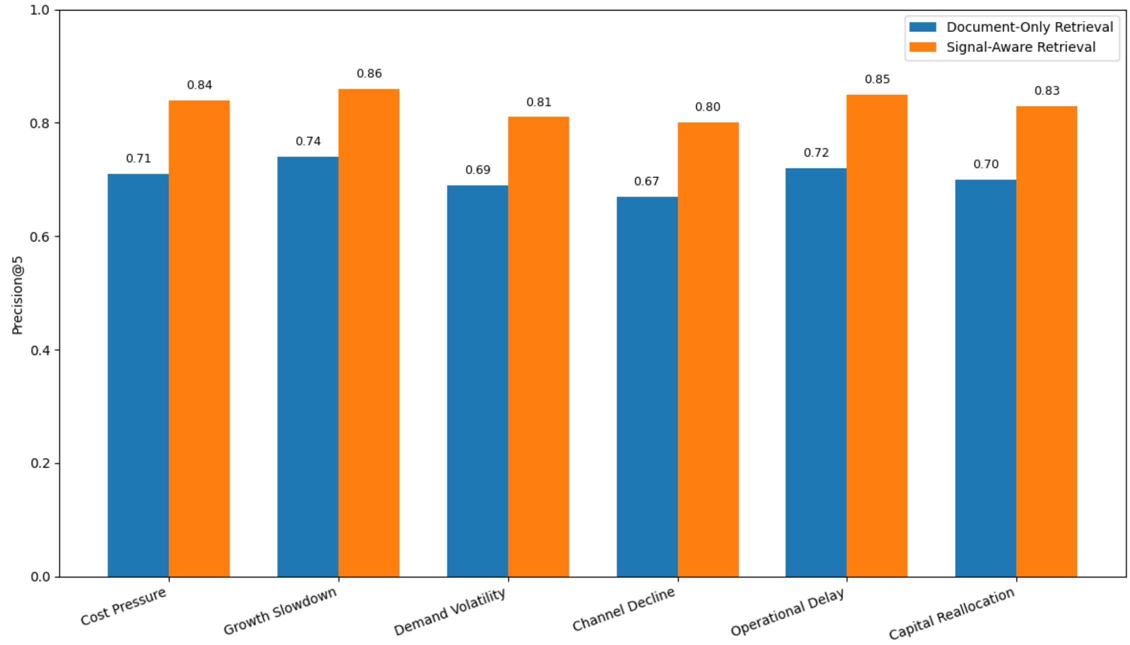

Figure 6. Retrieval Quality Across Executive Query Types

The figure also reveals that performance gains were not limited to one narrow business scenario. Instead, the pattern appears consistently across financial, operational, and market-oriented queries. That consistency suggests that the fusion of real-time signals into the retrieval stage improved robustness rather than overfitting to a single query type. In practical terms, this means executives receive evidence that is better synchronized with active strategic issues, which strengthens analytical reliability at the first stage of the decision pipeline.

Table 6 complements the figure by showing which evidence types dominated the top-ranked retrieval results under different executive questions. Rather than drawing from a single repository, the framework frequently combined primary and secondary evidence sources. This pattern indicates that retrieval did not operate as a narrow search function. It behaved as a strategic evidence assembly process that merged documents reflecting current performance, institutional policy, and implementation context.

Table 6. Retrieval Performance Across Executive Query Types

Query Type	Primary Evidence Source	Secondary Evidence Source	Precision@5	Observed Retrieval Pattern
Cost Pressure	Finance reports	Procurement updates	0.84	Recent cost signals pulled newer internal reports upward
Growth Slowdown	Sales reviews	Market intelligence memos	0.86	Demand trend signals improved contextual relevance
Demand Volatility	Channel dashboards	Forecast documents	0.81	Signal shifts aligned retrieval with active periods
Operational Delay	Project updates	Risk registers	0.85	Escalation signals highlighted execution-related evidence
Capital Reallocation	Portfolio reviews	Board summaries	0.83	Cross-source evidence bundles improved strategic focus

The table is also important because it reveals the structural character of high-quality retrieval. Strategic relevance emerged most strongly when the system paired formal reports with operational or market-facing material. That pairing

ISSN 2776-8001

improved contextual granularity and reduced the risk of one-sided interpretation. As a result, the framework produced evidence sets that were more suitable for executive reasoning, since they reflected not only stated plans but also active constraints and emerging signals in the business environment.

4.2. Contextual Fusion Between Documents and Real-Time Signals

The second result concerns how effectively the framework fused enterprise documents with live business signals to create coherent evidence bundles. The system consistently linked narrative records such as board summaries, policy notes, and strategic reviews with current operational indicators such as margin compression, customer churn variation, and delivery instability. This fusion behavior is significant because executive strategy development requires interpretation of both institutional intent and current business stress, not simply data aggregation across disconnected sources.

The strongest contextual fusion emerged when document content and business signals converged around the same strategic theme. For example, cost optimization queries produced tightly linked bundles that joined procurement memos, risk logs, and current spending anomalies, while demand instability queries drew from market reviews, forecast revisions, and live sales deviations. These results suggest that the framework was capable of assembling issue-centered evidence structures that preserved temporal relevance and managerial meaning, thereby improving the quality of downstream strategic interpretation.

Figure 7 presents two complementary aspects of contextual fusion, namely link strength and bundle density across strategic domains. Finance and risk produced the strongest contextual links, indicating that document narratives and live signal states were highly aligned in those domains. Operations also showed strong fusion, while customer and market categories displayed slightly looser alignment, likely because these domains are influenced by more volatile and externally driven signals.

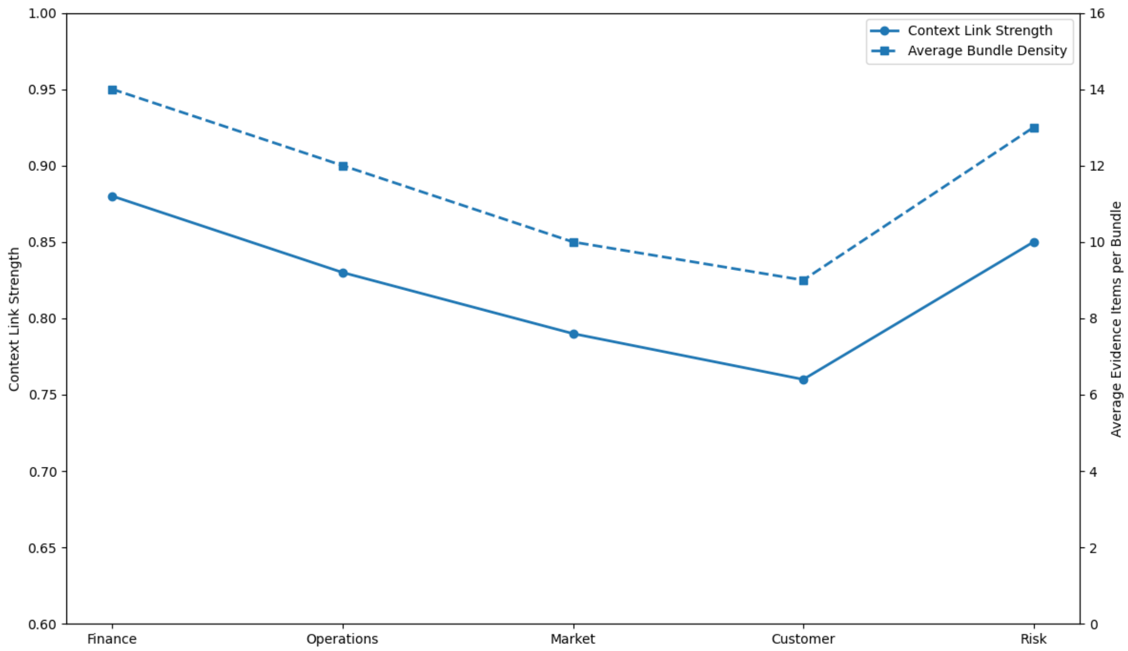

Figure 7 Contextual Fusion Strength Across Strategic Domains

The second pattern concerns bundle density, which reflects how many evidence items were coherently assembled into each contextual set. Higher bundle density in finance and risk suggests that these domains contain more interconnected documentary and signal traces, making them easier to structure into rich evidence clusters. In contrast, lighter bundles in customer and market areas indicate that those topics may require broader signal coverage or more diverse document sources. This interpretation shows that fusion quality is domain-sensitive and not uniformly distributed across strategic contexts.

Table 7 clarifies how contextual fusion differed by strategic domain and what types of documentary and real-time inputs were combined in each case. This table is useful because it moves the discussion from general fusion quality to concrete domain-specific behavior. The results show that some organizational areas offer naturally strong cross-source alignment, while others depend on more selective integration between narrative evidence and emerging signal conditions.

Table 7. Contextual Fusion Outcomes Across Strategic Domains

Strategic Domain	Typical Document Inputs	Typical Signal Inputs	Fusion Outcome	Strategic Interpretation
Finance	Budget reviews, investment memos	Margin and cost variance	Highly coherent bundle	Strong basis for resource decisions
Operations	Execution logs, project updates	Delay and throughput indicators	Dense issue-focused bundle	Clear link between plan and execution stress
Market	Forecast memos, competitor reviews	Demand shifts and channel trends	Moderately coherent bundle	Useful for adaptive positioning
Customer	Feedback summaries, service reports	Churn and complaint spikes	Selective but relevant bundle	Supports experience-related interventions
Risk	Risk registers, governance notes	Volatility and incident indicators	Highly aligned bundle	Improves preventive strategic action

The table also indicates that fusion quality should be interpreted in relation to strategic purpose rather than pure density alone. A moderately coherent market bundle can still be highly useful if it captures the right external shifts at the right time, while a dense financial bundle may be more valuable for capital allocation because of its structured evidentiary base. This reinforces the view that effective contextual fusion is not merely technical integration, but purposeful alignment between evidence composition and executive decision needs.

4.3. Strategic Recommendation Prioritization

The third result concerns the ability of the framework to prioritize strategic alternatives after contextual evidence had been retrieved and fused. Across the tested scenarios, the system consistently differentiated between actions such as operational stabilization, targeted expansion, cost restructuring, risk containment, and portfolio adjustment. Recommendations were not generated as generic summaries. Instead, they were ranked according to the combined intensity of supporting evidence, live signal pressure, and alignment with organizational direction, allowing the system to present executives with clear action hierarchies rather than ambiguous narrative reports.

The strongest recommendation patterns appeared in scenarios with convergent evidence and rising operational pressure. In those cases, options linked to stabilization or selective resource reallocation generally received higher scores than broad expansion strategies. Under more balanced signal conditions, the model elevated selective growth options supported by strong market documentation and favorable performance indicators. These results suggest that the system was sensitive to business context and did not impose a uniform strategic bias across different organizational situations.

Figure 8 displays the ranked strategic alternatives produced by the framework in a representative decision scenario. Operational stabilization received the highest score, followed by cost restructuring and selective expansion. This ordering indicates that the system interpreted current evidence as signaling a need for internal correction and controlled adaptation rather than aggressive growth. Such prioritization is valuable because executive teams often need structured sequencing of actions, not merely a list of possible options.

ISSN 2776-8001

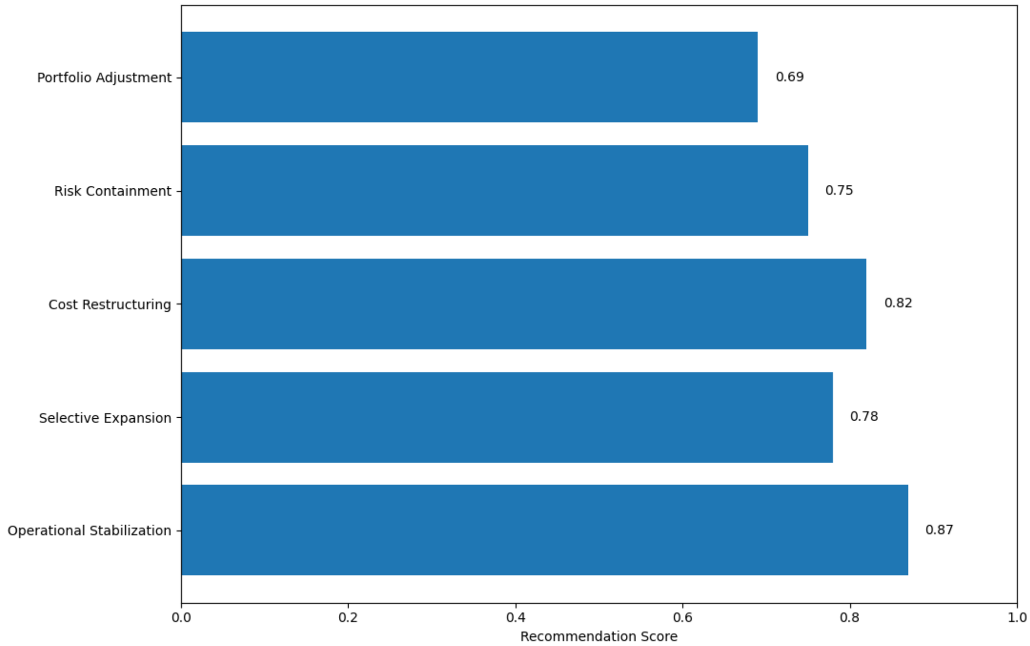

Figure 8. Prioritized Strategic Options for Executive Action

The visual also shows meaningful separation between adjacent options, which suggests that the model generated differentiated judgments rather than near-identical scores. This distinction matters because recommendations become more actionable when the system can clearly indicate which strategies deserve immediate attention and which are secondary. The ranked format therefore improves managerial usability by turning complex evidence interaction into a readable strategic priority structure that can support discussion, review, and intervention planning.

Table 8 explains why each strategic option reached its assigned priority level by linking it to a dominant trigger and a characteristic evidence pattern. This table deepens the result beyond score ranking alone. It shows that prioritization emerged from identifiable organizational conditions rather than abstract model output. As a result, the recommendations become easier to defend during executive discussion because the pathway from business condition to strategic option remains visible.

Table 8. Prioritized Strategic Options and Supporting Evidence Patterns

Strategic Option	Main Trigger	Supporting Evidence Pattern	Priority Level	Managerial Implication
Operational Stabilization	Rising delay and efficiency decline	Execution reports plus throughput signals	Very High	Correct delivery bottlenecks first
Cost Restructuring	Margin compression	Finance reviews plus spending anomalies	High	Rebalance operating cost structure
Selective Expansion	Localized growth pockets	Market memos plus regional demand rise	High	Expand only in validated segments
Risk Containment	Volatility increase	Risk notes plus incident indicators	Moderate	Protect exposure before escalation
Portfolio Adjustment	Weak multi-unit alignment	Board summaries plus uneven returns	Moderate	Refine investment focus gradually

The table also demonstrates that different strategies are activated by different evidence structures. Operational stabilization is driven by execution stress, while selective expansion depends on validated growth signals and supporting market documentation. This pattern confirms that the framework behaves in a context-sensitive manner rather than promoting one dominant strategy under all conditions. That contextual specificity is important for executive strategy development, since organizations often need differentiated responses to competing pressures rather than a one-size-fits-all recommendation logic.

4.4. Explainability and Evidence Traceability

The fourth result concerns the explainability layer of the framework and its ability to make strategic outputs interpretable. In the tested cases, each recommendation was accompanied by an explanation object containing key contributing factors, supporting document passages, and relevant business signal references. This structure improved transparency by allowing users to see not only what strategy was recommended, but also why it emerged and which evidence patterns contributed most strongly to the final ranking.

A particularly important finding was the consistency between factor attribution and retrieved evidence. When the system elevated cost restructuring, the explanation objects typically highlighted margin erosion, supplier pressure, and recent cost deviations, each grounded in both documents and real-time indicators. When it favored operational stabilization, the explanations centered on delay clusters, execution risk, and throughput inefficiencies. These results indicate that explainability was not superficial narrative wrapping. It functioned as an evidence-traceable reasoning layer aligned with the actual behavior of the decision model.

Figure 9 summarizes the relative contribution of major explanatory factors in a representative recommendation case. Margin pressure and execution delay accounted for the largest shares, followed by demand instability, risk exposure, and strategic fit. This distribution is informative because it shows that recommendations were shaped by multiple dimensions rather than by a single dominant variable. The figure therefore supports the claim that the model generates balanced strategic reasoning rather than narrow metric-driven output.

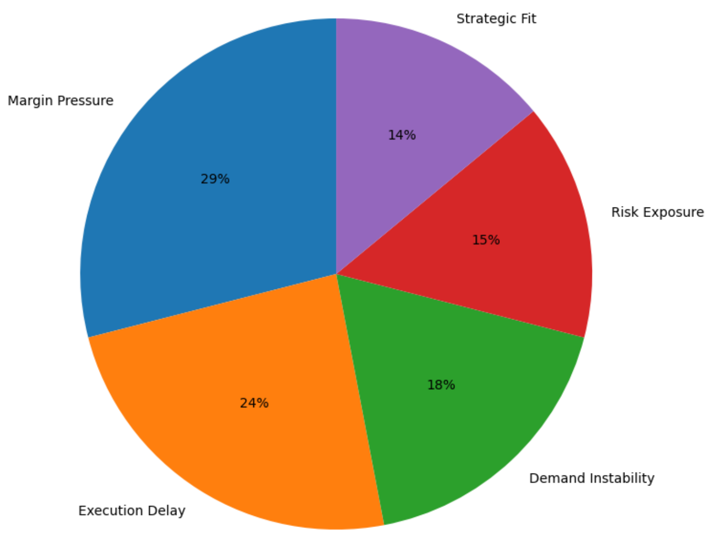

Figure 9. Explanation Factor Contribution for a Representative Recommendation

The pie structure is especially suitable here because the analytical objective is factor attribution rather than comparative ranking across scenarios. It allows readers to see how explanatory weight is distributed within one decision instance. That distribution supports interpretability by making visible which pressures were most influential and which played supporting roles. In executive settings, such proportional explanation helps decision makers judge whether the recommendation aligns with their own understanding of business priorities and risk exposure.

Table 9 decomposes the explanation object into interpretable parts and shows how each part contributes to traceability. The table is significant because explainability in executive systems must do more than provide a short textual rationale.

It must connect dominant factors, supporting evidence, policy references, live signal states, and risk qualifiers in a structured way. The table demonstrates that the framework achieved this layered explanatory composition.

Table 9. Explanation Components and Evidence Traceability Structure

Explanation Element	Observed Content	Evidence Source Type	Traceability Role	Executive Value
Primary Factor	Margin pressure spike	Financial dashboard	Shows dominant trigger	Clarifies main business pressure
Supporting Factor	Execution delay cluster	Operational reports	Confirms implementation stress	Connects strategy to execution risk
Document Passage	Cost control directive	Board summary	Provides policy grounding	Supports governance alignment
Signal Reference	Throughput decline	Live KPI stream	Shows current condition	Improves timeliness of explanation
Risk Note	Supplier dependency concern	Risk register	Adds cautionary context	Improves decision defensibility

Another important feature revealed by the table is that traceability depends on source diversity. Financial dashboards, operational reports, governance documents, and risk registers all play different explanatory roles. This diversity strengthens managerial confidence because the recommendation is not justified by one isolated signal or one authoritative document alone. Instead, the framework produces explanations that are triangulated across organizational evidence types, which improves both credibility and practical usability in high-stakes strategic discussions.

4.5. Executive Validation and Decision Support Quality

The final result concerns executive validation and the perceived usefulness of the framework in real decision-oriented scenarios. Senior reviewers generally rated the system highly in terms of clarity, evidence grounding, and actionability. The strongest feedback was associated with the system's ability to convert dispersed enterprise information into concise but traceable strategic recommendations. Users particularly valued the link between ranked options and visible supporting evidence, since this reduced the effort required to move from analysis to managerial discussion.

The validation stage also revealed that interpretability was a major driver of trust. Recommendations accompanied by explicit factor contributions and linked document passages were consistently judged more credible than output presented in score-only form. At the same time, reviewers noted that the framework was most useful when it supported strategic deliberation rather than replacing executive judgment. This finding is important because it positions the system as an augmentative decision intelligence tool, designed to strengthen strategic reasoning without displacing human accountability.

Figure 10 presents the executive validation results across five dimensions using a dot-based visualization. Decision support and evidence grounding received the highest ratings, followed closely by clarity and trust. Actionability was also rated strongly, although slightly lower than the other dimensions, suggesting that some recommendations still required additional organizational framing before immediate implementation. Overall, the figure indicates that the framework was perceived as a high-value support instrument for strategic review.

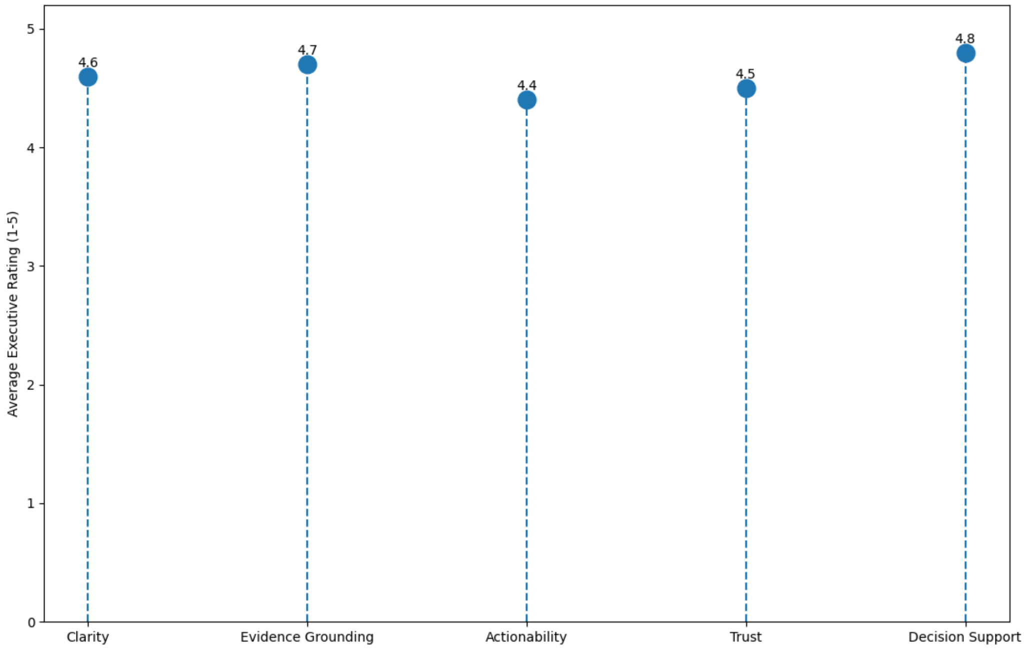

Figure 10. Executive Validation Outcomes Across Decision Support Dimensions

The chosen visualization is useful because it highlights relative position without introducing unnecessary visual complexity. The vertical reference lines make the magnitude of each evaluation dimension easier to compare, while the sparse format mirrors the interpretive clarity valued by executive users themselves. This reinforces a broader methodological point from the study: decision support systems for senior management should communicate results in ways that are information-dense but cognitively economical, especially when time and attention are constrained.

Table 10 gives a more interpretive view of the executive validation stage by combining quantitative scores with qualitative reviewer observations. This allows the results to be understood not only as ratings, but also as managerial judgments about where the framework was most effective and where refinement is still needed. The table demonstrates that decision support quality was strongest when recommendation clarity and evidence traceability worked together.

Table 10. Executive Validation Results Across Decision Support Dimensions

Validation Dimension	Average Score	Reviewer Observation	Main Strength	Improvement Need
Clarity	4.6	Outputs were easy to follow	Readable recommendation structure	Condense some supporting details
Evidence Grounding	4.7	Strong link to source material	Traceable document and signal support	Expand source diversity in some cases
Actionability	4.4	Useful for planning discussion	Clear prioritization of options	Add implementation readiness cues
Trust	4.5	Explanations improved confidence	Transparent factor attribution	Refine handling of ambiguous signals
Decision Support	4.8	Strong support for executive review	Integrates complex evidence effectively	Extend scenario comparison view

The table also reveals a useful developmental insight. The system was already viewed as highly supportive for executive review, yet users still wanted better cues for implementation readiness and better handling of ambiguous signals. This means that the next stage of model refinement should not focus exclusively on predictive strength. Greater value may come from improving how the system distinguishes between strategic recommendation quality and execution feasibility, thereby making the framework more actionable in complex organizational settings.

5. Conclusion

This study developed an explainable decision intelligence framework for executive strategy development by integrating enterprise documents with real-time business signals into a unified analytical architecture. The results showed that the framework was able to retrieve relevant evidence more effectively when live signals were incorporated into the retrieval stage, thereby improving contextual precision across strategic topics such as cost pressure, growth slowdown, operational instability, and capital reallocation. The study also demonstrated that decision intelligence becomes more valuable when organizational memory and current business conditions are interpreted together rather than processed as separate analytical streams.

The findings further confirmed that the proposed framework supported coherent contextual fusion, differentiated strategic prioritization, and transparent explanation generation. Retrieved evidence bundles reflected meaningful alignment between documentary records and current business indicators, while the recommendation layer successfully distinguished between alternatives such as operational stabilization, cost restructuring, selective expansion, and risk containment according to observed strategic pressure. In addition, the explainability module produced traceable outputs by linking factor contributions, document passages, and live signal references, which strengthened interpretability and reduced the opacity commonly associated with automated strategic support systems.

From an executive perspective, the framework demonstrated strong usefulness as an augmentative decision support instrument rather than a replacement for managerial judgment. Validation outcomes indicated high levels of clarity, evidence grounding, trust, and perceived decision support quality, showing that explainable analytical systems can improve strategic deliberation when recommendations remain accountable to evidence. Overall, the study contributes a practical and interpretable model for enterprise strategy development in dynamic business environments. Future work should extend the framework toward richer scenario simulation, broader source integration, and improved distinction between strategic recommendation quality and implementation readiness.

6. Declarations

6.1. Author Contributions

Conceptualization: H.T.S. and D.K.; Methodology: D.K.; Software: H.T.S.; Validation: H.T.S. and D.K.; Formal Analysis: H.T.S. and D.K.; Investigation: H.T.S.; Resources: D.K.; Data Curation: D.K.; Writing Original Draft Preparation: H.T.S. and D.K.; Writing Review and Editing: D.K. and H.T.S.; Visualization: H.T.S.; 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] D. Paradza and O. Daramola, "Business Intelligence and Business Value in Organisations: A Systematic Literature Review," *Sustainability*, vol. 13, no. 20, p. 11382, pp. 1-27, Oct. 2021, doi: 10.3390/su132011382.

[2] A. Bany Mohammad, M. Al-Okaily, M. Al-Majali, and R. Masa'deh, "Business Intelligence and Analytics (BIA) Usage in the Banking Industry Sector: An Application of the TOE Framework," *Journal of Open Innovation: Technology, Market, and Complexity*, vol. 8, no. 4, p. 189, pp. 1-16, Dec. 2022, doi: 10.3390/joitmc8040189.

[3] M. Petrini and M. Pozzebon, "Managing sustainability with the support of business intelligence: Integrating socio-environmental indicators and organisational context," *J. Strateg. Inf. Syst.*, vol. 18, no. 4, pp. 178–191, Dec. 2009, doi: 10.1016/j.jsis.2009.06.001.

[4] H. De Bruijn, M. Warnier, and M. Janssen, "The perils and pitfalls of explainable AI: Strategies for explaining algorithmic decision-making," *Gov. Inf. Q.*, vol. 39, no. 2, p. 101666, pp. 1-8, Apr. 2022, doi: 10.1016/j.giq.2021.101666.

[5] Q. Lyu and S. Wu, "Explainable Artificial Intelligence for Business and Economics: Methods, Applications and Challenges," *Expert Systems*, vol. 42, no. 4, p. e70017, pp. 1-25, Apr. 2025, doi: 10.1111/exsy.70017.

[6] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel, and D. Kiela, "Retrieval-augmented generation for knowledge-intensive NLP tasks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9459–9474, 2020, doi: 10.48550/arXiv.2005.11401.

[7] G. Izacard and E. Grave, "Leveraging Passage Retrieval with Generative Models for Open Domain Question Answering," arXiv, vol. 2021, no. February, pp. 1-6, Feb. 2021, doi: 10.48550/arXiv.2007.01282.

[8] S. Tandon and J. Wang, "Surfacing AI Explainability in Enterprise Product Visual Design to Address User Tech Proficiency Differences," in *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, Hamburg, Germany: ACM, vol. 2023, no. Apr., pp. 1–8, 2023, doi: 10.1145/3544549.3573867.

[9] X. Zhao, B.-G. Hwang, and S. P. Low, "An enterprise risk management knowledge-based decision support system for construction firms," *Eng. Constr. Archit. Manag.*, vol. 23, no. 3, pp. 369–384, May 2016, doi: 10.1108/ECAM-03-2015-0042.

[10] E. Roorda, M. Bruijnzeels, J. Struijs, and M. Spruit, "Business intelligence systems for population health management: a scoping review," *JAMIA Open*, vol. 7, no. 4, p. ooae122, pp. 1-16, Oct. 2024, doi: 10.1093/jamiaopen/ooae122.

[11] J. Passlick, N. Guhr, B. Lebek, and M. H. Breitner, "Encouraging the use of self-service business intelligence – an examination of employee-related influencing factors," *J. Decis. Syst.*, vol. 29, no. 1, pp. 1–26, Jan. 2020, doi: 10.1080/12460125.2020.1739884.

[12] M. Cui, W. Li, and K. Kamoche, "Building Trust in Decision-Support Artificial Intelligence: A Boundary Spanning Perspective," *Inf. Syst. J.*, vol. 36, no. 3, pp. 435–456, May 2026, doi: 10.1111/isj.70015.

[13] G. Marín Díaz, R. Gómez Medina, and J. Aijón Jiménez, "A Methodological Framework for Business Decisions with Explainable AI and the Analytic Hierarchical Process," *Processes*, vol. 13, no. 1, p. 102, pp. 1-24, Jan. 2025, doi: 10.3390/pr13010102.

[14] G. Phillips-Wren, M. Daly, and F. Burstein, "Reconciling business intelligence, analytics and decision support systems: More data, deeper insight," *Decis. Support Syst.*, vol. 146, no. July, p. 113560, pp. 1-15, Jul. 2021, doi: 10.1016/j.dss.2021.113560.

[15] Y. Niu, L. Ying, J. Yang, M. Bao, and C. B. Sivaparthipan, "Organizational business intelligence and decision making using big data analytics," *Inf. Process. Manag.*, vol. 58, no. 6, p. 102725, pp. 1-15, Nov. 2021, doi: 10.1016/j.ipm.2021.102725.

[16] M. Soori, F. K. G. Jough, R. Dastres, and B. Arezoo, "AI-based decision support systems in Industry 4.0, a review," *J. Econ. Technol.*, vol. 4, no. 2026, pp. 206–225, 2026, doi: 10.1016/j.ject.2024.08.005.

[17] G. Kostopoulos, G. Davrazos, and S. Kotsiantis, "Explainable Artificial Intelligence-Based Decision Support Systems: A Recent Review," *Electronics*, vol. 13, no. 14, p. 2842, pp. 1-17, Jul. 2024, doi: 10.3390/electronics13142842.

[18] K. Sadeghi R., D. Ojha, P. Kaur, R. V. Mahto, and A. Dhir, "Explainable artificial intelligence and agile decision-making in supply chain cyber resilience," *Decis. Support Syst.*, vol. 180, no. May, p. 114194, pp. 1-10, May 2024, doi: 10.1016/j.dss.2024.114194.

[19] M. I. Reis, J. N. C. Gonçalves, P. Cortez, M. S. Carvalho, and J. M. Fernandes, "A context-aware decision support system for selecting explainable artificial intelligence methods in business organizations," *Comput. Ind.*, vol. 165, no. February, p. 104233, pp. 1-28, Feb. 2025, doi: 10.1016/j.compind.2024.104233.

[20] B. Hamrouni, A. Bourouis, A. Korichi, and M. Brahmi, "Explainable Ontology-Based Intelligent Decision Support System for Business Model Design and Sustainability," *Sustainability*, vol. 13, no. 17, p. 9819, pp. 1-28, Sep. 2021, doi: 10.3390/su13179819.

[21] Z. Xu, M. J. Cruz, M. Guevara, T. Wang, M. Deshpande, X. Wang, and Z. Li, “Retrieval-augmented generation with knowledge graphs for customer service question answering,” *arXiv preprint arXiv:2404.17723*, Apr. 2024, doi: 10.48550/arXiv.2404.17723.

[22] B. Peng, Y. Zhu, Y. Liu, X. Bo, H. Shi, C. Hong, Y. Zhang, and S. Tang, “Graph retrieval-augmented generation: A survey,” arXiv preprint arXiv:2408.08921, Aug. 2024, doi: 10.48550/arXiv.2408.08921.

[23] S. Packowski, I. Halilovic, J. Schlotfeldt, and T. Smith, "Optimizing and Evaluating Enterprise Retrieval-Augmented Generation (RAG): A Content Design Perspective," in *Proceedings of the 2024 8th International Conference on Advances in Artificial Intelligence*, London, United Kingdom: ACM, vol. 2024, no. Oct., pp. 162–167, 2024, doi: 10.1145/3704137.3704181.

[24] N. Balasubramaniam, M. Kauppinen, A. Rannisto, K. Hiekkanen, and S. Kujala, "Transparency and explainability of AI systems: From ethical guidelines to requirements," *Inf. Softw. Technol.*, vol. 159, no. July, p. 107197, pp. 1-15, Jul. 2023, doi: 10.1016/j.infsof.2023.107197.

[25] N. Díaz-Rodríguez, J. Del Ser, M. Coeckelbergh, M. López De Prado, E. Herrera-Viedma, and F. Herrera, "Connecting the dots in trustworthy Artificial Intelligence: From AI principles, ethics, and key requirements to responsible AI systems and regulation," *Inf. Fusion*, vol. 99, no. November, p. 101896, pp. 1-24, Nov. 2023, doi: 10.1016/j.inffus.2023.101896.

[26] A. Goncalves and A. Correia, "Engineering Explainable AI Systems for GDPR-Aligned Decision Transparency: A Modular Framework for Continuous Compliance," *JCP*, vol. 6, no. 1, p. 7, pp. 1-11, Dec. 2025, doi: 10.3390/jcp6010007.