Recommending Actionable Strategies: A Semantic Approach to Integrating Analytical Frameworks with Decision Heuristics

Renato Ghisellini , Remo Pareschi * , Marco Pedroni , Giovanni Battista Raggi

Article

Recommending Actionable Strategies: A Semantic Approach to Integrating Analytical Frameworks with Decision Heuristics

Renato Ghisellini [1], Remo Pareschi [2],*, Marco Pedroni [1] and Giovanni Battista Raggi [1]

[1]   Institute for Generative Strategy
[2]   University of Molise
*   Correspondence: remo.pareschi@unimol.it

Abstract: We present a novel approach for **recommending actionable strategies** by integrating strategic frameworks with decision heuristics through **semantic analysis**. While strategy frameworks provide systematic models for assessment and planning, and decision heuristics encode experiential knowledge, these traditions have historically remained separate. Our methodology bridges this gap using **advanced natural language processing (NLP)**, demonstrated through integrating frameworks like the 6C model with the Thirty-Six Stratagems. The approach employs **vector space representations** and **semantic similarity calculations** to map framework parameters to heuristic patterns, supported by a computational architecture that combines deep semantic processing with constrained use of Large Language Models. By processing both **primary content** and **secondary elements** (diagrams, matrices) as complementary linguistic representations, we demonstrate effectiveness through corporate strategy case studies. The methodology **generalizes** to various analytical frameworks and heuristic sets, culminating in a **plug-and-play architecture** for generating **recommender systems** that enable cohesive integration of strategic frameworks and decision heuristics into actionable guidance.

Keywords: recommender systems; semantic analysis; strategic frameworks; decision heuristics; natural language processing; plug-and-play architecture

1. Introduction

Organizations today rely on two primary but historically separate textual traditions for **strategic management and decision-making**: **analytical frameworks** and **decision heuristics**. Both traditions encode strategic knowledge in natural language—often complemented by diagrams or structured representations—yet differ in scope and style. Analytical frameworks such as Porter's Five Forces, SWOT Analysis, and Value Chain Analysis offer systematic lenses for situational assessment and long-term planning. Meanwhile, decision heuristics—from early military wisdom (e.g., the Thirty-Six Stratagems) to modern "rules of thumb"—provide concise, actionable insights forged through real-world experience.

In practice, **combining these two traditions** holds tangible advantages: a more balanced approach to strategic planning, clearer avenues for *evidence-based recommendations*, and reduced time spent on exhaustive analyses. However, frameworks and heuristics rarely interact in a unified process. Frameworks excel at comprehensiveness and rigor but risk *analysis paralysis*, whereas heuristics are more agile but *can oversimplify* complex scenarios. **Bridging this gap** would enable decision-makers to reap the complementary strengths of each method, yielding recommended strategies that are both thorough and swiftly implementable.

Recent advances in **artificial intelligence (AI) and natural language processing (NLP)** offer a powerful way to integrate these traditions. By applying **semantic analysis** to uncover linguistic patterns, rhetorical structures, and conceptual interdependencies within strategic texts, we can construct automated mappings between the methodical constructs of analytical frameworks and the concise

action steps of heuristics. This paper proposes a **recommender-system architecture** that leverages these mappings to generate **actionable strategic recommendations**, ultimately expediting decision-making and enhancing strategic insight.

To illustrate this **semantic integration** concretely, we focus on two representative models: the 6C framework and the Thirty-Six Stratagems. The 6C framework synthesizes recurring strategic themes (offensive/defensive strength, relational capacity, potential energy, temporal availability, and context-fit), drawn from military and business literature. In contrast, the Thirty-Six Stratagems—rooted in Chinese political, military, and civil discourse—encapsulate centuries of heuristic insight in pithy expressions. Through **advanced NLP** tools—such as vector-space embeddings, topic modeling, and pattern recognition—we demonstrate how linguistic cues in each stratagem correlate with specific 6C parameters. This systematic analysis then drives an automated pipeline that matches any given strategic situation with suitable heuristics, producing **evidence-based, context-aware recommendations**.

Two key innovations underscore our approach. First, we embed the system in an **interactive simulation environment**, prompting decision-makers to express scenarios in natural language. The environment analyzes these textual inputs, computes relevance scores, and returns **recommendations** for how best to combine or select heuristics in light of the chosen strategic framework. Second, we employ **Large Language Models (LLMs)** in a controlled manner to produce coherent, narrative-style reports that clarify the rationale behind each recommendation. By integrating LLMs as interpreters *rather* than autonomous decision-makers, we preserve analytical rigor while providing accessible explanations.

In what follows, we detail our **semantic methodology**, discuss the **computational architecture** that enables framework-heuristic integration, and illustrate its real-world relevance through **case studies** in corporate strategy. We then show how this **plug-and-play architecture** generalizes beyond 6C and the Thirty-Six Stratagems, adapting to other widely known frameworks such as Porter's Five Forces and SWOT. Ultimately, we aim to demonstrate how **organizations can deploy a recommender-system approach** to merge comprehensive strategic analysis with proven heuristic insights, delivering **actionable guidance** that is both robust and readily applicable in complex environments.

The remainder of this paper is structured as follows:

- Section 2 provides background knowledge
- Section 3 details our language analysis methodology for framework integration
- Section 4 presents the computational architecture supporting this integration
- Section 5 demonstrates the approach through two case studies
- Section 6 provides empirical validation
- Section 7 discusses related work
- Section 8 discusses implications and future directions

2. Background

Integrating analytical frameworks with decision heuristics through semantic analysis represents a highly interdisciplinary endeavor, drawing from multiple domains, including strategic management, heuristics, computer science, and linguistics. This confluence of fields necessitates thoroughly examining key concepts and prior work across several dimensions. Specifically, we must understand (1) how strategic frameworks systematize decision parameters, as exemplified by the 6C model; (2) how decision heuristics encapsulate experiential knowledge, illustrated through the Thirty-Six Stratagems; (3) how semantic analysis enables framework-heuristic integration; (4) how mathematical formulations like Kullback-Leibler divergence support validation; and (5) how gamification principles and Large Language Models facilitate practical implementation. The following subsections provide this essential foundation, which underlies our novel recommender-system architecture for strategic decision support.

2.1. The 6C Framework as an Analytical Classification Tool

The 6C framework was conceived to provide a set of clear, well-defined parameters that would enable straightforward experimentation with our semantic integration approach. While established frameworks like SWOT Analysis and Porter's Five Forces offer comprehensive analytical tools, the 6C parameters were specifically designed to facilitate initial testing and validation of our methodology, with the understanding that the same principles could then be transferred to these well-known and widely adopted frameworks. The parameters were distilled from an extensive study of strategic literature, offering a simplified yet robust foundation for our initial framework-heuristic integration experiments.

The key parameters of the 6C model are as follows:

1. p_1 = **Offensive Strength:** The ability to proactively shape and influence the strategic landscape.
2. p_2 = **Defensive Strength:** The resilience to respond effectively to adversarial actions or challenges.
3. p_3 = **Relational Capacity:** The ability to manage and leverage relationships with external stakeholders.
4. p_4 = **Potential Energy:** The availability and strategic deployment of resources.
5. p_5 = **Temporal Availability:** The strategic use of time and timing in decision-making.
6. p_6 = **Contextual Fit:** The degree to which decisions align with the strategic context, ensuring they are well-informed and relevant.

These classification processes emerged from a comparative study of prominent military strategists—ranging from ancient figures such as Sun Tzu and Chanakya to modern thinkers like Machiavelli and Clausewitz and more contemporary theorists such as Beaufre and Liddell Hart. Over time, they have transcended their military origins to influence corporate strategy, echoing broader perspectives that integrate historical wisdom with rigorous managerial concepts [1].

The 6C framework functions as a classification system for organizing and analyzing data obtained through competitive intelligence and data analytics, a capability especially valuable in digitally transforming industries [2]. This classification provides a structured approach to understanding competitive landscapes, where the six parameters can be assessed quantitatively or qualitatively. While we do not delve into the specific mechanics of setting each parameter—this is handled through the gamified environment described in Section 2.6—the focus of this paper is on *how* these parameters (once established) can be integrated with heuristic decision patterns via semantic analysis.

2.2. The Thirty-Six Stratagems: Crystallized Decision Patterns

The Thirty-Six Stratagems comprise a collection of ancient Chinese military decision heuristics that have evolved into widely applicable strategic principles [3,4]. These stratagems are especially valuable for our study since they embody concise, experience-based rules, akin to "simple rules" in strategy-making [5], but encoded in metaphorical language with deep historical roots. The highly metaphorical nature of these stratagems, while encoding deep strategic wisdom, presents unique challenges for semantic analysis. Their idiosyncratic expressions require careful interpretation to bridge ancient military metaphors with modern strategic concepts. This linguistic complexity makes them an especially rigorous test case for our framework integration methodology—success here would suggest strong generalizability to more straightforward decision-making frameworks.

Traditionally, the thirty-six stratagems are divided into six categories:

1.   Stratagems for winning advantageous positions
2.   Stratagems for confrontation
3.   Stratagems for attack
4.   Stratagems for creating confusion
5.   Stratagems for gaining ground
6.   Stratagems for desperate situations

Each stratagem defines a concrete pattern of strategic action—often encapsulated in pithy maxims such as "Besiege Wei to rescue Zhao" (围魏救赵), which illustrates an indirect approach, or "Kill with a

borrowed knife" (借刀杀人), which suggests leveraging external resources. By codifying these patterns, the Thirty-Six Stratagems lend themselves to computational extraction and matching when we embed them in modern NLP pipelines.

2.3. Additional Strategic Frameworks

While the 6C framework and the Thirty-Six Stratagems serve as primary examples in our study, our approach generalizes to multiple analytical frameworks. Following the systematic nature of the 6C framework, other prominent analytical frameworks provide structured parameters for strategic assessment:

SWOT Analysis

SWOT Analysis, developed by Albert Humphrey at Stanford Research Institute, offers a systematic approach to evaluating internal and external factors [6]. Its clear parameter structure makes it particularly suitable for semantic analysis:

- **Strengths:** Internal capabilities and resources that provide competitive advantages
- **Weaknesses:** Internal limitations that may hinder strategic objectives
- **Opportunities:** External factors or trends that could benefit the organization
- **Threats:** External challenges that could negatively impact performance

Porter's Five Forces

This analytical framework, developed by Michael Porter, systematically dissects industry structure through five well-defined parameters [7]:

- **Competitive Rivalry:** Intensity of competition among existing players
- **Supplier Power:** Bargaining power of suppliers
- **Buyer Power:** Bargaining power of customers
- **Threat of New Entrants:** Ease with which new competitors can enter
- **Threat of Substitution:** Availability of alternative products or services

Like the 6C framework, SWOT and Five Forces provide clear analytical parameters that can be vectorized and processed through semantic analysis for framework-heuristic integration. The effectiveness of our approach with these widely adopted frameworks demonstrates its potential applicability to other strategic analysis tools, both contemporary and classical, a direction we discuss in our future work.

2.4. Semantic Analysis in Strategic Text Processing

Our approach employs several key techniques from Natural Language Processing (NLP) to analyze and connect strategic frameworks (like the 6C model) with decision heuristics (such as the Thirty-Six Stratagems). Recent advances in NLP, especially *Transformer*-based architectures [8], have shown that vector-space language representations can capture nuanced semantic meaning and contextual relationships.

- **Vector Space Representations:** We encode strategic concepts using word embeddings and sentence transformers, building on methods such as *BERT* [9], *Sentence-BERT* [10], that excel in capturing contextual nuances. These techniques enable mathematical operations on semantic meaning and allow us to compare entire passages or phrases in a high-dimensional embedding space.
- **Topic Modeling (Optional)**: We can also apply *Latent Dirichlet Allocation* (LDA) and related approaches [11] to identify high-level themes in strategic texts (e.g., "alliances," "resource optimization"). Although not the core driver in our current implementation, such thematic analysis can support explainability by highlighting relevant topics that link to each framework parameter or stratagem.

- **Semantic Similarity Metrics:** Using cosine similarity or similar measures, we quantify relationships between vectors representing framework parameters and stratagem patterns. In our implementation with the Thirty-Six Stratagems, this objectively measures how well each stratagem aligns with specific 6C parameters (e.g., *Offensive Strength, Relational Capacity*).

The semantic analysis process involves the following steps:

1. Preprocessing (if needed) and Concept Extraction: In scenarios where texts are unstructured, we may apply standard NLP techniques (tokenization, chunking, domain ontology extraction). However, this step becomes straightforward if frameworks like 6C or Porter's Five Forces are already delineated.
2. Creating vector representations of framework parameters and heuristic patterns;
3. Computing similarity matrices to link frameworks (e.g., 6C) with heuristics (e.g., Thirty-Six Stratagems);
4. Identifying significant semantic connections and ranking them for further interpretation.

Our work builds on emerging applications of BERT-like methods to nontraditional NLP contexts [12], where domain-specific texts require fine-grained semantic understanding. Adapting these cutting-edge techniques preserves both the textual richness of each strategic expression and the interpretability needed for real-world strategic decision-making. Having said that, it must be added that although these models have proven effective, they inherit statistical biases from their training corpora, which may affect how parameters and heuristics are aligned. Because we employ LLMs only for explanatory output, the overall recommendation pipeline remains largely heuristic- and framework-driven. Nevertheless, mitigating potential embedding biases—by fine-tuning domain-specific corpora, employing bias detection tools, or adopting interpretability frameworks—is an important direction for future development.

2.5. Kullback–Leibler Divergence

In information theory and statistics, the **Kullback–Leibler (KL) divergence** measures how one probability distribution diverges from a second, *reference* distribution [13]. Given two discrete probability distributions (P) and (Q) over the same outcome space (Ω), the KL divergence is defined as:

$$D_{\mathrm{KL}}(P,|,Q);=;\sum_{i\in\Omega};P(i),\log!\left(\frac{P(i)}{Q(i)}\right). \tag{1}$$

Intuitively, if (P) represents the "true" or expert-labeled distribution (e.g., the relative importance of each parameter in a scenario), while (Q) is the model's approximate distribution, the KL divergence quantifies how inefficient it is to use (Q) in place of (P). A lower

$$(D_{\mathrm{KL}}(P,|,Q))$$

indicates a closer match between the two distributions, while higher values signify greater disparity. Unlike many distance metrics, KL divergence is *not* symmetric (i.e.,

$$(D_{\mathrm{KL}}(P,|,Q) \neq D_{\mathrm{KL}}(Q,|,P)))$$

which means the direction of comparison matters. Because KL divergence also fails to satisfy the triangle inequality, it is not a true metric in the formal sense; however, it remains a widely used 'distance-like' measure for comparing probability distributions.

Why KL Divergence?

We selected KL divergence for three key reasons:

- **Interpretability:** It offers a straightforward interpretation of the "cost" of using an approximate distribution, aligning well with our need to validate semantic analysis against expert annotations.

- **Directionality:** Its asymmetric nature suits our context, where we specifically care about how well our approximate distributions match expert distributions, rather than vice versa.
- **Established Usage:** Its widespread adoption in machine learning and information theory provides a well-tested foundation for measuring distributional differences.

While metrics like the symmetric Jensen–Shannon divergence could be considered, KL divergence provides an efficient and intuitive tool for our semantic analysis workflow, where we compare system-discovered parameter distributions to expert-annotated ones. A lower divergence indicates our system captures expert priorities effectively, while higher values highlight areas needing further calibration.

Usage in This Study

In our context, we leverage KL divergence to compare the *discovered* parameter distributions (e.g., derived from the system's semantic analysis) to *expert-annotated* distributions. This comparison provides a quantitative measure of how closely our system's interpretations align with domain experts' judgments, thereby helping to validate and refine the *robustness* of our semantic mapping process. A lower divergence indicates our system is capturing expert priorities effectively, while higher values highlight areas needing further calibration.

2.6. Gamification of Strategic Decision-Making

To streamline the usage of these analytic results, we have developed a prototype *interactive simulation environment* that enables:

- Exploration of different strategic scenarios;
- Testing of various decision combinations;
- Immediate feedback on potential outcomes;
- Gradual learning from simulated experiences.

Such *gamification* introduces a user-centric interface through which decision-makers set or adjust the situation-specific scores of the chosen analytical framework. For instance, users can enter or modify the 6C parameters—*Offensive Strength*, *Defensive Strength*, and so on—based on real competitive intelligence data or hypothetical "what-if" explorations. The system then applies its semantic mappings (Section 2.1) to generate immediate feedback on how each parameter configuration impacts recommended strategies or outcomes.

Although this gamified environment is fully implemented and was used to produce the case studies described in this paper, its detailed interface design and mechanics lie beyond our current scope. Such a *user-interface* analysis merits its own dedicated treatment and can be assumed here without loss of relevant information about the underlying framework. By highlighting interactive experimentation and rapid feedback, our approach resonates with broader digital transformation trends [2] and encourages deeper engagement and practical experimentation among corporate decision-makers.

3. Language Analysis Methodology

Our methodology for integrating analytical frameworks with decision heuristics centers on semantic analysis of strategic texts. This section details the technical approach to discovering and quantifying relationships between framework parameters and heuristic patterns.

3.1. Vector Space Representation

The first step involves creating vector representations of both framework parameters and heuristic descriptions:

$$v(t) = \sum_{w \in t} \alpha_w \cdot e(w) \tag{2}$$

where $v(t)$ is the vector representation of text t, w represents individual words or phrases, $e(w)$ is the embedding vector for word w, and α_w is the weight assigned to word w. For framework parameters (in our case, the 6Cs), we create vectors from their definitions and associated descriptive text:

$$p_i = v(d_i) + \lambda \sum_j v(c_{ij}) \tag{3}$$

where p_i is the vector for parameter i, d_i is the base definition of parameter i, c_{ij} are associated contextual descriptions, and λ is a weighting factor for contextual information. Detailed calculations and examples of this vector space representation are provided in Appendix A.1.1.

3.2. Semantic Similarity Computation

We compute semantic similarity between framework parameters and heuristics using the cosine similarity measure. Specifically, for a given parameter vector p_i and heuristic vector h_j, the similarity is:

$$\mathrm{sim}(p_i, h_j) \;=\; \frac{p_i \cdot h_j}{\|p_i\| \, \|h_j\|}, \tag{4}$$

where $p_i \cdot h_j$ denotes the dot (scalar) product of the two vectors,

$$p_i \cdot h_j \;=\; \sum_k (p_i)_k \, (h_j)_k,$$

and $\|p_i\|$ (likewise $\|h_j\|$) is the Euclidean norm,

$$\|p_i\| \;=\; \sqrt{\sum_k \left((p_i)_k\right)^2}.$$

A higher cosine similarity value $\mathrm{sim}(p_i, h_j)$ indicates a closer semantic relationship between parameter p_i and heuristic h_j.

Computing $\mathrm{sim}(p_i, h_j)$ for all parameters $i \in \{1, \dots, m\}$ and all heuristics $j \in \{1, \dots, n\}$ produces a similarity matrix S, in which each element s_{ij} is defined by:

$$S \;=\; \begin{bmatrix} s_{11} & s_{12} & \cdots & s_{1n} \\ s_{21} & s_{22} & \cdots & s_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ s_{m1} & s_{m2} & \cdots & s_{mn} \end{bmatrix} \quad \text{with} \quad s_{ij} \;=\; \mathrm{sim}(p_i, h_j). \tag{5}$$

For example, when comparing parameter p_3 (Relational Capacity) with Stratagem 24 ("Use Allies' Resources"), our system computes a similarity score of 0.93, quantitatively capturing their strong semantic alignment. Detailed calculations and further examples of this process appear in Appendix A.2.1.

3.3. Distribution Discovery

For each heuristic h_j, our system generates a *discovered distribution* across the framework parameters (e.g., the 6C parameters). Concretely, we first compute similarity scores s_{ij} between parameter p_i and heuristic h_j (Section 3.2). We then normalize these scores to form a probability-like distribution:

$$d_{ij} = \frac{s_{ij}}{\sum_k s_{kj}}, \tag{6}$$

where d_{ij} represents the weight (or relative importance) of parameter p_i in heuristic h_j. This summation-based (L1) normalization treats each heuristic's parameter weights as if they were proba-

bilities that sum to 1. It thus naturally encodes the idea that a given heuristic distributes its "attention" across the available parameters.

Note on Alternative Normalization. An alternative would be to use the Euclidean (L2) norm, where

$$d_{ij} \;=\; \frac{s_{ij}}{\sqrt{\sum_k (s_{kj})^2}},$$

thereby turning each heuristic's parameter vector into a *unit vector* in ℓ_2 space. In our approach, we opt for L1 normalization to mirror a "probability-like" interpretation—each heuristic can be seen as distributing its "weight" over parameters in a manner analogous to probabilities. L2 normalization could be equally valid in other contexts, especially if one prefers strictly geometric interpretations of distance in the parameter space.

To validate these distributions, we compare them against expert-annotated distributions using the Kullback-Leibler (KL) divergence measure. For Stratagem 24, this comparison yielded a KL divergence of 0.0273, indicating strong alignment between system-discovered and expert-provided distributions. Detailed calculations and additional examples of this validation process are provided in Appendix A.2.2.

By comparing discovered distributions against expert-annotated ones for each heuristic, we obtain a numerical sense of alignment or mismatch. This process can be iterated: a large KL divergence flags a heuristic whose vector representation needs either textual refinements (e.g., additional synonyms or clarifications) or updates to the weighting scheme. Over time, **machine-based distribution discovery** converges with expert insights, yielding robust mappings that faithfully reflect how these strategic heuristics fit into analytical parameters.

As we already pointed out, the *experts* in this step are specialists in the *methodologies* (e.g., the 6C model, the Thirty-Six Stratagems) rather than sector-specific domain experts. This distinction ensures that the high-level semantic structure of each stratagem is validated by those who understand it conceptually, independent of particular industries or case studies.

3.4. Stratagem Selection Algorithm

After the *distribution discovery* phase (§3.3), we obtain an **invariant** distribution of analytical properties for each heuristic. This invariant distribution reflects how strongly each heuristic (e.g., a particular stratagem) aligns with each framework parameter (e.g., the 6C model) based on the text analysis and expert validation.

Situation-Specific Parameter Vector. In contrast, a **variable** (or *current*) situation vector, denoted by x, describes how the analytical parameters apply to the *present scenario*. For instance, if a certain strategic context demands high *Offensive Strength* (p_1) and moderate *Relational Capacity* (p_3), x will capture these intensities accordingly.

Matching Heuristics to the Situation. By comparing x with each heuristic's invariant distribution, we produce a *recommendation score* that indicates how well that heuristic fits the present conditions. This process is summarized in Algorithm 1.

Algorithm 1 Stratagem Selection

Require: Situation vector x, Similarity matrix S, Threshold θ
Ensure: Ranked list of relevant stratagems
  1: $scores \leftarrow \varnothing$
  2: **for** each stratagem j **do**
  3:     $d_j \leftarrow \text{normalize}(S[:,j])$ {Invariant distribution of parameters for heuristic j}
  4:     $score_j \leftarrow \text{similarity}(x, d_j)$ {Compare current situation vector to heuristic distribution}
  5:     **if** $score_j \geq \theta$ **then**
  6:         $scores.\text{append}((j, score_j))$
  7:     **end if**
  8: **end for**
  9: **return** $\text{sort}(scores, \text{descending=True})$

The algorithm produces a ranked list of relevant heuristics, from strongest to weaker matches, filtered by a minimum threshold θ to ensure only sufficiently strong alignments are proposed. When applied to strategic scenarios in our case studies, the algorithm successfully identified relevant stratagems matching the strategic context. For instance, in the hydrogen vs. electric vehicle competition case, it highlighted stratagems focused on indirect positioning and resource leveraging. Detailed examples of the algorithm's application, including specific calculations and case study connections, are provided in Appendix A.3.

3.5. Semantic Validation

To ensure *robustness* and *credibility* in the discovered semantic mappings, we conduct a three-pronged validation:

1. **Cross-Validation:** We compare parameter–heuristic distributions generated by multiple embedding approaches (e.g., different Transformer models, dimensionality settings). If the mappings remain consistent across these variations, it indicates resilience against model-specific biases or hyperparameter choices.

2. **Perturbation Analysis:** We introduce small textual modifications (e.g., synonyms, minor paraphrasing) to heuristic descriptions or framework definitions and observe whether the resulting distributions change drastically. A stable mapping under such perturbations implies that the system captures deeper semantic relationships rather than overfitting to exact word forms.

3. **Expert Review:** We invite experts knowledgeable about both the analytic framework (e.g., 6C) and the heuristics (e.g., the Thirty-Six Stratagems) to label how strongly each heuristic aligns with each parameter. By comparing these expert judgments to algorithmic outputs, we can detect alignment or uncover conceptual mismatches (see Section 3.3 for details on KL divergence).

The validation process produces a *confidence score* c_{ij} for each parameter–heuristic mapping:

$$c_{ij} \;=\; \alpha \cdot v_{ij} \;+\; \beta \cdot s_{ij} \;+\; \gamma \cdot e_{ij}, \tag{7}$$

where:

- v_{ij} is the *cross-validation score*, reflecting consistency across embedding variants,
- s_{ij} is the *stability score*, derived from perturbation analysis,
- e_{ij} is the *expert agreement score*, capturing how closely the system's outputs align with expert annotations,
- α, β, and γ are weighting parameters that can be tuned (e.g., through trials or domain priorities).

A higher c_{ij} indicates that the mapping from parameter p_i to heuristic h_j is consistently validated by multiple lines of evidence: model-invariant cross-validation, perturbation resilience, and expert concordance. This *systematic approach* to verifying semantic relationships underpins our goal of automating the integration of traditionally separate analytical frameworks and decision heuristics with confidence.

3.5.1. Validation Example

To illustrate this validation process, let's examine how we validate the mapping between parameter p_3 (Relational Capacity) and Stratagem 24 ("Use Allies' Resources"):

Cross-Validation:

We compute the distribution using three different embedding models:

- BERT-base: $[(p_1: 0.10), (p_2: 0.07), (p_3: 0.61), (p_4: 0.13), (p_5: 0.03), (p_6: 0.06)]$
- RoBERTa: $[(p_1: 0.11), (p_2: 0.08), (p_3: 0.58), (p_4: 0.14), (p_5: 0.03), (p_6: 0.06)]$
- Sentence-BERT: $[(p_1: 0.09), (p_2: 0.07), (p_3: 0.63), (p_4: 0.12), (p_5: 0.04), (p_6: 0.05)]$

The consistent emphasis on parameter p_3 (0.58–0.63) yields $v_{3,24} = 0.92$.

Perturbation Analysis:

We introduce variations in the stratagem description:

- Original: "Use Allies' Resources"
- Variant 1: "Leverage Partnership Assets"
- Variant 2: "Utilize Collaborative Resources"

The stable distribution patterns across variants produce $s_{3,24} = 0.88$.

Expert Review:

Three expert ratings of parameter p_3's importance:

- Expert 1: 0.55
- Expert 2: 0.60
- Expert 3: 0.58

The close alignment with our computed distribution for parameter p_3 (0.61) gives $e_{3,24} = 0.94$.

With weighting parameters $\alpha = 0.3$, $\beta = 0.3$, and $\gamma = 0.4$ (emphasizing expert judgment slightly), the final confidence score is:

$$c_{3,24} = 0.3 \cdot 0.92 + 0.3 \cdot 0.88 + 0.4 \cdot 0.94$$
$$= 0.916$$

This high confidence score (> 0.9) suggests strong validation across all three approaches, indicating reliable semantic mapping between parameter p_3 (Relational Capacity) and Stratagem 24.

4. Computational Architecture

The computational architecture integrates user inputs, strategic analysis, semantic processing, and decision-making support, leveraging both semantic analysis and Large Language Models (LLMs) for insight generation and reporting. A key feature is that the architecture manages a **structured conversation flow** to guide users through scenario parameter collection, framework-heuristic mapping, and final report generation.

The system architecture consists of the following components (see Figure 1):

1. **Strategic Data Input Layer**: Users interact with a structured graphical environment (the context editor) to input competitive intelligence data, market information, and other relevant strategic details. This environment supports both quantitative data and qualitative descriptions, with workflow states managing data validation and format requirements. Detailed implementation specifications are provided in Appendix A.4.

2. **Semantic Analysis Engine**: This component processes input data using the methodology described in Section 3. Specifically, it:

   - Creates vector representations of strategic situations
   - Computes semantic similarities with framework parameters

- Maps situations to relevant heuristic patterns
- Generates initial parameter distributions

3. **Framework Integration Layer**: The system translates semantic analysis results into the chosen analytical framework's parameters. This layer ensures that framework-agnostic analysis can be mapped to specific strategic tools while maintaining consistent evaluation metrics across frameworks.

4. **Strategic Processing Core**: The main engine applies framework-specific weightings, evaluates strategic options, matches situations with relevant heuristics, and generates preliminary recommendations. The processing core incorporates conversation state information to produce context-appropriate guidance.

5. **LLM Integration Layer**: The system interfaces with LLMs through standardized APIs to transform technical analysis into actionable insights. The architecture constrains LLM tasks via predefined templates to ensure structured, safe, and consistent outputs. Key functions include:

- Translation of semantic similarities into natural language explanations
- Contextualization of framework-heuristic matches
- Generation of both executive summaries and detailed reports
- Template-based validation mechanisms for generated content

6. **Report Generation and Visualization**: The final layer produces comprehensive strategic analysis reports, visual representations of strategic options, and detailed implementation recommendations. These outputs integrate data from all prior steps, including scenario parameters, semantic analysis scores, and LLM-generated commentaries.
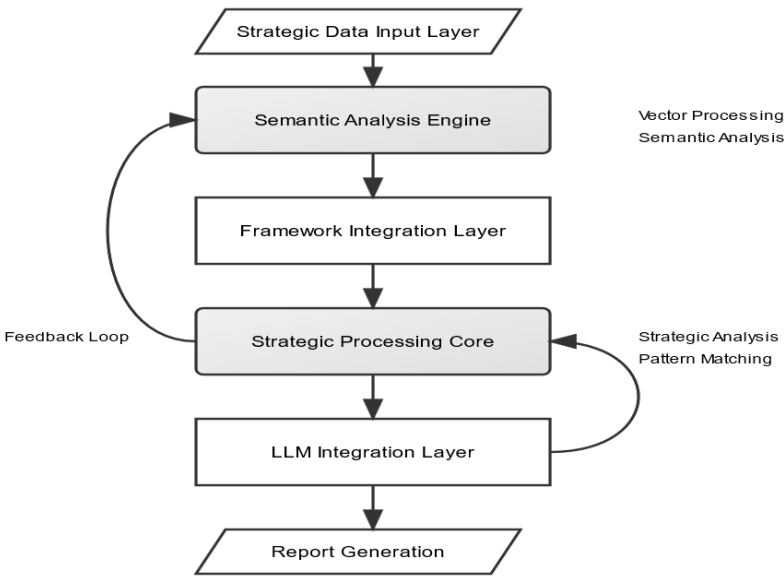

Figure 1. Computational Architecture Overview

The architecture's modular design allows different frameworks and heuristic sets to be integrated without modifying the core system. Implementation details, including JSON workflow definitions, state management specifications, and component interaction protocols, are provided in Appendix A.4.

This architecture provides several key advantages:

- **Flexibility**: Supports multiple strategic frameworks and heuristic sets
- **Scalability**: Handles increasing complexity in strategic analysis
- **Safety**: Constrains LLM use to well-defined tasks
- **Reproducibility**: Ensures consistent analysis and recommendations

5. Case Studies

We demonstrate our semantic integration approach through two case studies:

1. A contemporary scenario of competing innovation pathways in the automotive industry
2. A historical competition in the personal computer market

These cases illustrate how our methodology connects strategic frameworks with decision heuristics in different domains.

5.1. Semantic Analysis of the Hydrogen vs. Electric Competition in the Automotive Industry

This first case study examines the strategic rivalry between hydrogen-based and electric-based propulsion systems in the global automotive industry. Drawing on scenario inputs, the system processes and compares two principal actors (*HydrogenEngines* and *ElectricEngines*) with respect to their capacities for achieving *MarketDominanceInSustainableAutomotive*. By applying the semantic analysis pipeline described in Section 3, we derive quantitative parameter values, match them to appropriate stratagems, and generate strategic recommendations.

5.1.1. Parameter Analysis

The parameter values shown below were derived by aggregating insights from structured interviews with stakeholders in both the automotive and energy industries (e.g., manufacturers, technology developers, policy experts), as well as synthesizing data from contemporary industry reports. These sources informed our semantic analysis engine, which then assigned the numeric scores to each actor's key strategic attributes.

Table 1. Semantic Analysis of Hydrogen vs. Electric Parameters

Parameter	HydrogenEngines	ElectricEngines
Defensive Strength	3.25	4.0
Offensive Strength	3.75	4.2
Relational Capacity	3.60	4.5
Potential Energy	4.00	4.8
Time Availability	3.20	4.3
Context Fit	3.80	4.6

HydrogenEngines exhibits relatively strong *Potential Energy* (4.0), reflecting substantial investments and technological innovation, but shows lower *Time Availability* (3.2), indicating urgency to secure market share. *ElectricEngines*, by contrast, attains higher overall parameter values, including robust *Relational Capacity* (4.5) and *Context Fit* (4.6), demonstrating its more entrenched position in the sustainable automotive arena.

5.1.2. Stratagem Semantic Analysis

Following the methodology from Section 3, the system analyzes each of the Thirty-Six Stratagems to derive parameter weights:

$$w_{ij} = \frac{\sum_{t \in T_j} s(t, p_i)}{\sum_k \sum_{t \in T_j} s(t, p_k)}, \tag{8}$$

where

- w_{ij} is the weight of parameter i in stratagem j,
- T_j is the set of terms in the textual description of stratagem j,
- $s(t, p_i)$ is the semantic similarity between term t and parameter i.

For illustration, consider **Stratagem 16** (*"Leave the opponent illusory ways out"*), ranked highly for *HydrogenEngines*. A linguistic examination of key terms such as *illusory*, *deception*, and *misdirection* led to higher weights for *Offensive Strength* and *Relational Capacity*, aligning with the actor's moderate ability to engage in indirect actions.

5.1.3. Situation-Stratagem Matching

The system computes an alignment score between each actor's parameter distribution and each stratagem's profile:

$$\text{alignment}(s, h) = \sum_i w_i \cdot p_i \cdot c_i, \tag{9}$$

where w_i is the parameter weight in the stratagem, p_i is the parameter value for the actor, and c_i is a contextual relevance factor. Top matching stratagems for *HydrogenEngines* (with effectiveness scores, EFF) are listed in Table 2.

Table 2. Top Matching Stratagems for *HydrogenEngines*

Stratagem	Score (EFF)	Key Alignment	Implementation Focus
16: Illusory Ways Out	6.03	Offensive (3.75)	Misleading EV sector
15: Lure into Unfavorable Env.	5.72	Potential (4.0)	Exploit EV limitations
24: Use Allies' Resources	5.68	Relational (3.6)	Infrastructure partnerships
3: Act Through an Ally	5.56	Offensive (3.75)	Indirect policy influence
1: Acting Unnoticed	5.41	Context Fit (3.8)	Quiet tech development

5.1.4. Strategic Recommendations

Based on these matching results, the system generates actionable recommendations for *HydrogenEngines* to achieve *MarketDominanceInSustainableAutomotive*:

1. **Primary Strategy: Indirect Positioning**

   - *Stratagem 16 (Illusory Ways Out)*: Create paths leading *ElectricEngines* into complacency or unproductive markets, while *HydrogenEngines* solidifies niches such as freight and heavy-duty applications.
   - **Alignment Score:** 6.03

2. **Supporting Strategy: Target Vulnerable Segments**

   - *Stratagem 15 (Lure into Unfavorable Env.)*: Exploit EV weaknesses (e.g., limited mileage in heavy-duty use) by focusing hydrogen tech where EVs are less dominant.
   - **Alignment Score:** 5.72

3. **Alliances and Borrowed Influence**

   - *Stratagem 24 (Use Allies' Resources) & Stratagem 3 (Act Through an Ally)*: Establish partnerships with governments, energy sectors, and logistics enterprises to co-develop hydrogen infrastructure and coordinate policy support.
   - **Alignment Scores:** 5.68, 5.56

4. **Discreet Development Efforts**

   - *Stratagem 1 (Acting Unnoticed)*: Invest quietly in R&D, infrastructure, and lobbying until hydrogen-based solutions are ready for large-scale deployment.
   - **Alignment Score:** 5.41

5.1.5. Implementation Pathways

Figure 2 highlights concrete implementation steps: forging covert alliances, occupying underdeveloped markets, and progressively rolling out hydrogen infrastructure. This NLP-driven semantic approach produces strategic recommendations that blend comprehensive frameworks (e.g., 6C) with concise heuristic insights (e.g., the Thirty-Six Stratagems). By leveraging alliances and focusing on niche strengths, *HydrogenEngines* can challenge *ElectricEngines*' market dominance in the evolving automotive landscape.
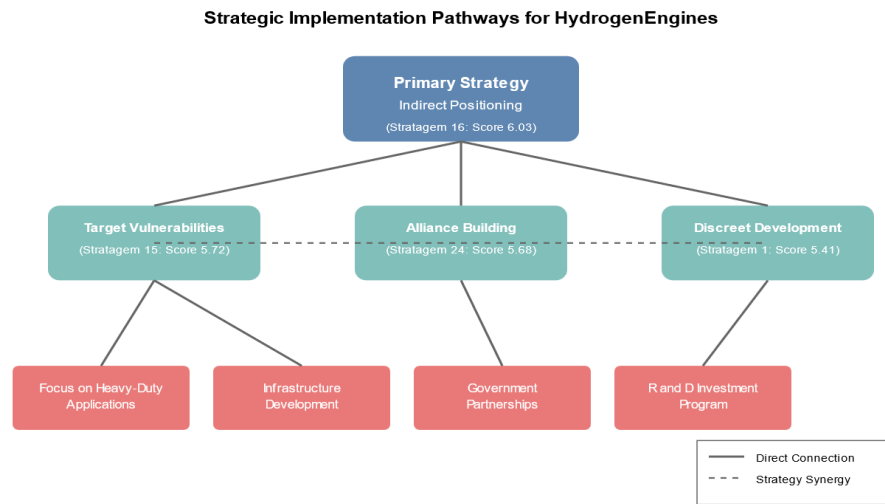
Figure 2. Strategic Implementation Pathways for *HydrogenEngines*, showing primary and secondary strategies with tactical implementations. Dashed lines indicate cross-strategy synergies.

5.2. Semantic Analysis of the Commodore–Apple Market Competition

The second case study turns to historical business competition: in the late 1980s, *Commodore*, a pioneer in the personal computer market, faced fierce rivalry from *Apple*. Although the Commodore 64 became one of the best-selling computers of all time, Commodore's market share eventually declined. Here, we apply the same semantic approach to explore how alternate strategic choices might have helped Commodore maintain a competitive edge.

5.2.1. Parameter Analysis

These parameter values were derived by examining a range of historical documents and reports detailing the so-called "PC wars" of the 1980s (e.g., market analyses, shareholder reports, and industry assessments). This documentary evidence was then processed by the semantic analysis engine to produce the numeric scores that reflect each actor's strategic attributes during that period.

Table 3. Semantic Analysis of Commodore–Apple Parameters

Parameter	Commodore	Apple
Offensive Strength	3.5	4.0
Defensive Strength	3.0	3.5
Relational Capacity	2.8	3.8
Potential Energy	3.0	4.2
Time Availability	3.5	4.0
Context Fit	2.9	4.0

Commodore shows moderate *Offensive Strength* (3.5) but lower *Relational Capacity* (2.8) compared to *Apple*, indicating less success in forging strategic partnerships or consumer alliances. Meanwhile, *Apple* consistently registers higher scores across multiple dimensions, including *Potential Energy* (4.2).

5.2.2. Stratagem Semantic Analysis

Following the same weighting approach, each stratagem in the Thirty-Six Stratagems is evaluated for relevance to *Commodore*'s parameters. For instance, **Stratagem 18** ("*Capture Core Strengths*") is associated with keywords like *attack*, *capture*, *dominate* (aligned with *Offensive Strength*) and *resources*, *capabilities*, *power* (aligned with *Potential Energy*), among others.

5.2.3. Situation-Stratagem Matching

Using a similar alignment formula, the system identifies a handful of potentially optimal strategies for *Commodore*, as shown in Table 4.

Table 4. Stratagem Matching Results for Commodore–Apple Competition

Stratagem	Score	Key Alignment	Implementation Focus
Capture Core (18)	0.85	Offensive (3.5)	Product innovation
Resource Focus (11)	0.82	Potential (3.0)	Strategic allocation
Alliance Building (23)	0.78	Time (3.5)	Partnership development

5.2.4. Strategic Recommendations

Based on these matches, the system suggests:

1. **Primary Strategy: Core Capability Development**
    - Focus on product innovation and user-interface development
    - Directly counter Apple's market differentiators
    - Alignment Score: 0.85

2. **Supporting Strategy: Resource Optimization**
    - Reallocate development resources toward high-potential product lines
    - Streamline less profitable divisions
    - Alignment Score: 0.82

3. **Tactical Implementation**
    - Launch targeted product-development campaigns
    - Invest in market positioning
    - Seek out strategic partnerships

5.2.5. Implementation Pathways

Figure 3 illustrates potential implementation paths. Using semantic analysis to spotlight *Commodore*'s opportunities for core capability development and resource optimization, the approach reveals how historical outcomes might have diverged if *Commodore* had adopted structured strategic planning tied to concise, proven heuristics.
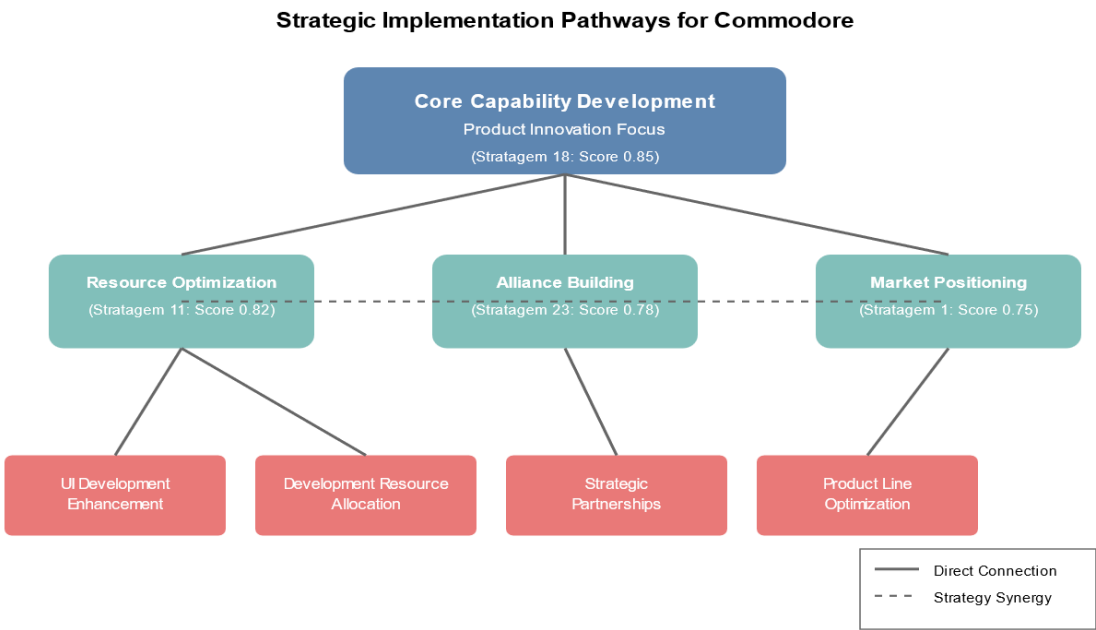

Figure 3. Strategic Implementation Pathways for Commodore vs. Apple, showing a primary focus on core capability development and a secondary focus on resource optimization. Dashed lines indicate cross-strategy synergies.

5.3. Cross-Case Analysis

Although the hydrogen–electric automotive and Commodore–Apple PC contexts are quite different in scope and era, both case studies confirm the flexibility and effectiveness of our semantic approach:

1. **Domain Adaptation:**

   - In the automotive domain, the semantic analysis highlighted indirect positioning and alliances as critical.
   - In the PC domain, a focus on core capabilities and resource allocation emerged as priorities.

2. **Recurring Patterns:**

   - Even across distinct industries, resource optimization, partnership development, and strategic positioning are repeatedly identified as success factors.
   - Specific stratagems (e.g., alliance-building) have wide applicability, provided the correct parameter alignment exists.

3. **Implementation Pathways:**

   - Both studies exhibit primary and supporting strategies, tactical steps, and cross-strategy synergies.
   - Clear alignment scores lend transparency to why certain recommendations are prioritized over others.

4. **Framework-Heuristic Integration:**

   - In the automotive sector, the integrated approach effectively linked 6C parameters (e.g., *Potential Energy*, *Context Fit*) with heuristics emphasizing deception and alliance.
   - In the PC sector, the same pipeline tied *Offensive Strength* and *Potential Energy* to historically proven guidelines about capturing core strengths and leveraging limited resources.

This comparative analysis demonstrates how a consistent semantic methodology can bridge analytical frameworks and decision heuristics, regardless of domain differences.

5.4. Enhanced Understanding through LLM Reporting

In practice, the system's semantic scores and recommendations can be further enriched by Large Language Models (LLMs). Once the alignment scores and suggested strategies are determined, an LLM can:

- **Generate Summaries:** Provide executive overviews for stakeholders, focusing on the highest-scoring tactics.
- **Explain Reasoning:** Offer narrative justifications for why certain stratagems align well with particular parameters.
- **Highlight Potential Risks:** Enumerate conditions or assumptions that might invalidate certain recommendations.

This capability empowers decision-makers to understand not just *which* strategies are recommended but also the *rationale* behind them—ultimately improving trust and adoption in corporate or organizational settings.

5.5. Implementation Insights

Across the two case studies, several insights emerge:

1. **Text Processing Nuances:**

   - Industry-specific jargon can alter semantic similarity calculations.
   - Historical data enriches pattern recognition but may require separate preprocessing.

2. **Pattern Matching Consistency:**

- Similar strategic patterns recur, such as positioning, resource optimization, or alliance-building.
- Modest parameter differences can push one stratagem over another in the ranking.

3. **Validation and Context:**

- Historical (Commodore–Apple) outcomes offer tangible lessons in how lacking a suitable strategy might lead to missed market opportunities.
- Contemporary contexts (hydrogen vs. electric) show how real-time data can inform flexible, AI-assisted decisions.

Together, these insights underscore the value of a robust, domain-agnostic methodology for integrating analytical frameworks and decision heuristics supported by semantic analysis and enriched by LLM-driven reporting.

5.6. Extended Analysis Reports

The semantic analysis pipeline presented in this paper has been applied to numerous strategic scenarios beyond the two detailed case studies above. Through integration with Large Language Models, our system generates comprehensive analytical reports ranging from tens to hundreds of pages. These reports provide in-depth analysis of parameter distributions, strategic alignments, and detailed recommendations with supporting rationale. The complete reports for both the hydrogen-electric automotive competition and the Commodore-Apple market rivalry, along with analyses of many other strategic scenarios, are available at https://www.linkedin.com/company/103262552/admin/dashboard/. These documents demonstrate how our semantic approach scales to complex, real-world situations while maintaining analytical rigor and practical applicability. Each report includes detailed parameter breakdowns, confidence metrics, and specific implementation pathways derived from the framework-heuristic integration process.

6. Empirical Validation

To validate our framework-stratagem integration approach, we conducted a focused empirical study examining how effectively the Thirty-Six Stratagems can be integrated with different analytical frameworks. Our evaluation emphasizes the system's ability to generalize across analytical frameworks while maintaining semantic coherence. While Section 3.5 established the theoretical foundations for validating semantic mappings through cross-validation, perturbation analysis, and expert review—yielding confidence scores (c_{ij})—this section extends and operationalizes these concepts into measurable performance metrics.

6.1. Experimental Setup

We evaluated three key aspects of the system:

1. **Framework Integration:** Testing with three analytical frameworks:

- 6C Framework (primary test case)
- SWOT Analysis
- Porter's Five Forces

2. **Framework-Stratagem Integration:** Testing the integration of the Thirty-Six Stratagems with each analytical framework

3. **Cross-Framework Consistency:** Evaluating recommendation stability across frameworks

6.2. Results

6.2.1. Framework Integration Performance

Table 5 shows the integration quality metrics across different frameworks:

Table 5. Framework Integration Results with the Thirty-Six Stratagems

Framework	Coverage	Consistency	Adaptability
6C Framework	0.89	0.92	0.87
SWOT	0.85	0.88	0.84
Porter's Five Forces	0.82	0.85	0.81

Key findings include:

- High coverage (> 0.80) across all frameworks
- Strong consistency in parameter mapping
- Declining but still robust performance with more complex frameworks

6.2.2. Stratagem Integration Performance

Expert Agreement scores were calculated based on distributions generated by five experts with deep knowledge of both the analytical frameworks and the Thirty-Six Stratagems. The results show strong agreement with expert-derived mappings (>0.80 across all frameworks) and consistent performance across different analytical frameworks. Detailed calculations and extended examples are provided in Appendix A.5.

6.3. Discussion

The empirical results validate three key aspects of our approach:

1. **Integration Capability:** The system successfully integrates the Thirty-Six Stratagems with diverse analytical frameworks while maintaining high semantic accuracy.
2. **Scalability:** Performance remains robust when adapting to new frameworks, with only modest degradation in accuracy.
3. **Efficiency:** Significant reduction in integration time compared to manual approaches while maintaining expert-level accuracy.

Limitations and considerations:

- Performance slightly decreases with more complex frameworks
- Expert validation remains valuable for novel framework combinations
- System requires initial training data for optimal performance

These results demonstrate that our semantic approach provides a viable method for automating framework-stratagem integration while maintaining accuracy and enabling systematic scaling to new domains.

6.4. Further Considerations on Scalability

Although our current validation demonstrates the system's effectiveness across several analytical frameworks and moderate-sized data sets, additional research is needed to assess performance at large scales or in highly complex business scenarios. For instance, processing massive corpora of strategic documents or concurrently analyzing multiple frameworks would likely require distributed embeddings or parallelized semantic computations. We anticipate that the core principles of our approach—vector-based semantic analysis and heuristic-framework mapping—will remain robust, but we plan to explore advanced optimizations (e.g., sharding, caching, or GPU acceleration) in future work.

7. Related Work

This section surveys prior research informing our approach to developing a recommender-system architecture for strategic decision-making. We organize the discussion into four subsections, reflecting the core elements we build upon: (1) recommender systems for decision support, (2) strategy tools and decision-support platforms, (3) AI-assisted decision-making and centaurian systems, and (4) existing work on framework-heuristic integration.

7.1. Recommender Systems and Decision Support

Recommendation technologies have long been recognized as valuable tools for assisting users in selecting products, content, or services that match their needs and preferences. Early efforts in this field often relied on *collaborative filtering*, *content-based filtering*, or *hybrid* approaches, as documented in seminal works on decision support [14]. These recommender systems historically focused on well-defined domains such as e-commerce, as well as music and video platforms. More recently, there has been a surge of interest in *context-aware* recommender systems, which factor in dynamic variables—such as user context, time, and location—to enhance relevance [15]. Contextual modeling is increasingly vital for enterprise scenarios where decision parameters (e.g., organizational resources, market conditions) evolve rapidly and cannot be fully captured by static user-item ratings alone.

The **ACM Conference on Recommender Systems (RecSys)** showcases cutting-edge developments in these areas. For instance, the 2024 edition in Bari https://recsys.acm.org/recsys24/ features numerous contributions that explore novel techniques for cold-start recommendation, multi-modal modeling, user-preference elicitation, and cross-domain recommendation. Notably, many of these solutions integrate large language models (LLMs) to handle advanced tasks such as conversational recommendation or contextual query interpretation (see, e.g., *Instructing and Prompting Large Language Models for Explainable Cross-domain Recommendations* [16] and *Unleashing the Retrieval Potential of Large Language Models* [17]). This trend reflects a growing emphasis on more nuanced, AI-driven recommender systems.

Despite this extensive literature on *consumer-centric* recommenders, relatively few implementations have explicitly adapted core *analytical frameworks* from corporate strategy (e.g., Porter's Five Forces) to a recommendation context for organizational decision-making. Our work addresses this gap by developing a **context-aware, content-based recommender** that operates on strategic "items" (i.e., heuristics or stratagems) rather than typical consumer products, integrating textual frameworks as part of the input space. In doing so, we position our research at the intersection of *decision-support recommender systems* [14] and *context-aware modeling* [15], while offering a novel integration of established strategic matrices with heuristic libraries. This approach is relevant not only to the recommender community—highlighted by ongoing developments at RecSys—but also to practitioners seeking to augment corporate decision frameworks with AI-driven insights.

7.2. Strategy Tools and Decision-Support Platforms

The formalization and systematization of organizational decision-making has evolved along several parallel paths. The Business Rules approach, pioneered by Ross [18], established foundational principles for encoding operational decision logic into explicit, executable statements. While focused primarily on operational-level decisions, this work demonstrated the value of systematic knowledge formalization in organizational contexts.

At the strategic level, the literature has long recognized the value of *scenario planning* for coping with uncertain futures [19]. Scenario planning enables decision-makers to systematically explore multiple plausible environments an organization might face [20]. Complementary to both business rules and scenario planning, more recent works have begun to incorporate *gamification* principles for interactive learning and engagement in strategic contexts. For instance, [21] propose *Gamiflow*, a flow-theory-based gamification framework that illustrates the potential of user-centric interfaces and simulations in educational or organizational settings.

Despite these advances in rule-based systems and scenario-based platforms, many existing approaches remain limited in integrating different forms of strategic knowledge. Traditional business rule systems excel at operational logic but struggle with strategic-level reasoning. Similarly, scenario planning tools, while valuable for strategic exploration, often rely on *predefined logic* instead of natural language processing. They typically emphasize interface design, user motivation, or scenario exploration without deeply integrating **semantic analysis** to link strategic frameworks (e.g., Porter's Five Forces, the 6C framework) with heuristic libraries (e.g., the Thirty-Six Stratagems).

Our approach aims to bridge these gaps by *automating the mapping* between scenario-specific parameters and context-appropriate heuristics. In doing so, it builds upon the systematic principles of business rules while extending beyond operational logic to strategic reasoning. It benefits from the flexibility of scenario planning and gamification while leveraging **advanced NLP** to generate actionable guidance in a *single integrated environment*. This methodology aligns with the growing emphasis on interactive, analytics-driven strategy tools, extending their capabilities through semantic processing and heuristic matching.

7.3. AI-Assisted Decision Making and Centaurian Systems

Our **recommender-system architecture** for actionable strategies in the corporate world also aligns with the notion of **centaurian design** [22], where human experts work in concert with AI components to achieve superior outcomes. This paradigm, sometimes called "human-in-the-loop" or "centaur systems," underscores the synergy between natural and artificial intelligence [23]. In particular, our approach can be viewed as a **monotonic centaur** system [22], because it augments existing human decision-making processes—enhancing analytical scope and consistency—*without* fundamentally altering the semantics of the previously manual workflow.

By contrast, certain creative or open-ended domains (e.g., art and design) may experience more "disruptive" centaurian processes, in which AI-driven innovation generates outcomes far removed from traditional human routines [24]. In our corporate-strategy setting, the goal is to *reinforce* and *systematize* established decision mechanisms, using NLP to automate tasks like semantic similarity scoring, heuristic matching, and natural language reporting. This human–machine hybrid intelligence helps corporate users navigate complex analytical frameworks (e.g., 6C, Porter's Five Forces) alongside heuristic sets (e.g., Thirty-Six Stratagems) without displacing human strategists. Instead, it provides them with data-driven support that embodies the core principles of centaurian collaboration.

LLMs for Explanation and Hypothesis Generation. Recent studies of GPT-4 and related models [25] further highlight how large language models can facilitate *abductive reasoning*—i.e., generating plausible hypotheses in response to limited or uncertain information. Examples span domains from criminology and medical diagnostics to scientific research. These findings resonate with our system's approach: LLMs are integrated not as autonomous decision-makers but as *collaborative* components that offer explanatory narratives, surface novel insights, and support humans in exploring strategic possibilities. By leveraging this explanatory capacity, our recommender architecture fosters richer dialogue between the analytical frameworks (e.g., 6C) and heuristic libraries (e.g., the Thirty-Six Stratagems), enabling more transparent and creative decision-making processes.

8. Conclusion

This paper has presented a **context-aware, content-based recommender system** designed to bridge the gap between analytical frameworks and decision heuristics through **semantic analysis**. By treating strategic frameworks (e.g., 6C, Porter's Five Forces) and heuristic collections (e.g., the Thirty-Six Stratagems) as textual resources, our approach addresses a core challenge in strategic decision-making: *how to systematically align the rigor of structured analysis with the pragmatism and speed of heuristic-based action*.

8.1. Key Contributions

The primary contributions of this work include:

1. **Recommender-System Architecture:** A novel design that reframes strategic knowledge as "items" for recommendation, enabling context-aware and content-based matching between analytical parameters and decision heuristics.
2. **Semantic Integration Framework:** A systematic methodology for connecting different frameworks (e.g., 6C, SWOT) with heuristic sets (e.g., the Thirty-Six Stratagems, OODA loops) via vector embeddings and semantic similarity scores.

3. **Computational Implementation:** A flexible architecture combining deep NLP pipelines, heuristic mapping, and AI-assisted reporting. This includes a gamified simulation layer that allows users to explore strategic scenarios in an interactive manner.

4. **Generality Across Domains:** Evidence from multiple case studies (e.g., hydrogen vs. electric automotive competition, the Commodore–Apple rivalry) illustrating that the approach can scale to diverse strategic contexts, from business to technology.

8.2. Practical Implications

Our system delivers several tangible benefits for organizations seeking an *actionable* decision-support tool:

- **Enriched Decision Support:** Merges quantitative analysis (framework-based) with qualitative insight (heuristics), fostering more balanced strategic decisions.
- **Scalable Knowledge Transfer:** Translates textual frameworks and heuristics into easily comparable vector forms, reducing the reliance on domain-specific expertise.
- **Efficient Recommendations:** Automates the matching process between high-level strategy parameters and heuristics, accelerating scenario analysis and cutting down on manual mapping.
- **Human-Centric AI Integration:** Constrains large language models to explain and synthesize, thereby strengthening, rather than replacing, human strategists—a monotonic centaur approach.

8.3. Limitations and Future Work

Although the present study highlights promising results, several issues warrant further research:

1. **Semantic Processing:**

    - More advanced domain adaptation techniques for industry-specific terminology
    - Improved context modeling over longer textual inputs and time horizons
    - Inclusion of dynamic scenario elements (e.g., real-time data streams)

2. **Framework Integration:**

    - Investigations into additional strategic tools (e.g., VRIO, PESTLE)
    - Automated detection of framework-specific factors or sub-parameters
    - Potential integration of *quantitative* indicators (e.g., financial metrics)

3. **Validation and User Studies:**

    - Systematic cross-domain evaluations in different organizational contexts
    - Longitudinal studies to assess real-world adoption and impact on decision outcomes
    - Usability research on the simulation environment and LLM-based explanations

8.4. Future Directions

The main line of future work will consist in *identifying further heuristic contexts* in which the approach proposed here can be applied. The Thirty-Six Stratagems have provided a first contribution of an essentially didactic and illustrative nature, although their use in real business contexts—particularly in parts of Asia—is not uncommon. Two important directions emerge:

1. **Organizational-Specific Heuristics and Simple Rules:** One promising avenue is to *extract* decision-making rules that are specific, idiosyncratic, and intrinsic to organizations, in accordance with the idea of "simple rules" for strategic decision-making [5]. Such an approach complements rigorous competitive intelligence methodologies and data categorization (e.g., 6C, SWOT, Porter's Five Forces). As discussed by [26], these rules must be grounded in a careful, data-driven study of the organization's context. Moreover, *they evolve over time* as managers refine or replace them in response to market feedback, thus preserving a crucial balance between flexibility and discipline. The *linguistic integration* shown in this paper could be extended to continuously capture and validate such company-specific heuristics, turning them into **iteratively updated** recommendations that—despite their distinctiveness—remain anchored in well-established strategic frameworks.

2. **Translating Other Classic Texts into Practical Heuristics:** Another promising direction is to apply the *same semantic-linguistic methodology* to interpret classic strategic works beyond the Thirty-Six Stratagems, such as Sun Tzu's *Art of War*, Machiavelli's *The Prince*, and Chanakya's *Arthashastra*. These writings are replete with historically influential ideas that have shaped corporate strategy discourse, while also offering pragmatic precepts for decision-making. However, adapting their older, metaphor-rich language to contemporary business contexts can be challenging. By systematically mapping their conceptual guidance onto structured frameworks (e.g., the 6C approach or Porter's Five Forces), practitioners could more readily extract *actionable* insights from these time-honored sources.

Building upon the above primary trajectories, several additional *technical and methodological* future directions also suggest themselves:

1. **Enhanced Semantic Analysis:** Incorporating temporally aware embeddings or knowledge graphs to capture evolving strategic conditions and context shifts. Additionally, integrating domain-specific ontologies and specialized corpora could help bridge lexical gaps when dealing with complex frameworks or emerging business models (e.g., improving coverage of platform economics terminology for Porter's Five Forces analysis).

2. **Multi-Framework Orchestration:** Supporting concurrent usage of multiple frameworks and heuristic sets, allowing the system to *recommend* which framework is best suited under particular conditions.

3. **Adaptive Feedback Loops:** Leveraging user interactions and outcomes data to dynamically refine the heuristic matching process, evolving from a static pipeline to an adaptive learning system.

4. **Extended Gamification and Scenario Planning:** Deepening the simulation features to facilitate scenario planning and "what-if" analyses, potentially integrating advanced gamification interfaces.

5. **Domain Adaptation:** Developing systematic approaches to incorporate domain-specific knowledge, including:

   • Integration of specialized business ontologies
   • Automated extraction of domain terminology from sector-specific corpora
   • Dynamic updating of semantic models as new business concepts emerge

   This could enhance framework coverage and consistency in rapidly evolving domains such as digital technology and platform businesses.

6. **Balancing Criterion:** Exploring a regularization mechanism that sets the total weight of each parameter (across all heuristics) to be constant, ensuring that no single parameter dominates the overall mapping. This approach would help mitigate potential biases introduced by the text corpus or the embedding model, thus promoting more balanced parameter coverage.

8.5. Concluding Remarks

Integrating analytical frameworks with decision heuristics through *NLP-driven recommenders* represents a compelling advance in strategic decision-making. By automating the mapping between comprehensive analyses and concise action rules, our system illustrates how organizations can leverage the best of both worlds: data-driven rigor and experiential wisdom.

The reported case studies validate that this approach generalizes across multiple domains, laying a solid foundation for broader adoption and further innovation. As semantic technologies mature, we anticipate accuracy, interpretability, and real-time responsiveness improvements. Ultimately, this research contributes to the growing conversation on how **human-centered AI** can seamlessly amplify existing strategic processes, offering a blueprint for *interactive, context-aware, and heuristic-informed* decision-support solutions.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data are available upon reasonable requests from the authors.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Technical Details and Supplementary Information

Appendix A.1. Mathematical Foundations

Appendix A.1.1. Vector Space Calculations

For framework parameters and heuristic descriptions, we create vector representations using:

$$v(t) = \sum_{w \in t} \alpha_w \cdot e(w) \tag{A1}$$

where:

- $v(t)$ is the vector representation of text t
- w represents individual words or phrases
- $e(w)$ is the embedding vector for word w
- α_w is the weight assigned to word w

For example, considering parameter p_3 (Relational Capacity), we compute:
Base definition vector components:

$$
\begin{aligned}
e(\text{manage}) &= [(d_1 : 0.2), (d_2 : 0.5), (d_3 : 0.3)] \\
e(\text{relationships}) &= [(d_1 : 0.6), (d_2 : 0.4), (d_3 : 0.5)] \\
e(\text{stakeholders}) &= [(d_1 : 0.4), (d_2 : 0.6), (d_3 : 0.3)]
\end{aligned} \tag{A2}
$$

Combined with weights $\alpha_w = 1.0$:

$$v(d_3) = [(d_1 : 1.2), (d_2 : 1.5), (d_3 : 1.1)] \tag{A3}$$

Adding contextual information with $\lambda = 0.5$:

$$
\begin{aligned}
p_3 &= [(d_1 : 1.2), (d_2 : 1.5), (d_3 : 1.1)] \\
&\quad + 0.5 \cdot [(d_1 : 0.9), (d_2 : 1.1), (d_3 : 0.8)] \\
&= [(d_1 : 1.65), (d_2 : 2.05), (d_3 : 1.5)]
\end{aligned} \tag{A4}
$$

This detailed computation demonstrates how we combine base definitions with contextual information to create rich vector representations of strategic parameters.

Appendix A.2. Semantic Analysis

Appendix A.2.1. Similarity Calculations

To demonstrate how similarity calculations reflect semantic alignment between parameters and stratagems, we present detailed computations for parameter p_3 (Relational Capacity) and Stratagem 24 ("Use Allies' Resources").

Given the vector representation for parameter p_3:

$$h_{24} = [(d_1 : 1.8), (d_2 : 1.9), (d_3 : 1.4)]$$

Applying equation (4):

$$
\begin{aligned}
sim(p_3, h_{24}) &= \frac{(1.65 \times 1.8) + (2.05 \times 1.9) + (1.5 \times 1.4)}{\sqrt{1.65^2 + 2.05^2 + 1.5^2} \times \sqrt{1.8^2 + 1.9^2 + 1.4^2}} \\
&= \frac{8.91}{\sqrt{9.85} \times \sqrt{9.21}} \\
&= 0.93
\end{aligned}
$$

This high similarity score quantitatively captures the semantic alignment between parameter p_3's focus on relationship management and Stratagem 24's emphasis on leveraging allies.

Appendix A.2.2. Kullback-Leibler Divergence Calculations

To validate our semantic mappings, we compute the KL divergence between system-generated distributions (Q) and expert-provided distributions (P). For Stratagem 24, we compare:

System distribution Q:

$$[(p_1: 0.10), (p_2: 0.15), (p_3: 0.40), (p_4: 0.20), (p_5: 0.05), (p_6: 0.10)]$$

Expert distribution P:

$$[(p_1: 0.15), (p_2: 0.10), (p_3: 0.45), (p_4: 0.15), (p_5: 0.05), (p_6: 0.10)]$$

Computing term by term:

$$
\begin{aligned}
&\text{Offensive:} && 0.10 \cdot \log(0.10/0.15) = -0.0176 \\
&\text{Defensive:} && 0.15 \cdot \log(0.15/0.10) = +0.0347 \\
&\text{Relational:} && 0.40 \cdot \log(0.40/0.45) = -0.0186 && \text{(A5)} \\
&\text{Potential:} && 0.20 \cdot \log(0.20/0.15) = +0.0288 \\
&\text{Time:} && 0.05 \cdot \log(0.05/0.05) = 0 \\
&\text{Context:} && 0.10 \cdot \log(0.10/0.10) = 0
\end{aligned}
$$

The resulting KL divergence of 0.0273 indicates strong alignment between system-generated and expert distributions.

Appendix A.3. Algorithm Implementation

Appendix A.3.1. Selection Process Details

To illustrate how the stratagem selection algorithm works in practice, consider a strategic scenario where a company needs to expand its market presence while maintaining existing partnerships. The situation vector x might be:

$$x = [(p_1: 0.15), (p_2: 0.10), (p_3: 0.45), (p_4: 0.20), (p_5: 0.05), (p_6: 0.05)]$$

From our analysis, we have several heuristic distributions including Stratagem 24 ("Use Allies' Resources"):

$$d_{24} = [(p_1: 0.10), (p_2: 0.07), (p_3: 0.61), (p_4: 0.13), (p_5: 0.03), (p_6: 0.06)]$$

Applying cosine similarity:

$$score_{24} = \frac{(0.15 \times 0.10) + (0.10 \times 0.07) + \cdots + (0.05 \times 0.06)}{\sqrt{0.15^2 + 0.10^2 + \cdots + 0.05^2} \times \sqrt{0.10^2 + 0.07^2 + \cdots + 0.06^2}}$$
$$= 0.89$$

Similarly, for two other relevant stratagems:
Stratagem 15 ("Lure Into Unfavorable Position"):

$$d_{15} = [(p_1: 0.40), (p_2: 0.15), (p_3: 0.20), (p_4: 0.15), (p_5: 0.05), (p_6: 0.05)]$$

$$score_{15} = 0.76$$

Stratagem 3 ("Kill With Borrowed Knife"):

$$d_3 = [(p_1: 0.30), (p_2: 0.10), (p_3: 0.35), (p_4: 0.15), (p_5: 0.05), (p_6: 0.05)]$$

$$score_3 = 0.82$$

With threshold $\theta = 0.75$, the algorithm returns:

```
[(24, 0.89), (3, 0.82), (15, 0.76)]
```

Appendix A.4. System Architecture Details

Appendix A.4.1. Workflow Definitions

The system uses JSON state definitions to manage conversation flow and data validation. Below, we provide some illustrative excerpts from those definitions:

```
{
 "state": "parameter_collection",
 "transitions": {
   "complete": "framework_selection",
   "incomplete": "parameter_prompt"
 },
 "validation": {
   "required_fields": ["offensive_strength",
                        "defensive_strength"],
   "value_bounds": {
     "min": 0,
     "max": 5
   }
 }
}
```

Appendix A.4.2. Conversation Management

The conversation manager implements a state machine that governs user interactions:

```
{
 "states": {
   "initial": {
     "type": "input_collection",
     "required_params": ["scenario_type", "actor_count"],
     "next": "actor_details"
   },
   "actor_details": {
     "type": "parameter_collection",
     "validation": "validate_actor_params",
     "next": "framework_selection"
   },
   "framework_selection": {
     "type": "single_choice",
     "options": ["6C", "SWOT", "Porter"],
     "next": "analysis"
   }
 }
}
```

Appendix A.4.3. Component Interactions

Inter-component communication follows standardized protocols:

```
{
 "request": {
   "type": "semantic_analysis",
   "content": {
     "text": "strategic situation description",
     "framework": "6C",
     "parameters": {
       "offensive_strength": 4.2,
       "defensive_strength": 3.8
     }
   }
 },
 "response": {
   "vectors": {
     "situation": [0.8, 0.6, 0.7],
     "params": {
       "p1": [0.9, 0.5, 0.4]
     }
   },
   "similarities": {
     "stratagem_24": 0.85
   }
 }
}
```

Appendix A.4.4. LLM Integration

Template examples for LLM-generated content:

```json
{
 "template_type": "strategy_explanation",
 "components": {
   "context": {
     "framework": "{{framework_name}}",
     "key_parameters": "{{param_list}}",
     "scores": "{{similarity_scores}}"
   },
   "structure": {
     "introduction": "Based on {{framework_name}} analysis...",
     "rationale": "The recommended strategy aligns with...",
     "implementation": "Key steps include..."
   },
   "constraints": {
     "max_length": 500,
     "required_sections": ["context", "rationale", "steps"],
     "style": "professional"
   }
 }
}
```

Appendix A.5. Validation Analysis

The examples in this appendix are intended primarily for *illustrative* purposes, showing how our validation procedures work in practice. In the real system:

- The vectors representing parameters and heuristics generally have much higher dimensionality (e.g., hundreds of embedding components) rather than the 3D vectors shown here.
- Expert reviews may involve more participants (e.g., five or more) to ensure broader consensus, whereas we show a three-expert sample below for didactic clarity.

Despite these simplifications, the overall process remains the same at larger scales.

Appendix A.5.1. Perturbation Analysis

Our perturbation analysis evaluated system robustness by introducing controlled variations in input text. For example, considering Stratagem 24 ("Use Allies' Resources"), we tested these variations:

Original text:

```
"Use Allies' Resources"
```

Variations tested:

```
1. "Leverage Partnership Assets"
2. "Utilize Collaborative Resources"
3. "Deploy Allied Capabilities"
```

Results for parameter p_3 (Relational Capacity):

$$
\begin{aligned}
\text{Original:} \quad & 0.61 \\
\text{Variation 1:} \quad & 0.58 \\
\text{Variation 2:} \quad & 0.63 \\
\text{Variation 3:} \quad & 0.59
\end{aligned}
$$

The stable distribution patterns across variants (standard deviation < 0.03) demonstrate robust semantic mapping.

Appendix A.5.2. Cross-Validation Results

To ensure stability across multiple embedding models, we compare the system-generated distributions of parameter importance using BERT-base, RoBERTa, and Sentence-BERT. Although these examples focus on Stratagem 24 in a simplified format, the same approach extends to higher-dimensional embeddings and additional heuristics or frameworks.

Model Comparison

Below are sample distributions for Stratagem 24 from each embedding model:
BERT-base:
$$[(p_1: 0.10), (p_2: 0.07), (p_3: 0.61), (p_4: 0.13), (p_5: 0.03), (p_6: 0.06)]$$

RoBERTa:

$$[(p_1: 0.11), (p_2: 0.08), (p_3: 0.58), (p_4: 0.14), (p_5: 0.03), (p_6: 0.06)]$$

Sentence-BERT:

$$[(p_1: 0.09), (p_2: 0.07), (p_3: 0.63), (p_4: 0.12), (p_5: 0.04), (p_6: 0.05)]$$

Expert Ratings

For demonstration, we show three expert ratings of parameter p_3's importance in Stratagem 24:

- Expert 1: 0.55
- Expert 2: 0.60
- Expert 3: 0.58

Although our main study employs five experts for deeper validation, these three ratings illustrate the calculation process. With weighting parameters $\alpha = 0.3$, $\beta = 0.3$, and $\gamma = 0.4$, the final confidence score calculation is:

$$
\begin{aligned}
c_{3,24} &= 0.3 \cdot 0.92 + 0.3 \cdot 0.88 + 0.4 \cdot 0.94 \\
&= 0.916
\end{aligned}
$$

indicating strong agreement between the system's discovered distribution and the experts' judgments.

Appendix A.5.3. Framework-Specific Implementation

Below is a short example illustrating how we parse a stratagem's text into vectors and then apply framework-specific adjustments. In real usage, these vectors are higher-dimensional, and additional domain refinements may be employed.

Step 1: Text Preprocessing

```
Input: "Use Allies' Resources"
Tokens: ["use", "allies", "resources"]
```

Step 2: Vector Representation

$$v_{\text{use}} = [0.2, 0.3, 0.1], \quad v_{\text{allies}} = [0.4, 0.6, 0.5], \quad v_{\text{resources}} = [0.3, 0.4, 0.2]$$
$$v_{\text{total}} = [0.9, 1.3, 0.8] \quad (\text{combined})$$

Step 3: Framework-Specific Adjustments

6C Framework. Increase weights for relational/offensive terms; *SWOT Analysis.* Distinguish internal vs. external factors; *Porter's Five Forces.* Emphasize industry-structure terminology.

These illustrative factors ensure framework-specific nuances are captured. Despite the minimal 3D example here, the actual system incorporates substantially larger embedding vectors and more advanced weighting logic to handle more complex strategic texts.

References

1. Mintzberg, H. *Strategy Safari: A Guided Tour Through the Wilds of Strategic Management*; Simon and Schuster: New York, NY, 2005.

2. Hirt, M.; Willmott, P. Strategic principles for competing in the digital age. McKinsey & Company article, 2014. https://www.mckinsey.com/business-functions/strategy-and-corporate-finance/our-insights/strategic-principles-for-competing-in-the-digital-age.

3. Taylor, P. *The Thirty-Six Stratagems: A modern interpretation of a strategy classic*; Infinite Success, Infinite Ideas, 2013.

4. von Senger, H.; Gubitz, M. *The Book of Stratagems: Tactics for Triumph and Survival*; Viking, 1991.

5. Eisenhardt, K.M.; Sull, D.N. Strategy as Simple Rules. *Harvard Business Review* **2001**, *79*, 106–116.

6. Helms, M.; Nixon, J. Exploring SWOT analysis – where are we now? : A review of academic research from the last decade. *Journal of Strategy and Management - J Econ Manag Strat* **2010**, *3*, 215–251. https://doi.org/10.1108/17554251011064837.

7. Porter, M.E. The Five Competitive Forces That Shape Strategy. *Harvard Business Review* **2008**, *86*, 78–93.

8. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.; Polosukhin, I. Attention is All you Need. In Proceedings of the Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA; Guyon, I.; von Luxburg, U.; Bengio, S.; Wallach, H.M.; Fergus, R.; Vishwanathan, S.V.N.; Garnett, R., Eds., 2017, pp. 5998–6008.

9. Devlin, J.; Chang, M.; Lee, K.; Toutanova, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Proceedings of the Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers); Burstein, J.; Doran, C.; Solorio, T., Eds. Association for Computational Linguistics, 2019, pp. 4171–4186. https://doi.org/10.18653/V1/N19-1423.

10. Reimers, N.; Gurevych, I. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In Proceedings of the Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019; Inui, K.; Jiang, J.; Ng, V.; Wan, X., Eds. Association for Computational Linguistics, 2019, pp. 3980–3990. https://doi.org/10.18653/V1/D19-1410.

11. Blei, D.M. *Probabilistic topic models*; Vol. 55, ACM New York, NY, USA, 2012; pp. 77–84. https://doi.org/10.1145/2133806.2133826.

12. Di Pilla, P.; Pareschi, R.; Salzano, F.; Zappone, F. Listening to what the system tells us: Innovative auditing for distributed systems. *Frontiers in Computer Science* **2023**, *4*. https://doi.org/10.3389/fcomp.2022.1020946.

13. Cover, T.M.; Thomas, J.A. *Elements of Information Theory*; John Wiley & Sons: Hoboken, NJ, USA, 2006. First published 7 April 2005, https://doi.org/10.1002/047174882X.

14. Liang, T.P. Recommendation systems for decision support: An editorial introduction. *Decision Support Systems* **2008**, *45*, 385–386. Special Issue Clusters, https://doi.org/https://doi.org/10.1016/j.dss.2007.05.003.

15. Kulkarni, S.; Rodd, S.F. Context Aware Recommendation Systems: A review of the state of the art techniques. *Computer Science Review* **2020**, *37*, 100255. https://doi.org/https://doi.org/10.1016/j.cosrev.2020.100255.

16. Petruzzelli, A.; Musto, C.; Laraspata, L.; Rinaldi, I.; de Gemmis, M.; Lops, P.; Semeraro, G. Instructing and Prompting Large Language Models for Explainable Cross-domain Recommendations. In Proceedings of the Proceedings of the 18th ACM Conference on Recommender Systems, RecSys 2024, Bari, Italy, October 14-18, 2024; Noia, T.D.; Lops, P.; Joachims, T.; Verbert, K.; Castells, P.; Dong, Z.; London, B., Eds. ACM, 2024, pp. 298–308. https://doi.org/10.1145/3640457.3688137.

17. Yang, T.; Chen, L. Unleashing the Retrieval Potential of Large Language Models in Conversational Recommender Systems. In Proceedings of the Proceedings of the 18th ACM Conference on Recommender Systems, RecSys 2024, Bari, Italy, October 14-18, 2024; Noia, T.D.; Lops, P.; Joachims, T.; Verbert, K.; Castells, P.; Dong, Z.; London, B., Eds. ACM, 2024, pp. 43–52. https://doi.org/10.1145/3640457.3688146.

18. Ross, R.G. *Principles of the Business Rule Approach*; Addison-Wesley Professional: Boston, MA, 2003.

19. Schoemaker, P., Scenario Planning; 2016; pp. 1–9. https://doi.org/10.1057/978-1-349-94848-2_652-1.

20. Dean, M. Scenario Planning: A Literature Review. Technical report, ResearchGate Preprint, 2019. Available online at https://doi.org/10.13140/RG.2.2.12629.24802, https://doi.org/10.13140/RG.2.2.12629.24802.

21. Challco, G.; Bittencourt, I.; Reis, M.; Santos, J.; Isotani, S., Gamiflow: Towards a Flow Theory-Based Gamification Framework for Learning Scenarios; 2023; pp. 415–421. https://doi.org/10.1007/978-3-031-36336-8_65.

22. Pareschi, R. Beyond Human and Machine: An Architecture and Methodology Guideline for Centaurian Design. *Sci* **2024**, *6*. https://doi.org/10.3390/sci6040071.

23. Saghafian, S.; Idan, L. Effective generative AI: The human-algorithm centaur. *arXiv preprint arXiv:2406.10942* **2024**.

24. Pareschi, R. *Centaur Art: The Future of Art in the Age of Generative AI*, 1 ed.; Springer, Cham, 2024; pp. X, 91. https://doi.org/10.1007/978-3-031-69063-1.

25. Pareschi, R. Abductive reasoning with the GPT-4 language model: Case studies from criminal investigation, medical practice, scientific research. *Sistemi intelligenti, Rivista quadrimestrale di scienze cognitive e di intelligenza artificiale* **2023**, pp. 435–444. https://doi.org/10.1422/108139.

26. Sull, D.; Eisenhardt, K.M. *Simple Rules: How to Thrive in a Complex World*; Houghton Mifflin Harcourt, 2015.