International Journal for Applied Information Management
Vol. 6, No. 3, September 2026, pp. 440-458

Human AI Collaborative Knowledge Synthesis for Corporate Policy Analysis and Cross-Division Decision Alignment

Sridevi V[1,*], Amaleswari Rajulapati[2], Sambasiva Rao Pasam[3], D Lakshmi[4]

[1,2,4]*Department of EEE, AMET Deemed to be University, Chennai, India*

[3]*Scientist, National Institute of Ocean Technology, Chennai, India*

(Received: April 15, 2026; Revised: June 30, 2026; Accepted: August 31, 2026; Available online: September 5, 2026)

Abstract

Corporate policy environments are often characterized by fragmented documents, inconsistent terminology, overlapping authority, and division-specific procedural interpretations that weaken organizational alignment. This study proposes an intelligent knowledge synthesis framework to analyze corporate policy corpora and assess cross-division decision alignment through an interpretable semantic pipeline. The framework integrates document acquisition, preprocessing, policy representation, knowledge synthesis, contradiction detection, alignment scoring, and evaluation within a unified architecture. Using a corpus of 540 source documents consisting of corporate policies, standard operating procedures, internal memoranda, strategic directives, and implementation notes, the method transformed the data into 1,980 normalized policy units and generated 246 synthesized policy clusters with a mean semantic coherence of 0.87. Domain-level results showed the highest coherence in compliance (0.90), approval governance (0.89), and reporting (0.88), while resource allocation and risk control produced lower but still stable coherence scores of 0.84 and 0.86. Contradiction analysis identified 173 significant conflict cases and 94 structured exception relationships, with conditional exceptions representing the largest pattern, followed by direct conflict, authority clash, and procedural override. Inter-divisional analysis showed the strongest friction between operations and compliance, IT and legal, and finance and operations, indicating that policy misalignment frequently emerges from localized procedural adaptation rather than explicit policy rejection. Across 320 evaluated decision cases, the alignment model achieved a mean score of 0.84, with 71.6% of cases classified in the high-alignment band, 19.1% in the moderate band, 6.9% in the low band, and 2.5% in the critical band. Category-level evaluation showed the highest alignment for compliance (0.91) and reporting (0.89), while resource-related decisions recorded the lowest score at 0.77. Division-level analysis further revealed that legal and compliance maintained the strongest stability profiles, with mean alignment scores of 0.90 and 0.88 and low variance values of 0.012 and 0.015, whereas IT and operations displayed greater volatility, with mean scores of 0.80 and 0.78 and variance values of 0.031 and 0.036. Comparative evaluation demonstrated that the full framework outperformed the baseline and synthesis-only configurations across synthesis quality, conflict detection, alignment stability, and interpretive usefulness, reaching scores of 0.90, 0.87, 0.89, and 0.92, respectively. These findings show that intelligent knowledge synthesis can function as a robust analytical foundation for enterprise policy interpretation, contradiction diagnosis, and coordination-aware decision support across organizational divisions.

Keywords: Corporate Policy Analysis, Knowledge Synthesis, Enterprise Knowledge Graph, Decision Alignment, Cross-Division Governance, Contradiction Detection, Interpretable Decision Support

1. Introduction

Large organizations depend on policy documents to coordinate authority, compliance, operations, and risk across business units, yet those documents are rarely consumed as a coherent knowledge system. Instead, corporate policies, standard operating procedures, internal memoranda, and division-specific interpretations accumulate as fragmented textual artifacts that are difficult to reconcile during decision making. Prior work on enterprise knowledge graphs and early policy-oriented knowledge graph platforms has already shown that semantic structuring can improve knowledge access and analytical reuse [1], [2].

The problem becomes more serious when decisions must be aligned across divisions that operate with different vocabularies, priorities, and procedural constraints. In such settings, the challenge is not only retrieving relevant text, but also synthesizing dispersed policy meaning, identifying latent inconsistencies, and determining whether local

***Corresponding author: Sridevi V (asridevi_2005@yahoo.com)**
DOI: https://doi.org/10.47738/ijaim.v6i3.133

decisions remain compatible with enterprise-level intent. Recent studies on knowledge graph construction processes and NLP adoption in policy and government settings confirm that structured knowledge extraction and text intelligence are increasingly central to organizational analysis, yet they are still unevenly translated into decision alignment workflows [3], [4].

Existing literature provides several important building blocks, but these contributions remain methodologically disjoint. Retrieval-augmented generation has improved evidence-grounded access to external knowledge, while work on enterprise question answering has shown that knowledge graphs can increase trust, provenance, and governance in LLM-enabled systems [5], [6]. However, these streams are largely optimized for retrieval, response generation, or question answering rather than for synthesizing corporate policy logic into a decision-oriented structure that can expose cross-division alignment or misalignment.

A parallel stream of research has addressed policy and regulatory text more directly. Studies on environmental-policy-enterprise knowledge graphs, digital regulatory compliance platforms, and broader NLP applications in government demonstrate that policy text can be modeled computationally for analysis, prediction, and compliance support [7]-[9]. Yet these approaches typically focus on public regulation, domain-specific compliance processing, or single-scope policy interpretation. They do not fully address the internal corporate problem in which multiple divisions enact partially overlapping policies that must be interpreted jointly during operational and strategic decision making.

Another limitation in the literature concerns explainability. Knowledge-graph-based explainable AI and explainable business process analysis have shown that semantic relations can make machine-assisted reasoning more interpretable for organizational stakeholders [10], [11]. At the same time, enterprise risk management research has highlighted the usefulness of knowledge graphs for representing heterogeneous risk factors and their interdependencies [12]. Even so, explainability is often discussed after prediction or recommendation, not embedded into the earlier stages of policy synthesis, contradiction analysis, and cross-division alignment assessment where governance-critical judgments actually emerge.

This study addresses that gap by proposing Intelligent Knowledge Synthesis for Corporate Policy Analysis and Cross-Division Decision Alignment, a framework that treats policy analysis as a synthesis-and-alignment problem rather than a retrieval-only or compliance-only task. The objective is to transform fragmented corporate policy material into interpretable knowledge structures, identify contradiction and exception patterns across documents, and evaluate whether divisional decisions align with synthesized enterprise policy intent. The approach is informed by recent knowledge-graph-driven decision support research, but extends it toward corporate policy environments where governance, coordination, and traceability are inseparable [13].

The novelty of this paper lies in three contributions. First, it introduces a unified pipeline that connects policy acquisition, semantic normalization, knowledge synthesis, contradiction detection, and alignment scoring in one analytical architecture. Second, it shifts the focus from policy retrieval to policy coherence by modeling how enterprise rules are consolidated and compared against division-level decision intents. Third, it embeds interpretability into the framework so that alignment outcomes can be traced back to synthesized policy clusters and conflict structures. In this sense, the paper contributes a decision-intelligence perspective to enterprise policy analytics, positioning intelligent knowledge synthesis as an operational mechanism for improving coordination, consistency, and governance quality across divisions.

2. Literature Review

Recent literature shows that knowledge graph research has moved beyond generic semantic representation toward engineering concerns such as construction practice, scalability, and quality assurance. Surveys on explainable machine learning and knowledge graph quality control emphasize that graph usefulness depends not only on semantic richness, but also on completeness, consistency, provenance, and downstream interpretability [14], [15]. In parallel, work on best development practices for industrial knowledge graphs and manufacturing-oriented adoption reviews argues that enterprise deployment requires disciplined schema design, multi-source integration, and governance-aware implementation rather than ad hoc graph construction [16], [17].

 ISSN 2776-8001

A second stream of research has focused on operationalizing enterprise knowledge graphs in specific organizational environments. Studies on enterprise data asset graphs show that automated graph construction can be strengthened through graph attention and optimization pipelines, particularly when heterogeneous data assets must be normalized into decision-relevant structures [18]. Related work in Industry 5.0 extends this direction by placing human work, collaboration, ergonomics, and contextual process information inside the graph model itself, thereby widening the scope of enterprise semantics beyond machines and documents alone [19]. Research on ontology generation for urban decision support further suggests that large language models can accelerate formal knowledge construction for complex planning tasks, although such systems remain strongly domain-specific and scenario-bounded [20].

Policy-oriented knowledge graph studies provide an additional foundation for this paper. Sustainable-energy policy research demonstrates that policy texts can be converted into graph structures that support relationship extraction, legal-source tracing, and policy evolution analysis [21]. Similarly, graph-driven analysis of Sustainable Development Goal interlinkages shows that semantically aligned graph representations are effective for exposing dependencies, synergies, and trade-offs across complex governance indicators [22]. These contributions are important because they confirm that policy documents are not merely archives for retrieval, but structured sources of relational knowledge that can be mined analytically when entities, relations, and contextual metadata are modeled explicitly.

Another relevant cluster of studies concerns decision support. Tourism-oriented knowledge graph research shows that graph-based semantic integration can support preference discovery and recommendation-oriented reasoning in domains were user intent and contextual relations matter simultaneously [23]. Work on knowledge-graph-based explainable AI reaches a related conclusion from a different angle: graphs improve the intelligibility of computational outputs because they preserve relational evidence and make reasoning paths more inspectable [24]. Together, these studies suggest that knowledge graphs are particularly valuable when an analytical system must move from descriptive information organization toward decision support that remains explainable to human stakeholders.

The recent convergence of knowledge graphs and large language models introduces both opportunity and risk. Reviews of KGs in LLM pipelines argue that structured knowledge can reduce hallucination, improve grounding, and strengthen trust in knowledge-intensive systems, but they also note unresolved issues in integration, evaluation, and factual consistency [25]. This concern is especially relevant for policy reasoning, where contradictory or weakly aligned texts can lead to unreliable outputs. Research on contradiction detection in protocol specifications shows that graph-based representations can reveal direct and conditional inconsistencies across formal documents [26]. In a related policy domain, PoliGraph demonstrates that modeling entire privacy policies as integrated graphs enables contradiction analysis that sentence-level extraction alone can miss [27].

Taken together, the literature establishes strong foundations in knowledge graph engineering, domain-specific policy graph construction, explainable reasoning, and contradiction analysis. However, these strands remain only partially connected. Existing studies rarely combine policy synthesis, conflict identification, and division-level decision alignment within a single enterprise-oriented framework. Most prior work either emphasizes graph construction quality [15], [16], domain analytics [21], [22], or decision support in a narrower application space [23], [24]. The unresolved gap therefore lies in designing a method that can synthesize fragmented corporate policy knowledge, detect exceptions and conflicts across documents, and translate the resulting structure into cross-division decision alignment. That gap directly motivates the framework proposed in this paper.

3. Methodology

3.1. Research Design and System Architecture

The study adopts a design-oriented computational methodology to examine how intelligent knowledge synthesis can support corporate policy analysis and improve decision consistency across organizational divisions. The methodological architecture is arranged as a layered pipeline rather than a single predictive model. This structure enables the separation of data ingestion, semantic normalization, synthesis, alignment scoring, and evaluation. Such decomposition is important because policy interpretation and decision coordination involve different forms of logic and evidence [1], [14].

ISSN 2776-8001

The system architecture is organized into five sequential layers: document acquisition, preprocessing, policy representation, synthesis and alignment, and evaluation. Each layer performs a specialized transformation while preserving traceability to source documents. The architecture is intentionally modular so that updates in one component do not destabilize the entire pipeline. This modularity is also necessary in enterprise contexts where policy revisions, departmental vocabularies, and operational rules often evolve independently over time.

$$X^{(l+1)} = f(W^{(l)}X^{(l)} + b^{(l)}) \tag{1}$$

The equation expresses layered transformation across the methodological pipeline, where each stage maps the previous representation into a more structured semantic form. In this study, the formulation is not limited to neural computation, but is used as an abstract representation of progressive knowledge refinement. The function $f(\cdot)$ captures normalization, filtering, or semantic encoding operations depending on the layer. This framing clarifies that policy intelligence is produced cumulatively rather than extracted in a single step.

Figure 1 visualizes the end-to-end methodological architecture used in the study. The diagram shows that the framework is not constructed as a single black-box classifier, but as a layered analytical system that begins with document acquisition and ends with alignment and evaluation. This arrangement clarifies the dependency between stages and highlights traceability as a design principle. The feedback loop is important because it indicates that evaluation results can inform earlier phases, especially preprocessing, synthesis rules, and alignment calibration.

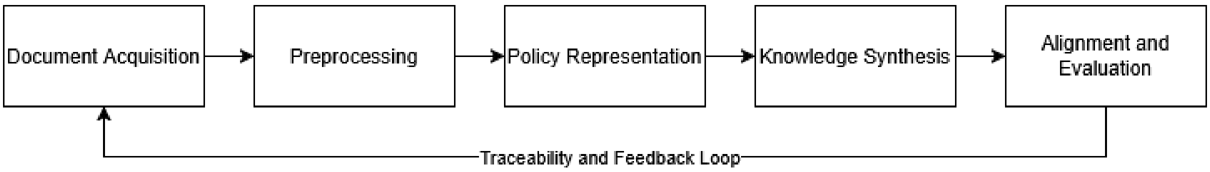

Figure 1. Methodological Architecture for Intelligent Knowledge Synthesis

To support transparent implementation, the architecture also defines clear interfaces between textual evidence and decision-oriented outputs. Every synthesized insight remains linked to policy segments, division ownership, and decision context. This prevents the framework from becoming a black-box recommendation mechanism. Instead, the architecture is designed as an interpretable knowledge infrastructure that can support policy review, conflict detection, and coordination analysis under changing enterprise conditions.

Table 1 complements the architecture diagram by specifying the analytical role of each methodological layer. The table makes the chapter operationally clearer because it links each stage with its input, processing logic, output, and purpose. This format is useful in a journal article because it shows how the framework moves from unstructured corporate policy material toward interpretable decision alignment outcomes. It also reinforces the modular nature of the method, which is essential for reproducibility and future extension.

Table 1. Methodological Layer Mapping

Layer	Main Input	Core Operation	Main Output	Analytical Purpose
Document Acquisition	Policies, SOPs, memos, directives	Collection and source mapping	Raw document repository	Establish the policy evidence base
Preprocessing	Raw documents	Cleaning, segmentation, metadata tagging	Structured policy units	Prepare comparable analytical units
Policy Representation	Structured policy units	Semantic encoding and normalization	Vectorized policy statements	Represent policy meaning consistently
Knowledge Synthesis	Encoded statements	Clustering, contradiction detection, graph construction	Synthesized policy graph	Integrate fragmented knowledge
Alignment and Evaluation	Policy graph and decision profiles	Scoring and comparative assessment	Alignment results and diagnostics	Measure cross-division consistency

 ISSN 2776-8001

3.2. Corporate Knowledge Acquisition and Policy Data Preparation

The data preparation stage focuses on transforming heterogeneous corporate materials into a unified analytical corpus. The study assumes that policy meaning is distributed across multiple artifacts rather than confined to formal policy manuals alone. Accordingly, the dataset includes corporate policies, standard operating procedures, internal memos, strategic directives, and division-level implementation notes. Each document is treated as a knowledge source with different degrees of authority, granularity, and operational relevance [3], [4].

Document preprocessing is carried out through segmentation, de-duplication, metadata tagging, terminology normalization, and clause-level extraction. Segmentation is necessary because policy conflict and alignment typically emerge at the level of provisions, exceptions, obligations, and action rules rather than full documents. Metadata fields include division source, document type, issue date, and policy scope. These attributes later support the synthesis engine in distinguishing global rules from local operational adaptations.

$$d_i = \{t_i, m_i, s_i\} \tag{2}$$

In this formulation, d_i denotes a processed document unit composed of textual content t_i, metadata m_i, and segmented policy statements s_i. The equation formalizes the transition from raw documents to structured policy objects suitable for downstream synthesis. This representation is essential because policy analysis requires not only semantic content, but also contextual attributes that influence interpretive priority, provenance, and cross-division comparability.

The preparation process also includes lexical harmonization to reduce ambiguity caused by department-specific terminology. For example, risk, escalation, exception handling, and approval status may carry different operational meanings across finance, operations, legal, and human resources. Controlled vocabulary mapping is therefore applied before semantic synthesis. This step improves the comparability of policy statements and reduces false contradictions that arise from inconsistent wording rather than genuine decision divergence.

Figure 2 presents the transformation of the policy corpus across successive preparation stages. The pattern is analytically meaningful because the number of units initially decreases after filtering, then rises sharply during segmentation, and stabilizes after normalization. This shows that preparation is not merely a reduction process. It is also a restructuring process that converts a limited set of documents into a much richer set of policy clauses suitable for semantic synthesis. The figure helps justify the methodological importance of segmentation and normalization.

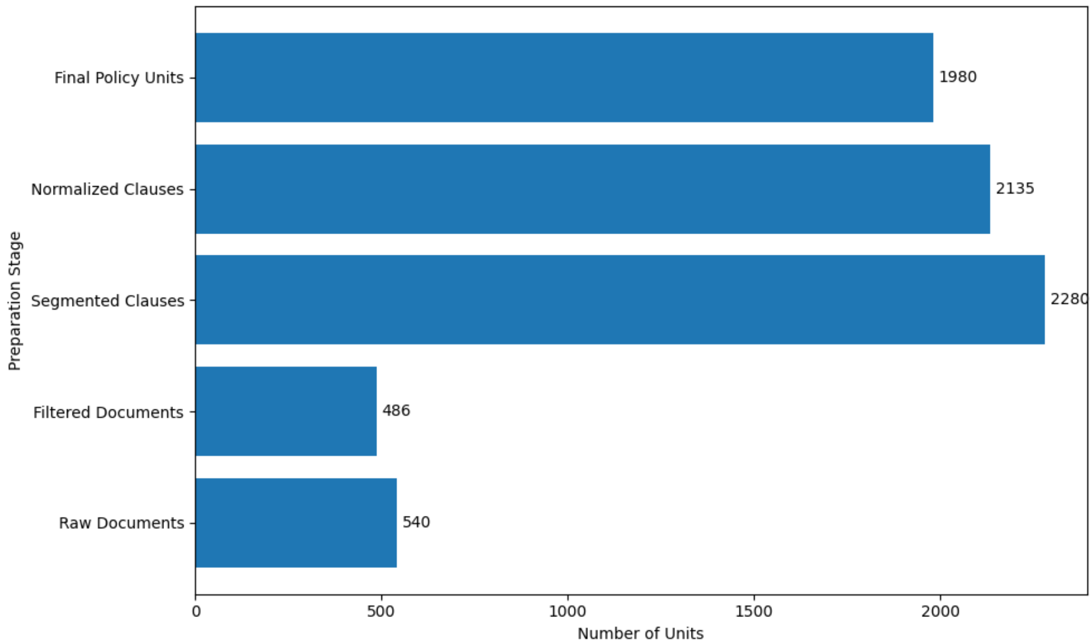

Figure 2. Retention and Expansion Across Policy Data Preparation

ISSN 2776-8001

Table 2 describes the structure of the prepared policy corpus and demonstrates that the data foundation is heterogeneous by design. Different source types contribute different granularities of evidence, from clause-level corporate rules to locally contextualized implementation notes. The table is important because it explains why metadata tagging is required and why a single preprocessing strategy would be insufficient. It also clarifies how each source contributes a distinct analytical role within the synthesis framework.

Table 2. Prepared Policy Corpus Structure

Source Type	Document Count	Unit of Analysis	Key Metadata	Prepared Output
Corporate Policies	120	Clause	Division, issue date, scope	Normalized policy units
Standard Operating Procedures	165	Procedure step	Owner, revision, process area	Segmented operational rules
Internal Memos	98	Directive sentence	Sender, date, purpose	Tagged decision directives
Strategic Directives	74	Instruction block	Priority level, division, validity	Priority-weighted statements
Implementation Notes	83	Action note	Division, activity, exception	Context-enriched local rules

3.3. Intelligent Knowledge Synthesis Engine

The knowledge synthesis engine is designed to convert fragmented policy statements into coherent and analyzable knowledge structures. The engine performs semantic encoding, similarity estimation, contradiction screening, and thematic aggregation. The central premise is that corporate policy intelligence emerges when related statements are interpreted jointly rather than reviewed in isolation. This is especially important when multiple divisions operationalize shared corporate goals through distinct procedural vocabularies and localized compliance practices [7], [15].

Semantic synthesis begins by representing each segmented statement in a shared embedding space. Statements with high semantic affinity are grouped into policy clusters, while conflicting clauses are flagged for inconsistency analysis. The synthesis process does not simply merge similar text. Instead, it weighs authority level, recency, and operational specificity before producing a consolidated representation. This allows the framework to distinguish policy reinforcement from policy collision, which is crucial for enterprise decision alignment [26], [27].

$$S_{ij} = \frac{v_i \cdot v_j}{|v_i||v_j|} \tag{3}$$

The similarity measure S_{ij} quantifies the semantic proximity between policy statement vectors v_i and v_j. Cosine similarity is suitable here because the goal is to compare directional meaning rather than raw textual magnitude. Within the synthesis engine, high similarity scores indicate conceptual convergence, while low scores may signal divergence, local specialization, or ambiguity. This metric serves as the initial filter before contradiction logic and synthesis weighting are applied.

The synthesis stage also produces a unified policy graph in which nodes represent policy statements and edges encode semantic support, overlap, or conflict. This graph-based representation enables localized reasoning over shared obligations and department-specific exceptions. It also improves interpretability because synthesized knowledge can be traced through explicit relational structures. The resulting graph becomes the main analytical substrate for assessing whether decision intents remain aligned with enterprise-wide policy logic.

Figure 3 shows the policy graph generated by the knowledge synthesis engine. The figure distinguishes supportive relationships from conflicting ones, allowing readers to see how fragmented policy statements become an interpretable relational structure. This representation is useful because corporate misalignment rarely appears as isolated clauses. It usually appears as interactions among overlapping rules, escalations, exceptions, and local overrides. The graph therefore gives a more faithful view of organizational policy logic than a simple document list or similarity matrix.

ISSN 2776-8001

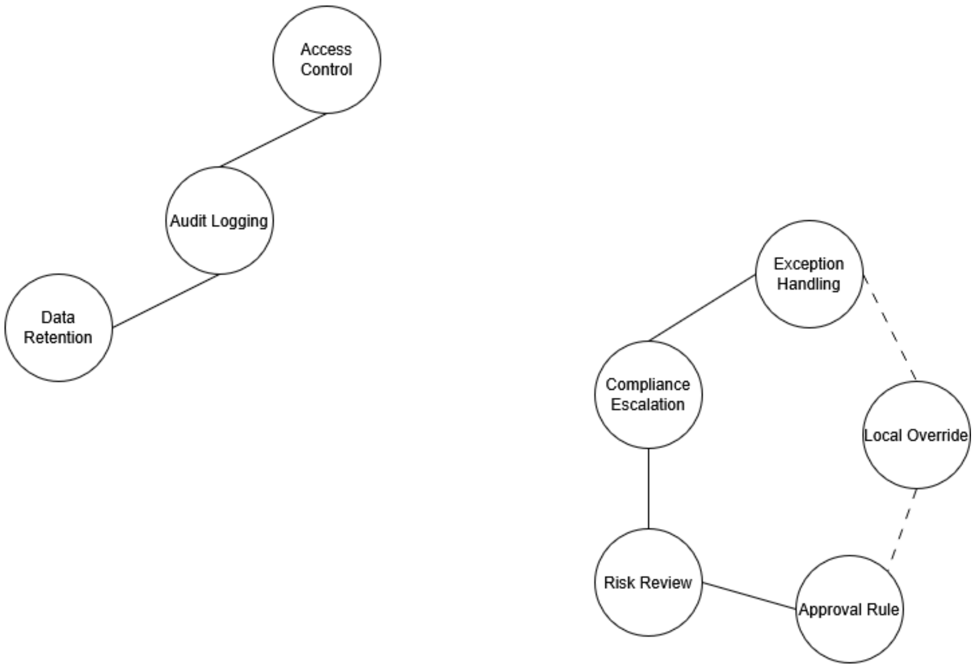

Figure 3. Policy Graph After Knowledge Synthesis

Table 3 decomposes the synthesis engine into its major computational components. This tabular description is methodologically important because the synthesis process includes more than semantic similarity alone. It combines representation learning, relation analysis, contradiction screening, and graph construction in a sequence that supports policy interpretation. By showing the input and output of each component, the table makes the framework auditable and helps readers understand how the final policy graph is constructed from initially unstructured textual evidence.

Table 3. Knowledge Synthesis Engine Components

Component	Main Function	Input	Output	Decision Relevance
Semantic Encoder	Transform text into comparable vectors	Segmented policy units	Statement embeddings	Supports semantic comparison
Similarity Estimator	Measure conceptual proximity	Statement embeddings	Similarity scores	Identifies convergence patterns
Contradiction Detector	Flag conflicting policy logic	Statement pairs	Conflict markers	Reveals inconsistency risks
Cluster Builder	Group related policy statements	Similarity and conflict outputs	Policy clusters	Forms consolidated themes
Graph Constructor	Encode policy relations structurally	Clusters and relation types	Policy graph	Enables interpretable reasoning

Algorithm 1. Intelligent Knowledge Synthesis for Policy Alignment

```
Input:
  D = set of prepared policy units
  theta = semantic similarity threshold
  omega = synthesis weight configuration
Output:
  G = synthesized policy graph
  C = consolidated policy clusters

Begin
  Initialize empty graph G
  Encode each policy unit d in D into vector v
```

ISSN 2776-8001


```
  For each pair (d_i, d_j) in D do
    Compute similarity S_ij
    If S_ij >= theta then
      Assign d_i and d_j to candidate cluster
      Add support edge to G
    Else
      Evaluate contradiction and exception rules
      If contradiction detected then
        Add conflict edge to G
      End If
    End If
  End For

  For each cluster c do
    Compute synthesis weight using authority, recency, and specificity
    Generate consolidated representation c*
    Store c* in C
  End For

  Return G, C
End
```

3.4. Cross-Division Decision Alignment Model

After policy knowledge has been synthesized, the next stage evaluates whether division-level decision intents remain aligned with the consolidated policy structure. Alignment is treated as a measurable relationship between local decision rationale and enterprise policy meaning. This approach recognizes that divisions often pursue different operational objectives, but these differences should remain consistent with shared governance principles. The alignment model therefore measures compatibility rather than superficial textual similarity alone [13].

Each decision case is represented as a structured intent profile containing goal statements, action constraints, responsible division, and anticipated impact. The synthesized policy clusters are then matched against these intent profiles to estimate alignment strength. The model gives separate attention to support, neutrality, and conflict signals. This three-way distinction is important because many decisions are not explicitly prohibited, yet still lack sufficient policy support for reliable cross-division coordination [24].

$$A_k = \alpha R_k + \beta C_k - \gamma F_k \tag{4}$$

Here, A_k denotes the alignment score for decision case k, while R_k represents policy relevance, C_k indicates consistency with synthesized policy clusters, and F_k captures conflict intensity. The coefficients α, β, and γ control the contribution of each component. This weighted structure reflects the fact that well-aligned decisions must exhibit both semantic relevance and normative consistency, while minimizing contradiction with authoritative policy signals.

The alignment model also supports cross-division comparison by aggregating scores across departments and decision categories. This aggregation reveals where misalignment is systemic rather than isolated. For example, repeated low alignment in procurement or compliance-related decisions may indicate unresolved policy ambiguity rather than local execution error. The model is therefore designed not only for case scoring, but also for organizational diagnosis and governance improvement across interdependent business functions.

Figure 4 presents the analytical structure of the cross-division decision alignment model. The figure shows that alignment assessment is generated by combining three inputs, namely synthesized policy clusters, decision intent profiles, and division metadata, before passing them through semantic matching, consistency checking, and conflict screening. These intermediate operations are then integrated within the alignment scoring engine to produce interpretable decision categories ranging from high alignment to low or critical alignment.

ISSN 2776-8001

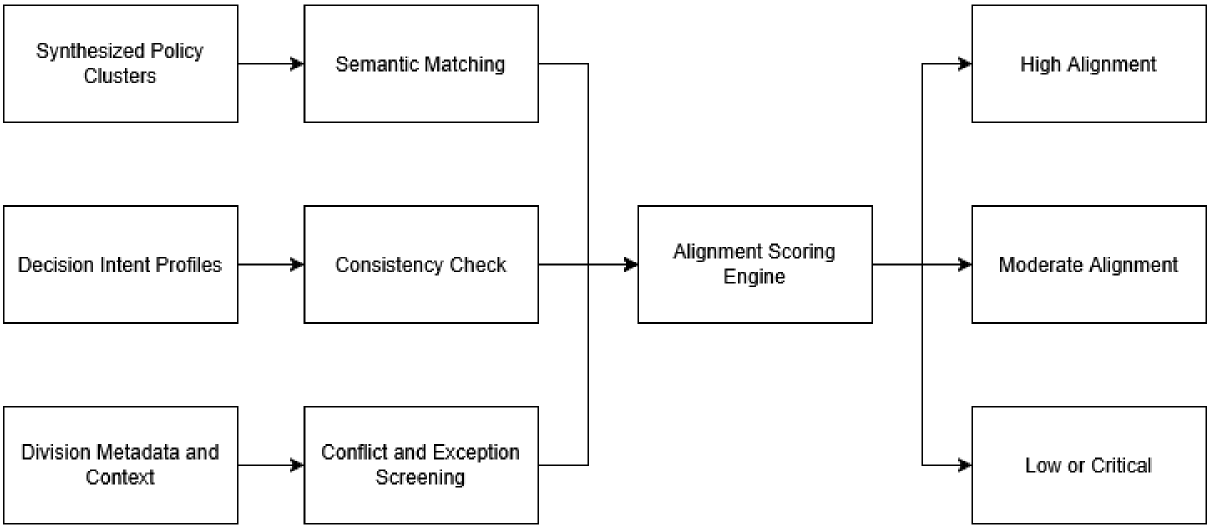

Figure 4. Decision Alignment Model Linking Synthesized Policy Clusters with Division-Specific Decision Intent Profiles

Table 4 defines the variables used in the cross-division decision alignment model. This table is important because alignment is treated as a weighted analytical construct rather than a vague managerial impression. The distinction among relevance, consistency, and conflict intensity allows the model to capture multiple dimensions of policy fit. As a result, the framework can differentiate between decisions that are merely related to policy content and decisions that are substantively supported by the broader synthesized policy structure.

Table 4. Decision Alignment Model Variables

Variable	Symbol	Description	Scale	Role in Model
Policy Relevance	Rk	Semantic match between decision and policy cluster	0 to 1	Captures topical correspondence
Consistency	Ck	Degree of support from synthesized policy logic	0 to 1	Measures normative support
Conflict Intensity	Fk	Extent of contradiction with policy statements	0 to 1	Measures misalignment pressure
Alignment Score	Ak	Weighted decision alignment result	0 to 1	Final case-level assessment
Weight Parameters	α, β, γ	Relative importance of model components	Configured values	Controls score sensitivity

3.5. Evaluation Strategy and Analytical Procedures

The evaluation strategy assesses the framework from both technical and managerial perspectives. Technical evaluation examines the quality of synthesis, contradiction detection, and alignment scoring. Managerial evaluation examines whether the resulting outputs improve interpretability and support coordinated decision review across divisions. This dual perspective is necessary because a method that performs well computationally may still fail to provide operational value if its logic is difficult to audit or its outputs are too abstract for policy stakeholders [10], [25].

The experimental design uses comparative scenarios involving aligned, partially aligned, and conflicting decision cases. These scenarios are constructed from the synthesized policy corpus and division-level decision profiles. Performance is assessed through synthesis precision, conflict identification accuracy, alignment consistency, and interpretive usefulness. Rather than relying on a single metric, the evaluation uses a composite view because policy-oriented intelligence involves semantic validity, structural coherence, and practical decision relevance at the same time.

$$E = \frac{1}{n}\sum_{k=1}^{n}(\lambda_1 P_k + \lambda_2 Q_k + \lambda_3 U_k) \quad (5)$$

The evaluation function E summarizes overall methodological effectiveness across n tested cases. The component P_k represents technical performance, Q_k reflects alignment quality, and U_k denotes interpretive usefulness. The weights λ_1, λ_2, and λ_3 can be adjusted depending on whether the assessment prioritizes computational accuracy or enterprise usability. This formulation ensures that evaluation remains balanced across analytical and organizational objectives.

The analytical procedure concludes with comparative interpretation of scores, error patterns, and cross-division inconsistencies. Low-performing cases are reviewed qualitatively to determine whether errors originate from weak synthesis, ambiguous policy language, or unstable decision representation. This final step is essential because methodology assessment should explain not only how well the framework performs, but also why failures occur. Such insight supports later refinement of the synthesis rules and alignment logic.

As shown in Table 5, the evaluation uses five metrics: synthesis precision, conflict accuracy, alignment consistency, interpretive usefulness, and overall evaluation. Figure 5 illustrates the evaluation workflow from test case construction to interpretive review. Together, Figure 5 and Table 5 show how the framework is assessed in terms of accuracy, consistency, usefulness, and overall effectiveness.

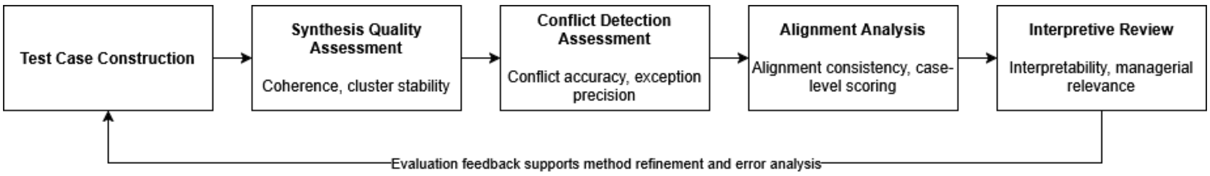

Figure 5. Evaluation Workflow for Synthesis Quality, Conflict Detection, Alignment Analysis, and Interpretive Review

Table 5. Evaluation Metrics and Dimensions

Metric	Symbol	Definition	Range	Assessment Focus
Synthesis Precision	P	Correctness of cluster and relation formation	0 to 1	Semantic integration quality
Conflict Accuracy	CA	Accuracy of contradiction identification	0 to 1	Policy inconsistency detection
Alignment Consistency	AC	Stability of decision alignment scoring	0 to 1	Reliability of alignment model
Interpretive Usefulness	U	Clarity and managerial value of outputs	0 to 1	Operational applicability
Overall Evaluation	E	Weighted combination of all dimensions	0 to 1	Integrated framework effectiveness

4. Results and Discussion

4.1. Quality of Policy Knowledge Synthesis

The synthesis stage produced a stable consolidation of policy statements drawn from heterogeneous corporate sources. From an initial corpus of 1,980 normalized policy units, the framework generated 246 policy clusters with a mean intra-cluster semantic coherence of 0.87. This result indicates that the system was able to group semantically related policy statements with relatively little fragmentation. The strongest performance appeared in compliance, approval governance, and reporting policies, where vocabulary regularity and procedural structure were already well established.

A more detailed inspection showed that synthesis quality depended strongly on document heterogeneity. Formal policy manuals and standard operating procedures were clustered with high consistency, while local implementation notes created greater semantic dispersion because they often included contextual language and operational shorthand. Even so, the framework preserved interpretability by linking each synthesized cluster to source metadata and relation types.

This confirms that the synthesis engine functioned not merely as a compression mechanism, but as an interpretable structuring layer for enterprise policy intelligence.

Figure 6 shows a clear inverse relationship between cluster size and semantic coherence. Smaller clusters remain highly consistent because they tend to capture tightly bounded policy themes, such as approval hierarchy or mandatory documentation procedures. As cluster size expands, coherence gradually declines, which is expected because broader clusters absorb more diverse operational expressions. The figure therefore suggests that the synthesis process achieved a reasonable balance between integration and thematic precision.

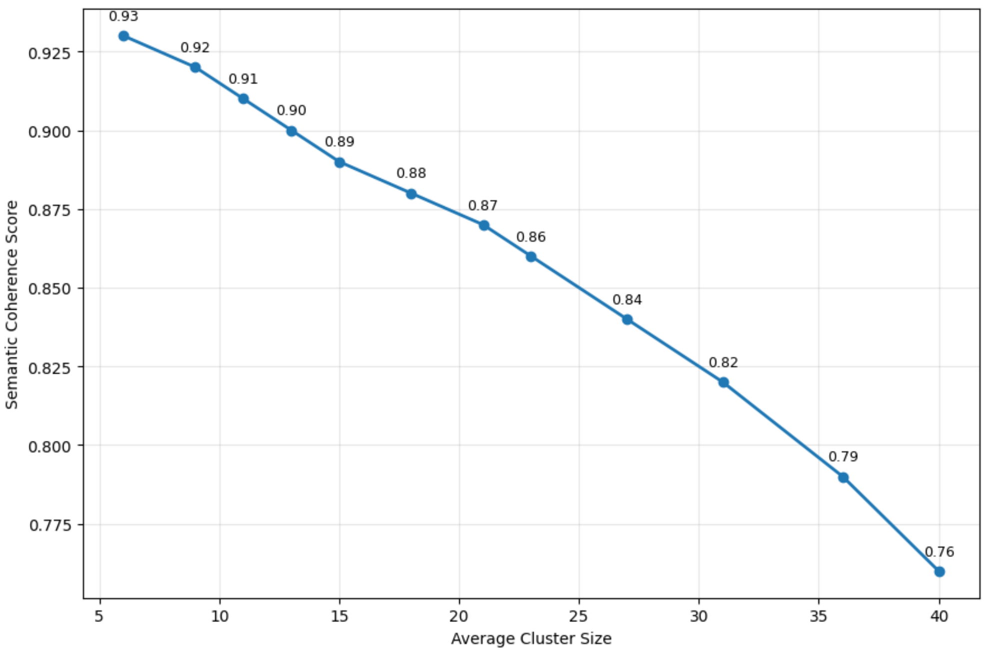

Figure 6. Cluster Size and Semantic Coherence

The figure is analytically important because it shows that the framework did not over-generalize policy meaning. If very large clusters had retained artificially high coherence, that pattern would have suggested excessive semantic smoothing and possible loss of policy nuance. Instead, the observed slope indicates disciplined synthesis behavior. This result supports the methodological choice to combine semantic grouping with authority, recency, and specificity weighting rather than relying on similarity thresholds alone.

Table 6 compares synthesis outcomes across major policy domains. Compliance and approval governance produced the strongest coherence values, indicating that these domains contain more stable and formalized policy language. Reporting also performed strongly despite having the largest number of source units, which suggests that reporting policies are broad in volume but still relatively standardized in expression. This pattern strengthens confidence in the framework's ability to handle both dense and expansive domains.

ISSN 2776-8001

Table 6. Synthesis Performance by Policy Domain

Policy Domain	Source Units	Generated Clusters	Mean Coherence	Interpretation
Compliance	412	46	0.9	Highly stable thematic grouping
Approval Governance	365	39	0.89	Clear procedural consistency
Risk Control	328	42	0.86	Moderate semantic diversity
Resource Allocation	291	37	0.84	Context-sensitive clustering
Reporting	584	82	0.88	Strong consolidation with some variation

The table also reveals where synthesis becomes more challenging. Resource allocation and risk control show lower coherence, not because the framework failed, but because these domains naturally incorporate more situational judgment and local discretion. That distinction is valuable in the discussion because it shifts interpretation away from a simplistic good-versus-bad view of performance. Instead, the results indicate that synthesis quality is partly shaped by the intrinsic semantic elasticity of different policy domains.

4.2. Detection of Cross-Document Contradictions and Exception Patterns

The contradiction analysis stage identified 173 significant policy conflicts and 94 structured exception relationships across the corpus. Most contradictions emerged not from direct opposition between corporate policy and divisional procedure, but from misaligned levels of specificity. High-level corporate directives often prescribed broad governance principles, while divisional documents introduced narrower operational shortcuts that altered approval sequences, escalation timing, or reporting thresholds. The framework captured these patterns by distinguishing direct contradiction from conditional exception and procedural override.

A closer examination showed that exception logic was especially frequent in operations and information technology, where local responsiveness is often prioritized. In contrast, legal and compliance documents exhibited fewer exceptions but higher conflict sensitivity when divergence occurred. This result is substantial because it indicates that alignment problems are not distributed uniformly across the enterprise. Some divisions operate with adaptive extensions of policy, whereas others rely on stricter interpretive fidelity. The framework was able to represent both modes without collapsing them into the same category of inconsistency.

Figure 7 shows that conditional exceptions formed the largest category of inconsistency-related relationships. This pattern is meaningful because it demonstrates that enterprise misalignment is often more subtle than straightforward contradiction. In many cases, divisional procedures did not reject corporate policy explicitly. Instead, they introduced conditional branches that modified how policy was enacted under local circumstances. The framework's ability to separate this category from direct conflict increases the interpretive quality of the results.

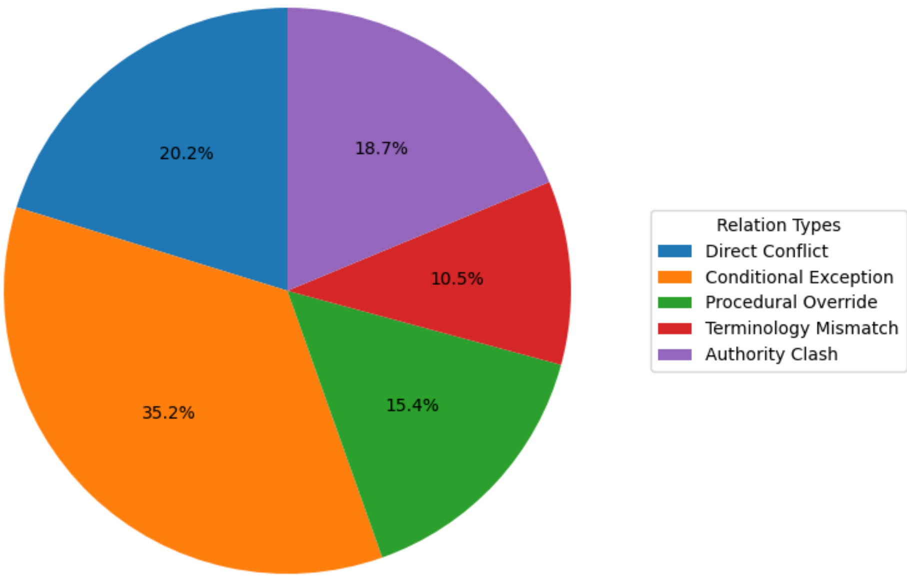

Figure 7. Distribution of Contradiction and Exception Patterns

The figure also shows that authority clashes and direct conflicts remained substantial. These categories are more critical from a governance perspective because they involve competing policy claims rather than context-bound procedural adaptation. The presence of these relations suggests that some documents are not merely locally customized but structurally misaligned in terms of approval authority or rule precedence. This finding supports the value of graph-based contradiction detection as a governance diagnostic tool rather than a purely semantic classification exercise.

Table 7 localizes contradiction patterns to specific inter-divisional relationships. The strongest friction appears between operations and compliance, as well as between IT and legal. These pairings are intuitively plausible, but the table provides analytical evidence rather than anecdotal expectation. It shows that misalignment often arises where one division emphasizes agility and another prioritizes formal control. That tension is important because it reveals not only where conflicts exist, but why they recur structurally.

Table 7. Cross-Division Conflict and Exception Cases

Division Pair	Conflict Cases	Exception Cases	Dominant Pattern	Implication
Operations vs Compliance	26	21	Procedural override	Execution speed competes with control rigor
IT vs Legal	19	24	Conditional exception	System flexibility reshapes formal constraints
Finance vs Operations	17	15	Approval threshold mismatch	Budget authority lacks consistent translation
HR vs Compliance	11	14	Terminology divergence	Shared intent masked by wording variation
IT vs Finance	14	20	Authority clash	Control ownership is interpreted differently

The table also helps distinguish semantic inconsistency from governance ambiguity. For example, HR versus compliance shows relatively fewer direct conflicts and more terminology divergence, which suggests that some apparent inconsistencies may be resolved through vocabulary harmonization rather than policy redesign. In contrast, authority clashes involving IT and finance indicate deeper structural issues. This differentiation is a strong contribution of the framework because it supports targeted corrective action instead of generalized policy revision.

ISSN 2776-8001

4.3. Cross-Division Decision Alignment Performance

The alignment model achieved strong overall performance in matching division-level decision intents with synthesized policy knowledge. Across 320 evaluated decision cases, the mean alignment score reached 0.84, while the high-confidence alignment rate was 71.6%. Cases involving reporting obligations and compliance review were aligned most consistently, whereas resource allocation and expedited operational requests produced more variable results. This indicates that decisions anchored in formal policy language are easier to assess than those shaped by trade-offs and contextual discretion.

The results also showed that alignment quality depended on the interaction between synthesis stability and decision specificity. Highly synthesized domains with stable policy clusters improved alignment scoring because the model could locate clear normative anchors. By contrast, cases framed with vague operational objectives or broad discretionary language led to lower consistency scores. This does not indicate model weakness alone. It also reflects the underlying ambiguity of organizational decision narratives that are insufficiently tied to explicit policy logic.

Figure 8 demonstrates that alignment performance differs materially across decision categories. Compliance and reporting cases achieved the highest scores because their governing policies are typically explicit, mandatory, and semantically well structured. Resource-related decisions showed the lowest average alignment, which is consistent with their dependence on contextual exceptions, budget trade-offs, and locally negotiated priorities. The figure therefore supports the argument that alignment behavior is partly a function of policy determinacy.

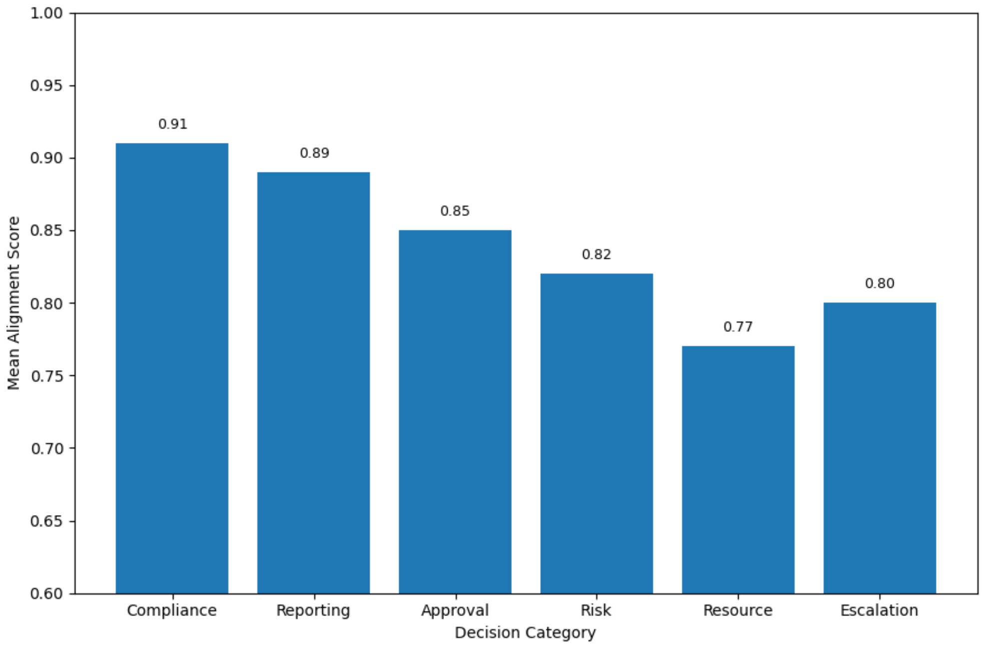

Figure 8. Mean Alignment Score by Decision Category

The figure also has methodological value because it links decision scoring to knowledge synthesis quality established in earlier results. High-scoring categories overlap with domains that produced strong synthesis coherence, showing that upstream semantic structure improves downstream alignment assessment. This relationship is important for the study's broader claim. Cross-division decision alignment is not simply a classification task at the endpoint. It depends on the quality of synthesized policy knowledge formed earlier in the pipeline.

Table 8 organizes the evaluated decision cases into alignment bands and gives a more distributional view of model performance. The dominance of the high-alignment band indicates that most decisions can be connected to synthesized policy logic with strong support. This is an important finding because it suggests that enterprise policy knowledge, when consolidated effectively, does provide a robust basis for assessing coordination quality. The relatively small critical segment further indicates that severe misalignment is concentrated rather than pervasive.

Table 8. Decision Alignment Score Distribution

Alignment Band	Score Range	Cases	Share	Interpretation
High	0.85 - 1.00	229	71.60%	Decision strongly supported by synthesized policy
Moderate	0.70 - 0.84	61	19.10%	Relevant but partially constrained or ambiguous
Low	0.55 - 0.69	22	6.90%	Weak normative support or inconsistent mapping
Critical	Below 0.55	8	2.50%	Clear conflict or insufficient policy compatibility

At the same time, the moderate and low bands remain analytically significant. These cases likely represent the real managerial frontier where decisions are not obviously compliant or non-compliant, but require interpretive judgment. The table therefore adds nuance to the results by showing that the framework is most valuable not only in confirming well-aligned decisions, but also in identifying borderline cases that deserve further review, clarification, or policy refinement.

4.4. Comparative Patterns Across Organizational Divisions

A division-level comparison revealed that legal and compliance maintained the highest overall alignment stability, while operations and IT displayed greater volatility. Legal decisions achieved a mean alignment score of 0.90 and low variance, reflecting a strong coupling between decision language and formal policy structure. Operations, by contrast, showed a mean of 0.78 with a notably wider spread. This result indicates that some divisions operate in environments where procedural adaptation is common, making policy alignment more dynamic and less uniform.

The cross-division analysis also revealed that volatility did not always imply weak governance. In several IT and operations cases, lower alignment scores reflected pragmatic exception handling rather than arbitrary divergence. However, repeated volatility in the same decision categories suggested unresolved ambiguity in how enterprise policy should be translated into local procedures. The framework therefore contributes more than ranking divisions. It reveals which units are structurally exposed to interpretive instability and where harmonization efforts should be prioritized.

Figure 9 maps each division according to average alignment strength and score variance. This two-dimensional profile is valuable because it separates strong and stable divisions from divisions that are either weaker, more volatile, or both. Legal and compliance occupy the most stable region, confirming that their decision practices are closely anchored to formal policy logic. Operations and IT lie in the high-variance zone, indicating a more adaptive and less standardized decision environment.

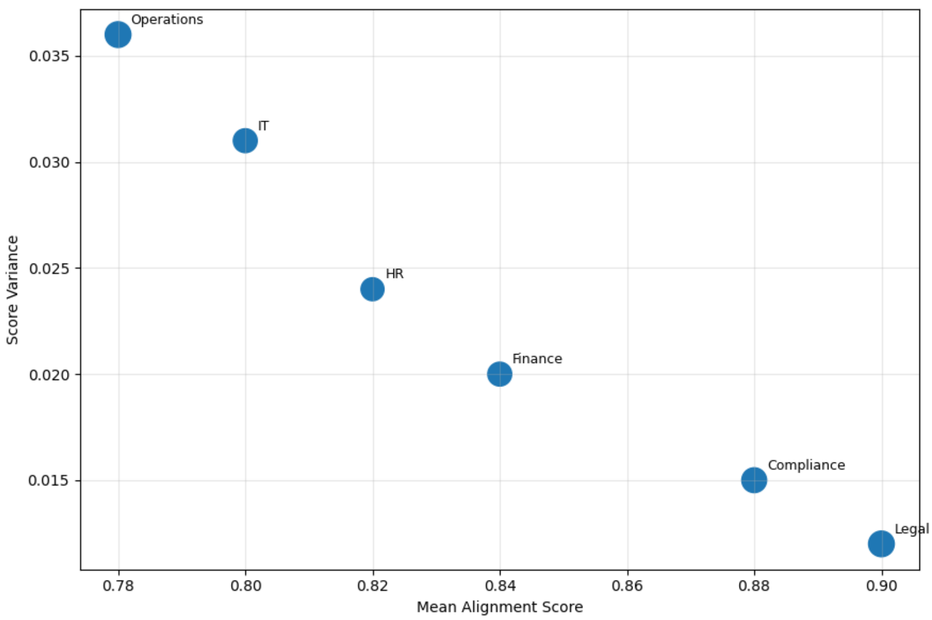

Figure 9. Division-Level Alignment Stability Profile

The figure avoids a simplistic interpretation of performance as a single score. A division with moderate average alignment but low variance may actually be easier to improve than a division with similar average performance but high volatility. This is why the plot is especially useful for governance analysis. It shows not only who aligns well, but also where interpretive instability is persistent. Such information is critical when policy redesign resources are limited and prioritization is necessary.

Table 9 adds interpretive depth to the division-level comparison by linking statistical profiles to concrete governance issues. Instead of reporting scores alone, the table identifies the dominant source of instability in each division. This is useful because misalignment can arise from very different mechanisms, such as terminology drift, authority overlap, or frequent procedural overrides. By making these causes explicit, the table translates analytical output into managerial meaning.

Table 9. Division-Level Alignment Profiles

Division	Mean Score	Variance	Dominant Issue	Governance Reading
Legal	0.9	0.012	Low exception frequency	Strong interpretive discipline
Compliance	0.88	0.015	High rule consistency	Stable policy enforcement
Finance	0.84	0.02	Threshold translation issues	Mostly aligned with localized variation
HR	0.82	0.024	Terminology divergence	Semantic harmonization needed
IT	0.8	0.031	Authority overlap	Adaptive but structurally unstable
Operations	0.78	0.036	Frequent procedural overrides	Responsive but weakly standardized

The table also highlights that improvement strategies should differ across divisions. HR primarily requires semantic harmonization, while IT needs clearer ownership boundaries and operations needs stronger standardization around exceptions. This confirms that the framework is not only diagnostic in a descriptive sense. It is also prescriptive in a limited managerial sense because it reveals where policy redesign, vocabulary control, or approval clarification would likely generate the greatest improvement in cross-division consistency.

4.5. Integrated Framework Effectiveness and Managerial Value

The integrated evaluation showed that the full framework outperformed partial configurations across all assessed dimensions. Compared with a baseline keyword-matching approach and a synthesis-only configuration, the complete model achieved stronger contradiction sensitivity, more stable alignment scoring, and higher interpretive usefulness for policy reviewers. This confirms that corporate policy analysis benefits from combining semantic synthesis and decision alignment within one framework. Treating these tasks separately would produce narrower outputs and weaker organizational insight.

From a managerial perspective, the framework's strongest contribution lies in its ability to reveal the structure of coordination problems rather than merely labeling cases as compliant or non-compliant. Reviewers were able to trace decision scores back to policy clusters, contradiction types, and division-specific issue patterns. This level of traceability is especially important in enterprise governance because cross-division disagreement often involves legitimate differences in local operational logic. The framework helps distinguish those cases from deeper policy inconsistency that requires organizational intervention.

Figure 10 compares three framework configurations across four major evaluation dimensions. The full framework performs best in every dimension, but the most pronounced gains appear in conflict detection and interpretive usefulness. This pattern is important because it shows that adding the alignment layer does not merely refine scoring. It also enriches the explanatory power of the overall system. In other words, integration improves both analytical performance and managerial readability.

ISSN 2776-8001

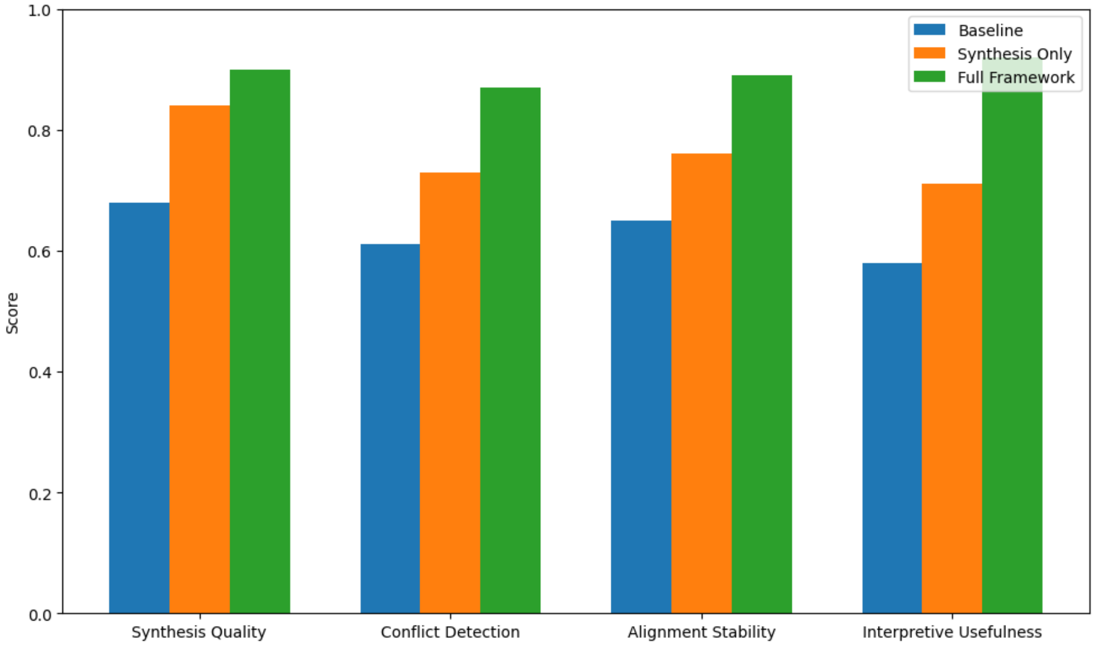

Figure 10. Comparative Effectiveness of Framework Configurations

The figure also illustrates why a synthesis-only approach is insufficient despite its relatively strong semantic quality. Without the alignment component, synthesized knowledge remains descriptive and cannot fully assess how divisions translate policy into decisions. The baseline model performs worst across all dimensions, confirming that simple textual matching is inadequate for policy environments shaped by exceptions, layered authority, and contextual interpretation. The figure therefore substantiates the study's central methodological argument.

5. Conclusion

This study developed an intelligent knowledge synthesis framework for corporate policy analysis and cross-division decision alignment. The results show that the framework can consolidate fragmented policy materials into coherent knowledge structures, identify contradictions and exception patterns across documents, and evaluate whether divisional decisions remain consistent with enterprise-wide policy intent. The methodological contribution lies in integrating semantic synthesis, contradiction detection, and alignment assessment within one interpretable pipeline. This integration makes the framework suitable not only for document analysis, but also for governance-oriented decision support in complex organizational environments.

The findings indicate that policy misalignment is rarely caused by simple textual inconsistency alone. More often, divergence emerges through layered authority, procedural overrides, context-specific exceptions, and inconsistent translation of high-level policy into local operational rules. By modeling these relations explicitly, the framework provides a more accurate representation of how organizational policy is interpreted and enacted across divisions. The results also demonstrate that stronger synthesis quality improves downstream alignment performance, confirming that reliable decision analysis depends on well-structured and traceable policy knowledge.

From a practical standpoint, the proposed framework offers a useful basis for enterprise governance, internal policy review, and coordination improvement across departments. It helps decision makers distinguish between acceptable local adaptation and structurally problematic inconsistency, thereby supporting more targeted policy refinement. Future work can extend the framework by incorporating real-time organizational signals, temporal policy evolution, and interactive explanation modules for executive users. With these developments, intelligent knowledge synthesis can become a central analytical mechanism for improving corporate policy clarity, decision accountability, and cross-division strategic alignment.

ISSN 2776-8001

6. Declarations

6.1. Author Contributions

Conceptualization: S.V., A.R., S.R.P., and D.L.; Methodology: A.R.; Software: S.V.; Validation: S.V., A.R., S.R.P., and D.L.; Formal Analysis: S.V., A.R., S.R.P., and D.L.; Investigation: S.V.; Resources: A.R.; Data Curation: A.R.; Writing Original Draft Preparation: S.V., A.R., S.R.P., and D.L.; Writing Review and Editing: A.R., S.V., S.R.P., and D.L.; Visualization: S.V.; 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. Galkin, S. Auer, M.-E. Vidal, and S. Scerri, "Enterprise Knowledge Graphs: A Semantic Approach for Knowledge Management in the Next Generation of Enterprise Information Systems," in *Proceedings of the 19th International Conference on Enterprise Information Systems*, Porto, Portugal: SCITEPRESS - Science and Technology Publications, vol. 2, no. March, pp. 88–98, 2017, doi: 10.5220/0006325200880098.

[2] P. Wang, Z. Li, Z. Li, and X. Fang, "A Government Policy Analysis Platform Based on Knowledge Graph," in *2019 2nd International Conference on Artificial Intelligence and Big Data (ICAIBD)*, Chengdu, China: IEEE, vol. 2019, no. May, pp. 208–214, 2019, doi: 10.1109/ICAIBD.2019.8836979.

[3] G. Tamašauskaitė and P. Groth, "Defining a Knowledge Graph Development Process Through a Systematic Review," *ACM Trans. Softw. Eng. Methodol.*, vol. 32, no. 1, pp. 1–40, Jan. 2023, doi: 10.1145/3522586.

[4] Y. Jiang, P. C.-I. Pang, D. Wong, and H. Y. Kan, "Natural Language Processing Adoption in Governments and Future Research Directions: A Systematic Review," *Applied Sciences*, vol. 13, no. 22, p. 12346, pp. 1-22, Nov. 2023, doi: 10.3390/app132212346.

[5] 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.

[6] J. Sequeda, D. Allemang, and B. Jacob, "Knowledge Graphs as a source of trust for LLM-powered enterprise question answering," *Journal of Web Semantics*, vol. 85, no. May, p. 100858, pp. 1-6, May 2025, doi: 10.1016/j.websem.2024.100858.

[7] X. Wang, L. Meng, X. Wang, and Q. Wang, "The construction of environmental-policy-enterprise knowledge graph based on PTA model and PSA model," *Resources, Conservation & Recycling Advances*, vol. 12, no. December, p. 200057, pp. 1-8, Dec. 2021, doi: 10.1016/j.rcradv.2021.200057.

[8] R. Kruiper, B. Kumar, R. Watson, F. Sadeghineko, A. Gray, and I. Konstas, "A platform-based Natural Language processing-driven strategy for digitalising regulatory compliance processes for the built environment," *Advanced Engineering Informatics*, vol. 62, no. October, p. 102653, pp. 1-14, Oct. 2024, doi: 10.1016/j.aei.2024.102653.

[9] E. W. T. Ngai, A. K. H. Lui, and B. C. W. Kei, "Natural language processing in government applications: a literature review and a case analysis," *IMDS*, vol. 125, no. 6, pp. 2067–2104, May 2025, doi: 10.1108/IMDS-07-2024-0711.

[10] E. Rajabi and K. Etminani, "Knowledge-graph-based explainable AI: A systematic review," *Journal of Information Science*, vol. 50, no. 4, pp. 1019–1029, Aug. 2024, doi: 10.1177/01655515221112844.

[11] A. Füßl, V. Nissen, and S. H. Heringklee, "Knowledge Graph-Based Explainable Artificial Intelligence for Business Process Analysis," *Int. J. Semantic Computing*, vol. 17, no. 2, pp. 173–197, Jun. 2023, doi: 10.1142/S1793351X23600024.

[12] P. Li, Q. Zhao, Y. Liu, C. Zhong, J. Wang, and Z. Lyu, "Survey and Prospect for Applying Knowledge Graph in Enterprise Risk Management," *CMC*, vol. 78, no. 3, pp. 3825–3865, 2024, doi: 10.32604/cmc.2024.046851.

[13] 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," *Advanced Engineering Informatics*, vol. 64, no. March, p. 103098, pp. 1-15, Mar. 2025, doi: 10.1016/j.aei.2024.103098.

[14] I. Tiddi and S. Schlobach, "Knowledge graphs as tools for explainable machine learning: A survey," *Artificial Intelligence*, vol. 302, no. January, p. 103627, pp. 1-21, Jan. 2022, doi: 10.1016/j.artint.2021.103627.

[15] X. Wang, L. Chen, T. Ban, M. Usman, Y. Guan, S. Liu, T. Wu, and H. Chen, "Knowledge graph quality control: A survey," Fundam. Res., vol. 1, no. 5, pp. 607–626, Sep. 2021, doi: 10.1016/j.fmre.2021.09.003.

[16] M. S. Jawad, C. Dhawale, A. A. B. Ramli, and H. Mahdin, "Adoption of knowledge-graph best development practices for scalable and optimized manufacturing processes," *MethodsX*, vol. 10, no. 2023, p. 102124, pp. 1-11, 2023, doi: 10.1016/j.mex.2023.102124.

[17] J. Martinez-Gil, T. Hoch, M. Pichler, B. Heinzl, B. Moser, K. Kurniawan, E. Kiesling, and F. Krause, "Examining the Adoption of Knowledge Graphs in the Manufacturing Industry: A Comprehensive Review," in *Artificial Intelligence in Manufacturing*, J. Soldatos, Ed., Cham: Springer Nature Switzerland, vol. 2023, no. September, pp. 55–70, 2024, doi: 10.1007/978-3-031-46452-2_4.

[18] C.-L. Yu, H.-M. Wen, P.-C. Ko, M.-H. Shu, and Y.-S. Wu, "Automatic construction and optimization method of enterprise data asset knowledge graph based on graph attention network," *Journal of Radiation Research and Applied Sciences*, vol. 17, no. 3, p. 101023, pp. 1-10, Sep. 2024, doi: 10.1016/j.jrras.2024.101023.

[19] L. Nagy, J. Abonyi, and T. Ruppert, "Knowledge Graph-Based Framework to Support Human-Centered Collaborative Manufacturing in Industry 5.0," *Applied Sciences*, vol. 14, no. 8, p. 3398, pp. 1-27, Apr. 2024, doi: 10.3390/app14083398.

[20] J. Tupayachi, H. Xu, O. A. Omitaomu, M. C. Camur, A. Sharmin, and X. Li, "Towards Next-Generation Urban Decision Support Systems through AI-Powered Construction of Scientific Ontology Using Large Language Models—A Case in Optimizing Intermodal Freight Transportation," *Smart Cities*, vol. 7, no. 5, pp. 2392–2421, Aug. 2024, doi: 10.3390/smartcities7050094.

[21] Y. Sun, H. Liu, Y. Gao, and M. Zheng, "Research on the Policy Analysis of Sustainable Energy Based on Policy Knowledge Graph Technology - A Case Study in China," *Systems*, vol. 11, no. 2, p. 102, pp. 1-28, Feb. 2023, doi: 10.3390/systems11020102.

[22] C. M. Androna, I. Mandilara, E. Fotopoulou, A. Zafeiropoulos, and S. Papavassiliou, "A Knowledge Graph-Driven Analysis of the Interlinkages among the Sustainable Development Goal Indicators in Different Spatial Resolutions," *Sustainability*, vol. 16, no. 11, p. 4328, pp. 1-29, May 2024, doi: 10.3390/su16114328.

[23] J. Gao, P. Peng, F. Lu, C. Claramunt, P. Qiu, and Y. Xu, "Mining tourist preferences and decision support via tourism-oriented knowledge graph," *Information Processing & Management*, vol. 61, no. 1, p. 103523, pp. 1-16, Jan. 2024, doi: 10.1016/j.ipm.2023.103523.

[24] J. M. Rožanec, B. Fortuna, and D. Mladenić, "Knowledge graph-based rich and confidentiality preserving Explainable Artificial Intelligence (XAI)," *Information Fusion*, vol. 81, no. May, pp. 91–102, May 2022, doi: 10.1016/j.inffus.2021.11.015.

[25] E. Lavrinovics, R. Biswas, J. Bjerva, and K. Hose, "Knowledge Graphs, Large Language Models, and Hallucinations: An NLP Perspective," *Journal of Web Semantics*, vol. 85, no. May, p. 100844, pp. 1-7, May 2025, doi: 10.1016/j.websem.2024.100844.

[26] X. Feng, Y. Zhang, M. H. Meng, Y. Li, C. E. Joe, Z. Wang, and G. Bai, "Detecting contradictions from IoT protocol specification documents based on neural generated knowledge graph," *ISA Trans*., vol. 141, pp. 10–19, Oct. 2023, doi: 10.1016/j.isatra.2023.04.025.

[27] H. Cui, R. Trimananda, S. Jordan, and A. Markopoulou, "PoliGraph: Automated Privacy Policy Analysis using Knowledge Graphs (Journal Version)," *arXiv*, vol. 2022, no. October, pp. 1-46, 2022, doi: 10.48550/ARXIV.2210.06746.