A Methodology for System Analysis Based on Semantic Digital Twins of Problem Domains

Dmitry Lande[1], Leonard Strashnoy[2]
[1] ORCID: 0000-0003-3945-1178
National Technical University of Ukraine – Igor Sikorsky Kyiv Polytechnic Institute
[2] ORCID: 0009-0008-5575-0286
University of California, Los Angeles (UCLA)

Abstract

This paper proposes a unified methodology for semantic system analysis that integrates Retrieval-Augmented Generation (RAG), semantic digital twins, Agentic AI, and bidirectional semantic graph construction into a common analytical framework. Unlike conventional large language model (LLM) applications that directly generate answers from prompts, the proposed methodology first constructs a semantic digital twin of the problem domain from dynamically retrieved evidence and then performs system analysis over this intermediate model. The methodology combines forward semantic expansion from the current state with backward semantic expansion from the target or forecast state, producing an integrated bidirectional semantic digital twin that serves as the basis for scenario extraction, multicriteria evaluation, ranking, and explainable recommendation generation.

A fundamental contribution of the proposed approach is the unification of decision support and forecasting within a single methodological framework. In decision-support problems, the target state is explicitly defined by the user, whereas in forecasting the target state is initially unknown and is represented by a set of alternative future hypotheses automatically generated by the LLM. Each hypothesis is transformed into its own bidirectional semantic digital twin, after which the same procedures of scenario extraction, evaluation, ranking, and explanation are applied. Consequently, decision support becomes a special case of the more general forecasting methodology corresponding to a single target state.

The methodology is implemented through an agent-based execution framework in which specialized agents cooperatively construct, refine, analyze, and interpret the semantic digital twin using iterative interactions with Retrieval-Augmented Generation. Experimental case studies demonstrate that the proposed approach provides interpretable, evidence-grounded, and reproducible solutions for both decision support and forecasting while preserving a common system-analysis workflow. The proposed methodology defines a new class of Agentic Semantic System Analysis systems, in which large language models function primarily as instruments for constructing and analyzing semantic digital twins rather than merely generating textual responses.

Keywords: System analysis; semantic digital twin; Agentic AI; Retrieval-Augmented Generation; decision support; forecasting; semantic networks; explainable AI

I. Introduction

System analysis has remained one of the fundamental methodological disciplines for investigating complex socio-economic, technical, and organizational systems for several decades. Its primary objective is to construct formal models of a problem domain, identify relationships among its components, analyze alternative development scenarios, and support informed decision-making under conditions of uncertainty. The classical toolkit of system analysis includes structural analysis, operations research, decision theory, scenario analysis, system dynamics, cognitive modeling, network analysis, multi-criteria optimization, statistical forecasting, and numerous other approaches, each of which is effective for a particular class of problems.

The rapid development of information technologies has significantly expanded the practical applicability of system analysis. The availability of large volumes of digital data, the emergence of open information resources, and advances in machine learning and computational technologies have enabled the development of modern Decision Support Systems (DSS), which combine mathematical models, expert knowledge, and automated information processing [1]. Nevertheless, most existing systems rely either on predefined domain models or on specialized algorithms designed for specific classes of problems [2]. Consequently, the construction of domain models, the generation of development scenarios, and the interpretation of analytical results still depend heavily on human experts.

In recent years, this situation has changed considerably owing to the rapid development of Large Language Models (LLMs) [3]. These models have demonstrated remarkable capabilities in text analysis, information summarization, natural language generation, and knowledge processing. The emergence of Retrieval-Augmented Generation (RAG) [4] has further enhanced these capabilities by enriching the internal knowledge of language models with up-to-date information retrieved from external information resources, thereby substantially improving the reliability of generated results. At the same time, multi-agent architectures have attracted increasing attention, enabling specialized agents to cooperate at different stages of the analytical process within a unified workflow.

Despite these advances, most existing studies employ LLMs primarily as universal answer generators or as tools for automating individual analytical procedures. Considerable research has focused on integrating LLMs with knowledge bases, developing agent-based systems, exploiting knowledge graphs, Retrieval-Augmented Generation, and Chain-of-Thought reasoning [5]. In the majority of these approaches, however, the language model directly produces the final response to a user's query. The processes of constructing a domain model, analyzing alternative scenarios, evaluating their consistency, and ranking possible solutions remain implicit within the internal reasoning of the language model and therefore cannot be formally analyzed, verified, or explained.

In classical system analysis, the situation is fundamentally different. Decision-making begins with constructing an explicit model of the system under investigation, identifying its structure, analyzing causal relationships, generating alternative development scenarios, comparing them according to predefined criteria, and selecting the most appropriate alternative [6]. Thus, model construction constitutes the central procedure of system analysis, whereas the final recommendation represents only the last stage of the analytical process.

This observation motivates the need to integrate the capabilities of modern large language models with the methodology of system analysis. The objective is not to replace classical analytical methods with LLMs, but rather to develop a new methodological framework in which language models are employed as instruments for constructing semantic models of problem domains that can subsequently be used for scenario analysis, alternative evaluation, and forecasting of possible future developments [7].

In this paper, we propose a methodology for semantic system analysis that integrates Retrieval-Augmented Generation, semantic digital twins, bidirectional semantic expansion, and agentic execution of structured prompts into a unified framework for constructing and analyzing alternative development scenarios of a problem domain. Unlike the conventional use of large language models, where the output is a directly generated textual response, the proposed methodology considers the semantic digital twin of the problem domain as the central analytical object. The semantic digital twin is automatically constructed from a relevant information context obtained through Retrieval-Augmented Generation. Its informational foundation is provided by a Retrieval-Augmented Generation framework integrated with the **InfoStream** open-source information monitoring system, enabling the use of current and verifiable information throughout the analytical process. The constructed semantic digital twin is iteratively refined through forward and backward semantic expansion, producing an integrated representation of the problem domain that serves as the basis for automatic scenario extraction, multi-criteria evaluation, scenario ranking, and the generation of explainable recommendations.

The proposed methodology is generic and can be applied to a broad class of system analysis problems. In this work, its applicability is demonstrated using two representative problem classes: decision support and forecasting. In decision-support problems, the target state is explicitly specified by the user, and the semantic digital twin is employed to identify the most effective scenarios leading to that state. In forecasting problems, a large language model first automatically generates a set of alternative future states representing forecast hypotheses. Subsequently, an independent bidirectional semantic digital twin is constructed for each hypothesis, followed by scenario analysis, probability estimation, and ranking of the most plausible future developments. Consequently, decision support and forecasting are treated as two special cases of a unified semantic system analysis methodology that differ only in the way the set of target states is defined.

The objective of this study is to develop a methodology for semantic system analysis based on semantic digital twins of problem domains, large language models, Retrieval-Augmented Generation, and agentic technologies, enabling the construction of explainable and reproducible scenarios for decision support and forecasting using current open information resources.

The scientific contribution of this work is the development of a new methodology for semantic system analysis that integrates Retrieval-Augmented Generation, semantic digital twins, bidirectional semantic expansion, and agentic execution of structured prompts into a unified framework for constructing, analyzing, and interpreting alternative scenarios. Unlike existing approaches, the central analytical object is neither a textual query, nor a collection of documents, nor a knowledge graph, but a semantic digital twin of the problem domain automatically constructed from the current information space. Furthermore, the paper demonstrates for the first time that decision support and forecasting can be formulated as two special cases of a unified semantic system analysis methodology that differ only in the mechanism used to define the set of target states.

2. Methodology of Semantic System Analysis

The proposed methodology is based on integrating the classical principles of system analysis with the capabilities of Large Language Models (LLMs) and Retrieval-Augmented Generation (RAG). Its distinguishing feature is that the object of analysis is not an individual user query but the information space of the problem domain represented by a collection of interconnected documents. Unlike conventional LLM-based approaches, in which a response is generated immediately after receiving a prompt, the proposed methodology first constructs a semantic model of the problem domain. All subsequent analytical procedures—including scenario generation, alternative evaluation, ranking, and recommendation generation—are then performed within this intermediate model.

The methodological foundation of the proposed approach follows the classical concept of system analysis, according to which every decision should be based on an explicit model of the system under investigation. Model construction precedes the analysis of alternatives, whereas recommendation generation constitutes only the final stage of the analytical process. In the proposed methodology, this role is performed by a semantic digital twin automatically constructed from a relevant information context obtained through Retrieval-Augmented Generation.

The informational foundation of the methodology is provided by the InfoStream open-source information monitoring system, which continuously collects, normalizes, indexes, and stores documents originating from news agencies, web resources, official portals, social media, and other open information sources. The architecture of InfoStream combines classical information retrieval technologies with modern analytical components, providing high relevance, timeliness, reliability, and reproducibility of information retrieval results. An important feature of the system is its continuously maintained retrospective news archive spanning more than thirty years, which enables not only real-time analysis but also long-term investigations of the evolution of information processes.

Unlike most existing Retrieval-Augmented Generation implementations, where the retrieval component merely selects several relevant documents for answer generation, the proposed methodology treats Retrieval-Augmented Generation as an integral part of the system analysis process. The information context is constructed using the InfoReS full-text search engine, which supports a formal query language, contextual operators, morphological processing, automatic detection of semantic duplicates, and linguistic-statistical document ranking. Consequently, the language model receives a cleaned, relevant, and well-structured document collection, significantly reducing information noise and the risk of generating unsupported statements.

Within the proposed methodology, Retrieval-Augmented Generation performs a substantially broader function than conventional prompt augmentation. The retrieved document collection is regarded as an empirical representation of the current state of the problem domain from which the language model constructs its semantic digital twin. Consequently, Retrieval-Augmented Generation provides not only access to current knowledge but also the informational basis for automatically constructing the structural model of the investigated system.

Subsequent analysis is performed not on the document collection itself but on the constructed semantic digital twin. Its nodes correspond to concepts of the problem domain, whereas its directed edges represent semantic, causal, temporal, functional, and other meaningful relationships between these concepts. Unlike traditional knowledge graphs, the proposed semantic digital twin is not derived from a predefined ontology. Instead, it is generated dynamically according to the current information context and the specific analytical task.

The next stage of the methodology consists of constructing two complementary semantic subnetworks. The first is generated through forward semantic expansion originating from the initial system state and represents possible directions of system evolution. The second is generated through backward semantic expansion starting from the target or forecast state and represents possible prerequisite conditions leading to that state. The integration of these subnetworks identifies their semantic intersection, which contains the most plausible transition scenarios between the initial and target states. This stage fundamentally distinguishes the proposed methodology from conventional LLM-based response generation.

The resulting integrated semantic digital twin serves as the foundation for subsequent scenario analysis. Alternative chains of interconnected concepts are automatically extracted from the graph and interpreted as possible development scenarios of the investigated situation. For each scenario, integral evaluation measures are computed by considering semantic consistency, evidence provided by Retrieval-Augmented Generation, structural connectivity of the semantic digital twin, and the results of multi-agent assessment. After scenario ranking, the language model generates natural-language explanations of the obtained results, thereby providing explainable recommendations.

The proposed methodology is sufficiently general to support a broad spectrum of system analysis problems. In this work, two representative implementations are considered. In decision-support problems, the target state is explicitly specified by the user, and the semantic digital twin is employed to identify the most effective transition scenarios leading to that state. In forecasting problems, a large language model first automatically generates a set of alternative future states. Subsequently, an independent bidirectional semantic digital twin is constructed for each forecast hypothesis, followed by scenario analysis, probability estimation, and ranking of alternative future developments. Therefore, both decision support and forecasting rely on the same methodological framework and differ only in the procedure used to define the target state.

3. Formal Model of the Semantic System Analysis Methodology

The proposed methodology considers the system analysis process as a sequence of interconnected information transformations, beginning with an unstructured information space and culminating in the generation of explainable recommendations or forecast scenarios. Unlike the conventional use of Large Language Models (LLMs), where answer generation constitutes the only processing stage, the proposed approach employs the LLM to construct an intermediate semantic model that becomes the central object of subsequent analysis.

Formally, the methodology can be represented by the following operator model

$$M=\langle R,\Phi,\Theta,\Pi,\Gamma,\Omega\rangle,$$

where

- R – Retrieval-Augmented Generation operator;
- Φ – semantic graph construction operator;
- Θ – bidirectional semantic subgraph construction operator;
- Π – scenario generation operator;
- Γ – scenario ranking operator;
- Ω – result interpretation operator.

Accordingly, the overall system analysis process is described by the following composition of operators

$$D \xrightarrow{R} C \xrightarrow{\Phi} G \xrightarrow{\Theta} G^{*} \xrightarrow{\Pi} P \xrightarrow{\Gamma} P^{*} \xrightarrow{\Omega} A \text{ where}$$

- D – the set of information resources;
- C – the relevant information context;
- G – the initial semantic graph;
- $G*$ – the integrated bidirectional semantic graph;
- P – the set of extracted scenarios;
- $P*$ – the ranked set of scenarios;

- A – an explainable recommendation or forecast.

Unlike conventional Retrieval-Augmented Generation architectures, where the retrieval stage terminates with transferring the retrieved textual context directly to a large language model, Retrieval-Augmented Generation in the proposed methodology constitutes only the initial stage of constructing the problem-domain model. The central information object becomes the integrated semantic graph G^*, upon which all subsequent system analysis procedures are performed.

4. Retrieval-Augmented Generation as the Information Foundation of System Analysis

Unlike most contemporary Retrieval-Augmented Generation (RAG) implementations, where the retrieval module is primarily used to obtain a limited set of documents required for answer generation, the proposed methodology considers Retrieval-Augmented Generation as the informational foundation of system analysis. Its role extends beyond supplying a large language model with relevant contextual information. Instead, RAG is employed to construct a representative model of the information environment that provides a sufficient basis for building a semantic digital twin of the problem domain.

The information infrastructure supporting Retrieval-Augmented Generation is provided by the InfoStream content monitoring system (https://infostream.ua), which performs continuous collection, normalization, indexing, and long-term storage of open information resources. The architecture of the system combines the classical InfoReS full-text search engine, automatic semantic duplicate detection, linguistic-statistical document processing, a long-term retrospective news archive, and modern mechanisms for integrating large language models. As a result, the generated document collections simultaneously satisfy the requirements of relevance, completeness, timeliness, and reliability.

The Retrieval-Augmented Generation procedure consists of several consecutive stages. First, a formalized information query is constructed, which may include logical operators, contextual constraints, and temporal conditions. The InfoReS search engine retrieves relevant documents, performs their morphological normalization, removes semantic duplicates, and produces a structured document collection. The retrieved documents are subsequently ranked according to topical relevance, informational significance, and temporal freshness. This ranked document collection is then supplied as input to the large language model.

Unlike conventional RAG architectures, where the retrieved corpus serves solely as contextual information for response generation, the proposed methodology interprets the document collection as an empirical representation of the current state of the investigated system. Consequently, Retrieval-Augmented Generation is no longer regarded as an auxiliary component of the LLM pipeline but becomes the first stage in constructing the semantic model of the problem domain.

The proposed use of Retrieval-Augmented Generation provides several important advantages. First, the semantic digital twin is constructed not from the internal parameters of the language model but from current information obtained from open information sources. Second, every concept and every semantic relationship within the model can be traced back to specific source documents, thereby ensuring explainability, transparency, and verifiability of the analytical results. Third, the retrospective archive maintained by InfoStream makes it possible to analyze not only the current state of the information environment but also its long-term evolution, providing an essential foundation for forecasting future developments.

5. Semantic Digital Twin as a Model of the Problem Domain

A semantic digital twin of a problem domain is defined as a dynamic directed semantic graph automatically constructed from the current information context provided by Retrieval-Augmented Generation. It represents the current state of the problem domain, captures the structure of semantic relationships among concepts, and provides the foundation for scenario analysis, decision support, and forecasting.

The central element of the proposed methodology is a semantic graph constructed from the information context obtained through Retrieval-Augmented Generation. Unlike the traditional use of semantic networks as knowledge representation structures or concept graphs, the proposed methodology assigns a substantially broader role to this graph. It is regarded as a semantic digital twin of the problem domain, i.e., a dynamic model representing its current informational state, the structure of semantic relationships among concepts, and possible directions of its future evolution.

Traditionally, the concept of a digital twin refers to a virtual representation of a physical object or technical system that is continuously synchronized with its real-world counterpart. In the proposed methodology, the same principle is transferred to information processes. The object of modeling is not a physical system but its informational representation constructed from continuously acquired open-source information collected by the InfoStream monitoring system. Consequently, the semantic digital twin is not a static repository of knowledge but a continuously evolving model that is automatically updated according to changes in the information environment, thereby reflecting the current state of the problem domain.

Construction of the semantic digital twin begins immediately after the completion of the Retrieval-Augmented Generation stage. At this point, the large language model receives not a single document but a representative corpus of relevant information resources that has been previously selected, filtered, and structured by InfoStream. By analyzing this corpus, the language model identifies concepts, establishes semantic, causal, temporal, functional, and associative relationships among them, and constructs a directed semantic graph representing the current state of the problem domain.

Formally, the semantic digital twin is defined as the directed graph

$$G = (V, E, W),$$

where

- $V = \{v_1, v_2, ..., v_n\}$ denotes the set of concepts representing the problem domain;
- $E \subseteq V \times V$ denotes the set of semantic relationships between concepts;
- $W : E \rightarrow R$ is the edge-weight function characterizing the strength or significance of each semantic relationship.

Unlike classical knowledge graphs, the graph vertices are not generated from a predefined ontology. Instead, they are determined dynamically by the current information context produced through Retrieval-Augmented Generation. This enables the model to adapt automatically to a specific problem domain and to incorporate newly emerging concepts, events, organizations, technologies, and other informational entities without requiring manual updates of a knowledge base.

The semantic relationships between graph vertices are likewise dynamic rather than predefined. They are established by the language model through analysis of the retrieved document corpus and represent various forms of interaction between concepts. Depending on the analytical objective, these relationships may correspond to causal dependencies, temporal transitions between system states, functional relationships, structural associations, semantic similarity, or other forms of semantic interaction. The associated edge weights reflect the degree of evidential support provided by Retrieval-Augmented Generation and may be further refined during the multi-agent analysis process.

An important characteristic of the proposed methodology is that the semantic digital twin is **not** considered the final output of the large language model. Instead, it serves as an intermediate analytical model upon which all subsequent system analysis procedures are performed. The structure of the semantic digital twin determines the possibility of identifying alternative development scenarios, exploring multiple transition paths toward target states, analyzing causal relationships, and generating explainable recommendations.

In essence, the semantic digital twin fulfills the same role that mathematical, cognitive, or simulation models traditionally play in classical system analysis. Unlike these conventional models, however, it is generated automatically from the current information environment, allowing rapid adaptation to changes in the problem domain and enabling the analysis of systems whose structure is incomplete, weakly formalized, or continuously evolving.

Consequently, the introduction of the semantic digital twin fundamentally changes the role of the large language model within the system analysis process. Rather than functioning as a generator of ready-made answers, the LLM becomes an instrument for automatically constructing a formal model of the problem domain. This model then becomes the central object of subsequent analysis, supporting the discovery of alternative scenarios, their evaluation and ranking, forecasting of possible developments, and the generation of explainable recommendations.

6. Bidirectional Construction of the Semantic Digital Twin

One of the fundamental characteristics of the proposed methodology is the abandonment of the traditional approach in which a single semantic graph is constructed first and paths between predefined vertices are subsequently identified. In classical graph-based methods, the graph structure is assumed to be known or constructed in advance, whereas the analytical task consists of finding optimal paths, causal chains, or other structural relationships within an already existing graph.

Definition 2 (Bidirectional Construction of the Semantic Digital Twin). *Bidirectional construction of the semantic digital twin* is defined as the process of independently performing forward and backward semantic expansion from the initial and target system states, followed by the integration of the resulting semantic subnetworks into a unified semantic digital twin of the problem domain.

The proposed methodology follows a fundamentally different principle. The semantic digital twin does not exist prior to the analysis but is constructed dynamically during the problem-solving process itself. Its construction proceeds simultaneously in two directions: from the current state of the investigated system and from the target state, which may either be explicitly specified by the user or automatically generated during the forecasting process. Consequently, model construction and solution search are performed as a single integrated analytical process rather than as two independent stages.

Let S_0 denote the initial state of the system and S_g the target state.

In decision-support problems, the target state S_g is explicitly specified by the user as the desired objective of the system. In forecasting problems, the set of possible target states is generated automatically by a Large Language Model on the basis of the current information context.

The construction of the semantic digital twin consists of two independent procedures.

The first stage performs forward semantic expansion

$$S_0 \xrightarrow{\Phi_f} G_f,$$

where $G_f = (V_f, E_f)$ denotes the semantic graph constructed through the successive expansion of the conceptual neighborhood of the initial state. At each iteration, the Large Language Model, supported by Retrieval-Augmented Generation, identifies concepts directly related to the already constructed portion of the graph and subsequently introduces new vertices and semantic relationships. The process continues until a predefined analysis depth is reached or one of the stopping criteria is satisfied.

In parallel, an analogous procedure is carried out in the reverse direction,

$$S_g \xrightarrow{\Phi_b} G_b,$$

where $G_b = (V_b, E_b)$ represents the semantic graph describing the set of possible prerequisite conditions for reaching the target state.

Unlike forward semantic expansion, the language model does not predict future developments during this stage. Instead, it reconstructs the possible informational, functional, and causal conditions that may lead to the specified target state. Conceptually, the LLM addresses the following question:

The two expansion procedures are completely independent and employ separate Retrieval-Augmented Generation cycles. Consequently, every newly generated graph vertex may be associated with its own information context, allowing its semantic meaning, relationships, and role within the semantic digital twin to be continuously refined. Therefore, the semantic digital twin is constructed not only through the internal reasoning capabilities of the language model but also through repeated interaction with the continuously evolving information space.

After both expansion procedures have been completed, the resulting semantic graphs are integrated,

$$G^* = G_f \cup G_b.$$

The graph integration process includes concept normalization, elimination of duplicate vertices, aggregation of semantic relationships, and reconciliation of edge weights. As a result, a unified semantic digital twin of the problem domain is obtained, simultaneously representing both the possible directions of system evolution and the prerequisite conditions for achieving the target state.

Thus, the semantic digital twin is not regarded as a predefined graph but as the result of integrating two independently constructed semantic expansions. This distinguishes the proposed methodology from conventional graph-based approaches, in which the graph structure is assumed to exist before the analytical process begins.

The algorithmic principle of Bidirectional Semantic Digital Twin Construction is illustrated in Fig. 1. Unlike conventional graph-based methods, where the graph structure is considered to be predefined, the proposed methodology constructs the semantic digital twin simultaneously in two directions—from the current system state and from the target (or forecast) state. The integration of the resulting semantic subnetworks produces a unified semantic digital twin that serves as the principal analytical object for scenario extraction, evaluation, and ranking of alternative scenarios.

.

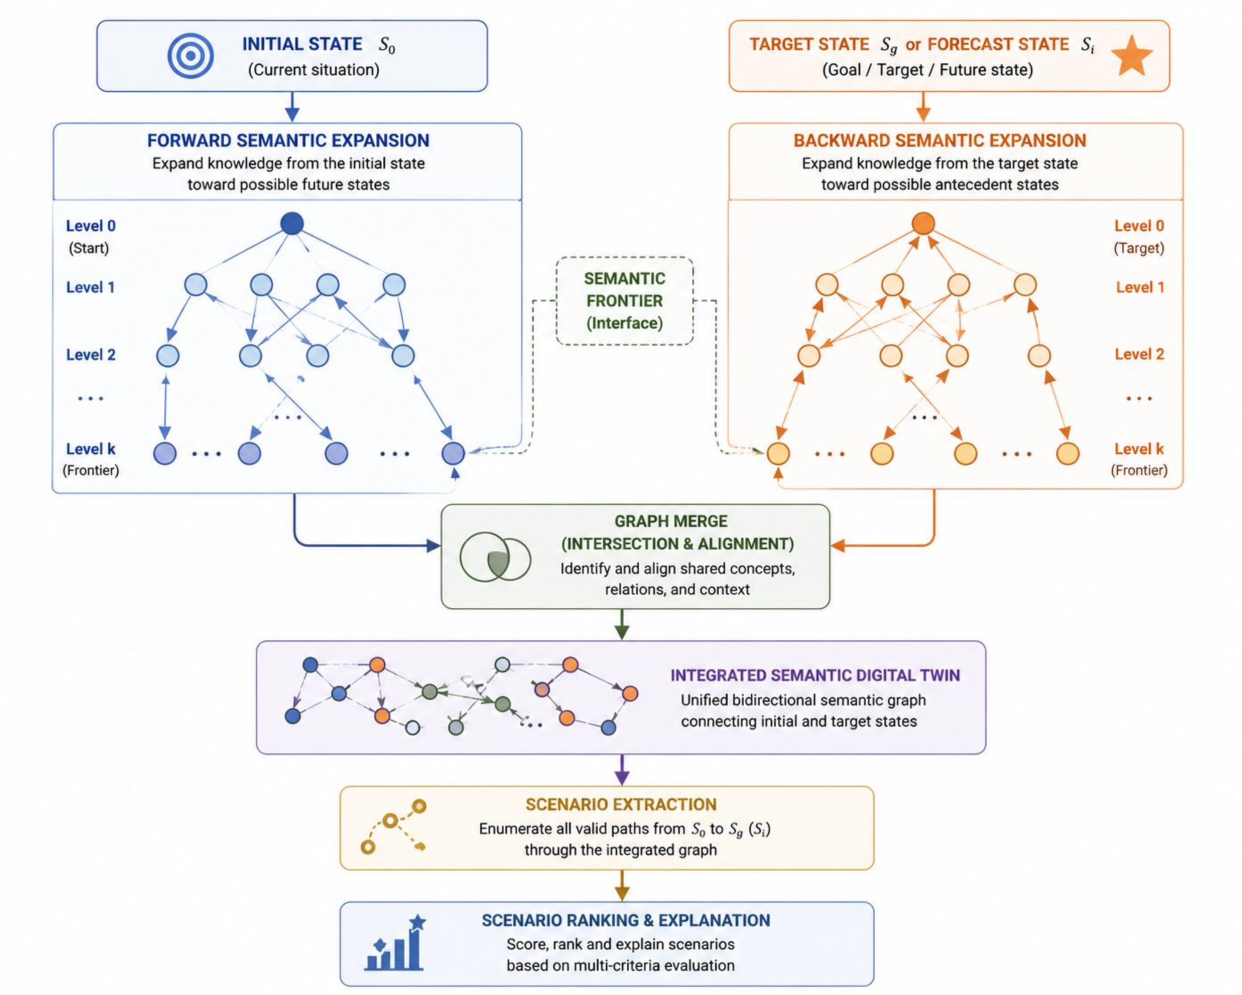

Figure 1. Algorithmic Principle of Bidirectional Semantic Digital Twin Construction

Bidirectional construction of the semantic digital twin consists of two independent semantic expansion processes. Forward semantic expansion describes the possible directions of system evolution starting from the current state, whereas backward semantic expansion reconstructs the set of possible prerequisite conditions required to reach the target or forecast state. The integration of these two semantic subnetworks produces a unified semantic space in which transition scenarios between system states are automatically extracted, evaluated, ranked, and transformed into explainable recommendations.

A particularly important role in the subsequent analysis is played by the intersection of the two semantic subnetworks,

$$V_f \cap V_b.$$

This intersection defines the set of concepts that simultaneously represent both the logical continuation of the initial state and the necessary prerequisites for achieving the target state. In other words, it corresponds to the most semantically consistent part of the digital twin, where the information obtained through forward and backward semantic expansion mutually reinforces and validates itself. This region subsequently serves as the basis for scenario construction, scenario ranking, and the generation of explainable recommendations.

Another important characteristic of the proposed bidirectional semantic digital twin construction is its ability to support continuous model refinement. Unlike classical graph algorithms that operate on a fixed graph structure, the proposed methodology allows the semantic digital twin to be dynamically updated at any stage of the analytical process. The arrival of new information within the open information space automatically modifies the Retrieval-Augmented Generation context, which may subsequently trigger reconstruction of individual fragments of the semantic digital twin. Consequently, the model is not static but continuously evolves together with the problem domain, reflecting its current informational state.

The methodological significance of the proposed approach lies in the integration of model construction and system analysis into a single unified analytical process. In conventional decision-support systems, the domain model is typically developed independently of a specific analytical task. In contrast, the proposed methodology constructs the semantic digital twin dynamically during the analysis itself and immediately orients it toward discovering alternative scenarios that lead to the desired objective or forecast state. Thus, model construction is no longer regarded as a preliminary stage preceding system analysis. Instead, both processes become inseparable components of a unified analytical methodology, in which the semantic digital twin is continuously constructed, refined, and analyzed throughout the entire decision-support or forecasting process.

7. Agentic Execution of System Analysis Procedures

Construction of the semantic digital twin constitutes only the initial stage of the proposed system analysis methodology. Subsequent solution of decision-support or forecasting problems requires the execution of a sequence of interconnected analytical procedures, including expansion of the information context, analysis of the semantic graph structure, generation of alternative scenarios, their evaluation, verification of the consistency of the obtained results, and the generation of explainable recommendations. In conventional decision-support systems, these procedures are performed either by individual algorithms or directly by human experts. In the proposed methodology, they are executed by a multi-agent system built upon Large Language Models.

A fundamental characteristic of the proposed approach is that the agents do not operate as independent intelligent modules producing alternative textual responses. Instead, each agent performs a well-defined system analysis procedure on the shared semantic digital twin of the problem domain. Consequently, the primary object of interaction among agents is not textual communication but a common structured model of the problem domain that is progressively refined throughout the analytical process.

Let $A=\{A_1,A_2,...,A_m\}$ – denote the set of agents participating in the analytical process. Each agent is associated with an operator

$$A_i : G^* \rightarrow G^*,$$

meaning that the agent performs a specific transformation of the semantic digital twin or computes additional characteristics describing individual fragments of the model. The collective operation of the multi-agent system is represented by the composition of agent operators

$$G_0^* \xrightarrow{A_1} G_1^* \xrightarrow{A_2} \ldots \xrightarrow{A_m} G_m^*, \; *,$$

where G_0^* denotes the initial semantic digital twin obtained after integrating the forward and backward semantic graphs, and G_m^* denotes the refined semantic digital twin prepared for subsequent scenario extraction.

The functional specialization of agents is determined by the specific system analysis task, whereas their roles remain conceptually unchanged. The principal analytical procedures include retrieving additional information from open sources, refining semantic relationships between concepts, expanding the semantic digital twin with newly identified vertices, verifying the logical consistency of the constructed graph fragments, evaluating the significance of concepts, generating alternative scenarios, and producing natural-language explanations of the obtained results. Since all agents operate on a shared semantic digital twin of the problem domain, inconsistencies between individual stages of the analytical process are avoided, ensuring the integrity and coherence of decision making.

An important characteristic of the proposed methodology is the iterative nature of agent interaction. After each analytical procedure, the obtained results are incorporated into the semantic digital twin, which subsequently becomes the updated information basis for the next analysis cycle. Whenever insufficient information or inconsistencies are detected in any fragment of the model, the multi-agent system initiates an additional Retrieval-Augmented Generation cycle to construct a refined information context. Consequently, the information space, the semantic digital twin, and the multi-agent system form a closed analytical loop that continues until the semantic model reaches the required level of completeness and internal consistency.

Unlike conventional multi-agent systems, where communication is primarily based on message exchange or independent computational results, the proposed methodology enables agents to interact through a shared semantic digital twin. This shared analytical model provides a common semantic workspace that supports consistent knowledge accumulation, facilitates reuse of intermediate analytical results, and significantly reduces the risk of generating logically inconsistent recommendations.

Another important advantage of the proposed agentic organization is its independence from any particular Large Language Model or agent framework. Individual agents may be implemented using different LLMs, specialized analytical modules, machine learning algorithms, expert systems, or conventional software components. The defining characteristic of an agent is therefore not its implementation technology but its ability to perform a formalized system analysis procedure on a common semantic digital twin of the problem domain.

For this reason, the multi-agent system is regarded not as an independent artificial intelligence technology but as the execution layer of the proposed methodology. It provides the computational environment in which the formalized procedures of semantic system analysis are executed while remaining independent of the specific problem domain, implementation platform, or underlying artificial intelligence technologies

Consequently, the proposed methodology establishes a clear separation between the methodology of system analysis and its computational realization, allowing the analytical procedures to be executed by heterogeneous artificial intelligence technologies while preserving a unified semantic model of the problem domain.

8. Scenario Extraction, Evaluation, and Ranking

Following the construction of the semantic digital twin and its subsequent refinement by the multi-agent system, the central objective of system analysis becomes the identification of possible scenarios describing the evolution of the investigated situation. Within the proposed methodology, a **scenario** constitutes the fundamental unit of analysis because it represents not merely individual semantic relationships between concepts but an ordered sequence of logically connected transitions between states of the problem domain leading either to a specified objective or to one of the possible future states.

Unlike classical graph algorithms, which search for the shortest or otherwise optimal paths within a static graph, the proposed methodology treats a scenario as a meaningful chain of interconnected concepts obtained through the integration of forward and backward semantic expansion. Consequently, a scenario represents not only a topological property of the semantic digital twin but also an information model describing a possible evolution of the investigated situation.

Formally, a scenario is defined as an ordered sequence of concepts

$$p = \left(v_1, v_2, ..., v_k\right),$$

where $v_i \in V$ are vertices of the semantic digital twin, and every pair of adjacent concepts is connected by one or more semantic relationships.

The set of all admissible scenarios is denoted by

$$P = \left\{p_1, p_2, ..., p_N\right\}.$$

In practical applications, the number of possible scenarios may become extremely large, making exhaustive enumeration computationally infeasible. Therefore, after constructing the semantic digital twin, a dedicated scenario extraction procedure is performed to identify only those conceptual chains that satisfy predefined semantic and structural constraints.

Scenario extraction is performed by the multi-agent system through the successive expansion of conceptual chains from the initial system state toward the intersection region of the forward and backward semantic subnetworks. At each iteration, the system evaluates the semantic consistency of each transition, the evidential support provided by Retrieval-Augmented Generation, the logical coherence of the partially constructed scenario, and its relevance to the analytical objective. Whenever multiple continuation alternatives are identified, the multi-agent system branches the construction process, thereby generating a set of alternative development scenarios.

Each extracted scenario is subsequently subjected to comprehensive evaluation. Unlike traditional graph-based methods, where optimality is typically determined by path length or the cumulative weight of graph edges, the proposed methodology employs a multi-criteria evaluation that considers the semantic characteristics of each scenario.

For each scenario p_i, an integral evaluation score $Q(p_i)$ is computed as an aggregation of multiple evaluation criteria,

$$Q(p_i) = F(q_1, q_2, ..., q_m),$$

where F(·) denotes the aggregation function and q_j are the individual evaluation criteria.

The principal evaluation criteria include:

- semantic consistency, characterizing the logical coherence of the scenario and the absence of contradictions between successive transitions;
- evidential support, determined by the quantity, reliability, and temporal relevance of the Retrieval-Augmented Generation documents supporting each stage of the scenario;
- structural connectivity, reflecting the degree to which the scenario is integrated into the semantic digital twin;
- agent-based assessment, obtained through the independent evaluation of the scenario by multiple analytical agents;
- target-state achievement, measuring the completeness with which the analytical objective is accomplished.

The set of evaluation criteria is not fixed and may be adapted according to the requirements of a particular system analysis problem. This flexibility ensures the generality of the proposed methodology and enables its application across a wide range of problem domains.

After all scenarios have been evaluated, they are ranked according to their integral scores,

$$\Gamma : P \rightarrow P^*,$$

where $P^* = (p_1, p_2, ..., p_N)$ denotes the ordered set of scenarios sorted in descending order of their integral evaluation score $Q(p)$.

Unlike conventional approaches to alternative ranking, the proposed methodology does not restrict the analytical outcome to a single optimal solution. Instead, the user is presented with a ranked collection of the most promising scenarios, enabling comparative analysis of alternatives, assessment of associated risks, and selection of the most appropriate solution while taking into account additional external factors that may not have been explicitly represented during construction of the semantic digital twin.

An essential feature of the proposed methodology is the explainability of the obtained results. For each scenario, the multi-agent system automatically generates a natural-language explanation describing the rationale underlying its construction, the key concepts governing the evolution of the situation, the principal causal relationships, and the information sources supporting the proposed scenario. Since every stage of scenario construction is explicitly associated with corresponding fragments of both the semantic digital twin and the Retrieval-Augmented Generation context, the user can trace the complete chain of reasoning from the original source documents to the final recommendation. This provides a high degree of transparency, reproducibility, and confidence in the analytical results.

The proposed mechanism for scenario extraction, evaluation, and ranking is independent of any particular application domain. The only difference between specific analytical tasks lies in the procedure used to define the target state of the semantic digital twin. In decision-support problems, the target state is explicitly specified by the user, whereas in forecasting problems it is generated automatically as a set of alternative future states. In all other respects, the procedures for semantic digital twin construction, scenario extraction, scenario evaluation, and scenario ranking remain identical. This methodological invariance makes it possible to regard the proposed approach as a general methodology for system analysis based on semantic digital twins, applicable to a broad spectrum of decision-support and forecasting problems.

9. Decision Support Method Based on the Semantic Digital Twin

The proposed methodology of semantic system analysis enables the implementation of the complete decision-support cycle, beginning with the construction of the information context and culminating in the generation of explainable recommendations. Unlike conventional Decision Support Systems (DSS), in which information retrieval, domain modeling, alternative evaluation, and recommendation generation are performed by separate and largely independent components, the proposed approach integrates all these stages into a unified process based on the construction and analysis of a semantic digital twin of the problem domain.

The decision-support problem is formally defined by the triple

$$D = \left\langle S_o, S_g, C \right\rangle,$$

where

- S_0 denotes the current state of the problem domain;
- S_g denotes the desired target state;
- C denotes the set of external constraints and decision-making criteria.

Unlike forecasting problems, the target state of the system is explicitly specified by the user. It may represent the desired operating state of the system, the achievement of a strategic objective, or the elimination of an undesirable condition. Consequently, the objective of decision support is not to predict future developments but to identify the most well-founded transition scenarios leading from the current system state to the specified target state.

Once the initial conditions have been defined, the multi-agent system formulates a task-specific Retrieval-Augmented Generation query. Unlike conventional information retrieval, the query is constructed not only from keywords but also by considering the semantics of the current state, the characteristics of the target state, and the constraints specified by the user. As a result, a document corpus is assembled that provides the most representative description of the problem domain for the analytical task under consideration.

Using this information context, the procedure of Bidirectional Semantic Digital Twin Construction is initiated. Forward semantic expansion starts from the current system state and models its possible directions of evolution. Simultaneously, backward semantic expansion is performed from the specified target state, reconstructing the potential prerequisite conditions required for its achievement. The integration of the resulting semantic subnetworks produces a unified semantic digital twin containing the most comprehensive structural representation of the investigated situation.

An important characteristic of the proposed methodology is that construction of the semantic digital twin is not a one-time operation. Following each cycle of agentic analysis, the model may be refined through additional Retrieval-Augmented Generation cycles, expansion of the conceptual space, or revision of individual semantic relationships. Consequently, the semantic digital twin progressively converges toward an increasingly complete representation of the current information environment.

The next stage consists of extracting alternative transition scenarios connecting the initial and target states. Unlike classical graph-search algorithms, the multi-agent system analyzes not only the structural properties of the semantic digital twin but also the semantic characteristics of concepts, the logical progression of events, causal dependencies, and the evidential support associated with individual transitions. The result is a set of alternative scenarios, each representing a distinct strategy for achieving the specified objective.

Once the scenarios have been generated, they undergo a multi-criteria evaluation. At this stage, the multi-agent system no longer constructs new concepts but evaluates the generated scenarios from multiple expert perspectives. Each agent performs its assessment according to its own specialization while operating on the same semantic digital twin as a shared analytical model. This ensures the independence of expert evaluations while preserving a common informational foundation for the entire decision-making process.

Let $A=\{A_1,A_2,...,A_m\}$ denote the set of expert agents. For each scenario p_i, an evaluation vector is constructed, $r(p_i)=(r_{i1},r_{i2},...,r_{im})$, where r_{ij} denotes the evaluation assigned to scenario p_i by the expert agent A_j.

The integral evaluation of the scenario is then computed using the aggregation function $Q(p_i)=F(r(p_i))$, where the aggregation operator F may implement different collective decision-making mechanisms, including weighted averaging, multi-criteria optimization, expert voting, or other aggregation methods.

The use of a multi-agent architecture enables the implementation of the collective expert assessment principle. Instead of relying on a single response generated by a large language model, the proposed methodology produces a set of independent expert evaluations that are subsequently integrated during the scenario-ranking stage. This approach increases the robustness of the analytical results against occasional errors of individual language models, incorporates multiple analytical perspectives, and significantly improves the explainability of the generated recommendations.

After the ranking procedure has been completed, the user is presented not with a single recommendation but with an ordered set of the most promising scenarios for achieving the specified objective. For each scenario, the system automatically generates a natural-language explanation describing the logical sequence of concept transitions, the evidential justification of individual analytical steps, the degree of support provided by open information sources, and the results of the multi-agent expert evaluation. Consequently, the final output of the proposed methodology is not merely a recommendation but a fully explainable analytical process leading to that recommendation.

An important property of the proposed approach is its adaptability. Since every stage of the analytical process is based on the current Retrieval-Augmented Generation context, the arrival of new information or changes in the information environment automatically trigger updates of the semantic digital twin, revision of alternative scenarios, and their subsequent re-ranking. This enables the methodology to be applied to real-time decision-support problems operating in continuously evolving information environments.

10. Agentic Implementation Using Structured Prompts

The practical implementation of the proposed methodology is based on the **agentic execution of a structured prompt**, which defines the sequence of system analysis procedures, the interactions among agents, and the rules governing the integration of intermediate results. Unlike conventional prompts, which directly formulate a request to a Large Language Model, a structured prompt specifies the algorithm for solving the analytical problem by distributing individual stages of the analysis among specialized agents.

Within the proposed methodology, the structured prompt serves as a **declarative specification** of the system analysis process. It defines the initial system state, the target state, the scenario evaluation criteria, the search parameters, the interaction rules among agents, and the format of the generated results. Consequently, the prompt is regarded not merely as a textual instruction for a Large Language Model but as a formal specification of the system analysis methodology.

Let the agentic system be defined as

$$A=\{A_1,A_2,...,A_m\},$$

where each agent performs a specific transformation of the semantic digital twin.

In the general case, an individual agent implements the operator

$$A_i:(G^*,C)\to(G^*,R_i),$$

where

- G^* denotes the current state of the semantic digital twin;
- C denotes the local analytical context of agent AiA_iAi;
- R_i denotes the result produced by the corresponding system analysis procedure.

Unlike conventional multi-agent systems, the agents do not operate on independent copies of the available information. Instead, they interact through a shared semantic digital twin, which serves as the common analytical workspace. As a consequence, the results produced by each agent become immediately available to all subsequent stages of the analysis without requiring information duplication or repeated interpretation of textual messages.

The practical implementation of the proposed agentic architecture comprises several functional groups of agents.

The first group consists of information-context construction agents. Their responsibilities include constructing the Retrieval-Augmented Generation context, refining search queries, assessing the completeness of the retrieved information, and, whenever necessary, initiating additional retrieval cycles through the InfoStream monitoring system to obtain supplementary documents.

The second group of agents is responsible for constructing the semantic digital twin. Its functions include concept extraction, identification of semantic relationships, concept normalization, integration of the results of forward and backward semantic expansion, and maintenance of the structural consistency of the semantic digital twin.

The third group of agents performs scenario analysis. Its responsibilities include the extraction of alternative scenarios, verification of their logical consistency, and evaluation of their completeness, evidential support, and internal coherence.

The final group of agents performs expert evaluation and explanation of the analytical results. These agents carry out multi-criteria scenario ranking, assess the degree to which the analytical objective has been achieved, generate natural-language explanations, and prepare the final recommendations presented to the user.

An important characteristic of the proposed architecture is the absence of a fixed agent configuration. The number of agents, their functional specialization, and their execution sequence are determined dynamically by the structured prompt according to the requirements of the particular system analysis task. Consequently, the structured prompt specifies not only the information request but also the configuration and execution logic of the agentic system.

The overall agentic analysis process can be represented as the sequence of iterations

$$G_0^* \to G_1^* \to \ldots \to G_t^*,$$

where each iteration corresponds to the completion of one execution cycle of the multi-agent system.

After each execution cycle, the completeness of the constructed semantic digital twin is evaluated. If the predefined information sufficiency criteria are not satisfied, the multi-agent system automatically initiates a new Retrieval-Augmented Generation cycle, refines the information context, and reconstructs the corresponding fragments of the semantic digital twin. In this way, the proposed methodology implements an **adaptive self-refinement mechanism** for the problem-domain model.

This architecture fundamentally differs from the conventional use of Large Language Models. In the traditional workflow, an LLM produces a response immediately after receiving a prompt, and this response represents the final output of the computation. In the proposed methodology, however, the final recommendation is generated only after completion of the entire analytical cycle, including construction of the semantic digital twin, scenario extraction, multi-agent evaluation, and scenario ranking. Consequently, the Large Language Model participates in every stage of the analytical process, whereas the final recommendation represents the collective outcome of the agentic system rather than the output of an individual LLM invocation.

The practical significance of this approach lies in its ability to reproduce the complete analytical process. Unlike conventional prompting, where the user receives only the final response, the proposed agentic implementation preserves the entire sequence of intermediate analytical results, making it possible to reconstruct the scenario-generation process, analyze the contribution of individual agents, and explain the rationale underlying the generated recommendations. This capability is particularly important for decision-support systems, where transparency, reproducibility, and auditability of analytical results are essential requirements.

Functional Decomposition Principle

The proposed methodology is founded on the principle of functional decomposition of system analysis, according to which every classical system analysis procedure is implemented by an individual agent or a dedicated group of agents. Consequently, the agentic architecture reflects the logical structure of the system analysis process rather than the software architecture of the implementation.

Unlike most contemporary Agentic AI architectures, where agents function primarily as general-purpose workers executing individual information requests, each agent in the proposed methodology performs a well-defined analytical function. These functions include information-context construction, semantic digital twin construction, structural verification of the model, generation of alternative scenarios, multi-criteria evaluation, scenario ranking, and interpretation of analytical results. Collectively, these functions constitute the complete system analysis workflow that has traditionally been performed either by human experts or by collections of independent software modules.

This functional decomposition provides several important advantages. First, individual analytical procedures remain independent and can be reused across different problem domains. Second, the architecture is inherently scalable, since introducing new analytical procedures does not require redesigning the overall system but only adding new specialized agents or extending existing ones. Consequently, the proposed agentic system should be regarded as a general execution environment for the semantic system analysis methodology rather than as a software implementation of a particular algorithm.

11. Semantic Forecasting Method Based on the Semantic Digital Twin

The proposed methodology is applicable not only to decision-support problems but also to forecasting the evolution of complex socio-economic, informational, technological, and cyber systems. Unlike conventional forecasting approaches, which rely on time-series extrapolation, statistical models, or simulation techniques, the proposed method employs the semantic digital twin of the problem domain as the foundation for generating, analyzing, and evaluating possible future system states.

The fundamental distinction between forecasting and decision support lies in the absence of a predefined target state. In the decision-support problem considered in the previous section, the target state is explicitly specified by the user. In forecasting problems, however, the future state is unknown and must be generated during the analytical process itself. For this reason, the proposed semantic system analysis methodology introduces an additional stage—generation of a set of potential future states—which precedes the construction of the semantic digital twin.

Let S_0 denote the current state of the problem domain. The first stage of the forecasting process consists of constructing a set of possible future states,

$$F = \{S_1, S_2, ..., S_N\},$$

where N denotes the number of forecasting hypotheses generated by the Large Language Model.

Unlike conventional forecasting methods, these states are not regarded as the final forecasting results. Instead, they represent alternative forecasting hypotheses generated by the Large Language Model on the basis of the current Retrieval-Augmented Generation context. Each hypothesis subsequently undergoes the complete semantic system analysis cycle described for the decision-support methodology.

For each forecasting state S_i, the Bidirectional Semantic Digital Twin Construction procedure is performed independently,

$$S_0 \rightarrow G_f(i),,$$

after which the integrated semantic digital twin is constructed,

$$G_i = G_f(i) \cup G_b(i)..$$

Consequently, each forecasting hypothesis gives rise to its own semantic model of the evolution of the problem domain. This fundamentally distinguishes the proposed methodology from most contemporary LLM-based forecasting systems, which typically generate only a single forecast without subsequently verifying its internal semantic consistency.

After constructing the semantic digital twin for each forecasting hypothesis, the multi-agent system performs scenario extraction, multi-criteria scenario evaluation, and scenario ranking. As a result, each forecasting hypothesis is characterized not by a single textual description but by a set of alternative realization scenarios, each possessing its own evidential basis, logical structure, and integral evaluation score.

Formally, the set of scenarios corresponding to the forecasting state S_i is denoted by

$$P_i = \{p_{i1}, p_{i2}, ..., p_{ik}\}.$$

After evaluating the scenarios, an integral assessment of the forecasting hypothesis, $R(S_i)$, is computed. This assessment takes into account not only the evaluation scores of the individual scenarios but also their mutual consistency, the evidential support provided by open information sources, the structural completeness of the semantic digital twin, and the results of the multi-agent expert assessment.

Unlike conventional forecasting models, which estimate only the probability of the final state, the proposed methodology also evaluates the credibility of the analytical process leading to that state. Consequently, the most probable forecast is not the one generated first by the Large Language Model but the one for which the most semantically consistent, information-supported, and logically coherent semantic digital twin can be constructed.

After completion of the analysis, all forecasting hypotheses are ranked according to the function $R(S_i)$, and the user is presented with an ordered list of the most well-founded future developments together with the corresponding realization scenarios and their natural-language explanations. Consequently, the output of the proposed methodology is not a single forecast but a ranked set of possible future states, each supported by its own semantic digital twin, realization scenarios, and Retrieval-Augmented Generation evidence base.

12. Generalization of the Semantic Systems Analysis Methodology

The proposed decision support method represents a specific implementation of the semantic systems analysis methodology presented in the preceding sections. Its distinguishing characteristic is that the target state of the system is specified by the user before the analysis begins. With respect to this target state, a bidirectional semantic digital twin is constructed, alternative transition scenarios are generated, subjected to multicriteria evaluation and ranking, and subsequently transformed into automatically generated explainable recommendations.

At the same time, an analysis of all stages of the proposed methodology demonstrates that the overwhelming majority of semantic systems analysis procedures are invariant with respect to the method used to determine the target state of the system. Regardless of the problem formulation, the following stages remain unchanged: the construction of a relevant information context using Retrieval-Augmented Generation (RAG), the construction of a semantic digital twin of the problem domain, bidirectional semantic expansion, agent-based refinement of the model, extraction of alternative scenarios, their evaluation and ranking, and the generation of explainable results. The distinction between different classes of problems arises only at the stage of defining the set of possible target states with respect to which the bidirectional semantic analysis is performed.

Let $F = \{S_{g1}, S_{g2}, ..., S_{gn}\}$ denote the set of possible target states of the problem domain.

Then, the proposed methodology can be represented as the universal operator

$$\Psi : (S_0, F) \rightarrow P^*,$$

where

- S_0 is the initial state of the problem domain;
- F is the set of possible target states;
- P^* is the ranked set of scenarios obtained as the result of semantic systems analysis.

In the general case, for each target state g$S_{gi} \in F$, the complete analysis cycle is performed independently. This cycle includes the construction of the relevant information context, the generation of a bidirectional semantic digital twin, the extraction of transition scenarios, their multicriteria evaluation, and subsequent ranking. After the analysis has been completed for all target states, the obtained results are aggregated, making it possible to construct a unified ordered set of alternative scenarios.

In decision support problems, the set of target states consists of a single element $F = \{S_g\}$, where S_g is specified by the user as the desired outcome of the system's operation. Consequently, $|F| = 1$, and the decision support problem represents a special case of the general semantic systems analysis methodology.

In forecasting problems, the situation is fundamentally different. The target state of the system is unknown a priori; therefore, the set $F = \{S_{g1}, S_{g2}, ..., S_{gn}\}$ is generated automatically by a Large Language Model based on the analysis of the current information context. The generated target states are interpreted as a set of alternative forecasting hypotheses, each of which undergoes the complete semantic systems analysis cycle independently. After the analysis has been completed, all generated scenarios are evaluated according to the specified criteria, their probabilities of realization are estimated, and a ranked list of the most probable future developments is produced.

A conceptual generalization of the proposed methodology is presented in Fig. 2. As illustrated in the figure, both decision support and forecasting rely on the same informational foundation provided by Retrieval-Augmented Generation, employ the same procedure for constructing a semantic digital twin, and use identical mechanisms for bidirectional semantic analysis, agent-based evaluation, scenario extraction, and scenario ranking. The only difference lies in the method used to construct the set of target states GGG. In decision support problems, this set is specified by the user and contains a single element, whereas in forecasting problems it is generated automatically by a Large Language Model as a set of possible future states of the system. Consequently, decision support represents a particular case of the more general forecasting formulation, for which the condition $|F| = 1$, holds, whereas forecasting corresponds to the case $|F| > 1$..

The proposed generalization has important methodological implications. It makes it possible to regard decision support and forecasting not as two independent classes of applied problems, but rather as two operational modes of a unified semantic systems analysis methodology. Within this framework, the semantic digital twin of the problem domain serves as the central object of analysis, while Large Language Models, Retrieval-Augmented Generation, and the agent-based architecture function as instruments for its automatic construction, refinement, and exploration. This organization of the analytical process provides the proposed methodology with explainability, reproducibility, and universality, making it applicable to a broad class of systems analysis problems.

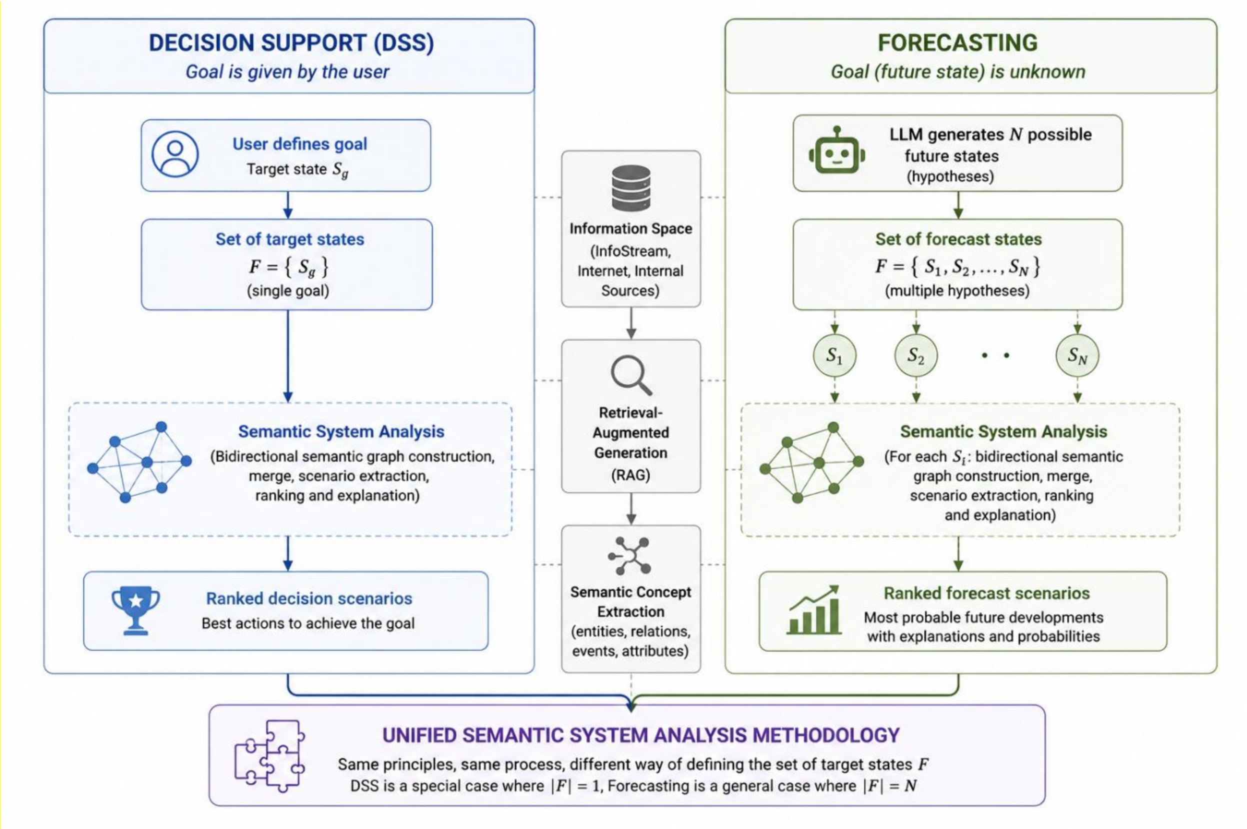

Figure 2. Generalization of Decision Support Systems and Forecasting as Two Special Cases of a Unified Methodology

This formulation makes it possible to move naturally from decision support problems to forecasting problems.

13. Semantic Forecasting Method

Unlike decision support problems, in which the desired target state of the system is known before the analysis begins, in forecasting problems the target state is unknown *a priori*. This uncertainty is the defining characteristic of forecasting and distinguishes its formulation from other classes of systems analysis problems.

In classical forecasting methods, this uncertainty is addressed through the construction of statistical models, time-series extrapolation, regression analysis, or simulation modeling. The proposed methodology adopts a fundamentally different approach. Rather than eliminating uncertainty at the outset, it explicitly incorporates it by constructing a set of possible future states of the system.

Consequently, the first stage of semantic forecasting is the generation of a set of forecasting hypotheses.

Let the current state of the system be denoted by S0S_{0}S0. Based on the information context provided by Retrieval-Augmented Generation and the semantic digital twin, the agent-based system constructs a set of possible future states,

$$F = \{S_1, S_2, \dots, S_N\},$$

where N is the forecasting parameter.

It is essential to emphasize that the set F is **not** the forecasting result itself.

Rather, it represents a set of forecasting hypotheses generated by the Large Language Model in accordance with the current information context. These hypotheses may describe different, including mutually exclusive, future developments. At this stage, the agent-based system does not attempt to identify the most probable forecast; instead, it seeks to provide the most comprehensive possible representation of the space of potential future states.

This step constitutes a fundamentally new element compared with conventional approaches to the use of Large Language Models. In most contemporary LLM-based systems, the forecast is produced as a single textual response generated directly by the model. In the proposed methodology, such a response is regarded merely as one possible forecasting hypothesis and is not accepted without further systems analysis.

For each forecasting hypothesis S_i, the complete semantic systems analysis cycle described in the preceding sections is executed independently. The agent-based system constructs a bidirectional semantic digital twin between the current S_0 and the corresponding forecasting hypothesis S_i, after which it performs scenario extraction, multicriteria evaluation, and scenario ranking.

Consequently, within the proposed methodology, forecasting is effectively transformed into a series of decision support problems that are solved in parallel for each forecasting hypothesis. After all analyses have been completed, the obtained results are integrated, and the forecasting hypotheses are ranked according to the quality of the generated scenarios, the strength of their supporting information, the structural consistency of the semantic digital twin, and the expert assessments produced by the agent-based system.

14. Experimental Validation of the Methodology

The objective of the experimental study is to evaluate the effectiveness of the proposed semantic system analysis methodology in solving decision support and forecasting problems. Unlike most studies devoted to the use of Large Language Models, the present work evaluates not the quality of an individual LLM but the influence of the organization of the system analysis process on the quality of the obtained solutions. Consequently, the subject of the experimental evaluation is the proposed methodology itself rather than a particular software implementation.

The experimental validation was conducted using the same Large Language Model, the same Retrieval-Augmented Generation framework, and a common information repository constructed by the InfoStream monitoring system. This experimental setup eliminated the influence of differences between language models or information resources and made it possible to focus exclusively on the impact of the proposed semantic system analysis methodology.

Two classes of analytical problems were investigated:

- decision support;
- semantic forecasting.

For each problem class, two experiments were conducted.

In the first experiment, a conventional approach was employed, in which the Large Language Model received a standard natural-language prompt and directly generated the final response.

In the second experiment, a structured agentic prompt implementing the complete cycle of the proposed methodology was employed. The workflow included Retrieval-Augmented Generation, construction of the semantic digital twin, bidirectional semantic analysis, multi-agent evaluation, extraction of alternative scenarios, and scenario ranking.

Thus, throughout all experiments, the only varying factor was the organization of the analytical process, whereas the language model, the information repository, and the input data remained identical.

14.1. Decision Support Experiment

The first series of experiments addressed the decision-support problem.

In the conventional approach, the user formulated the analytical task as an ordinary natural-language query. After receiving the Retrieval-Augmented Generation context, the Large Language Model directly generated a recommendation consisting of a list of possible actions accompanied by a brief explanation.

In the second experiment, the proposed agentic implementation of the methodology was employed. The structured prompt specified the initial system state, the target state, the scenario evaluation criteria, the rules governing Bidirectional Semantic Digital Twin Construction, the parameters of agent interaction, and the format of the generated results.

After constructing the information context, the multi-agent system generated the forward and backward semantic subnetworks, integrated them into a unified semantic digital twin of the problem domain, and extracted alternative transition scenarios connecting the initial and target states. The extracted scenarios were subsequently subjected to multi-agent evaluation according to the predefined criteria, ranked, and finally transformed into explainable recommendations.

The structured prompt, the corresponding AgentFlow agentic workflow [8], and representative fragments of the generated responses are presented in Appendix A.

Comparison of the obtained results demonstrated that the proposed methodology transforms the analytical process from generating a single recommendation into constructing a ranked set of alternative scenarios. Furthermore, each scenario is accompanied by a natural-language explanation of the rationale underlying its selection. In addition, every recommendation can be traced back to the corresponding concepts of the semantic digital twin and verified using the supporting evidence provided by the Retrieval-Augmented Generation context.

14.2. Semantic Forecasting Experiment

The second series of experiments was devoted to **semantic forecasting**.

In the conventional approach, the Large Language Model received a description of the current situation and directly generated a forecast of future developments. The resulting output effectively represented a single forecasting hypothesis produced by the language model without any subsequent analysis of alternative possibilities.

The proposed methodology employed a fundamentally different approach.

During the first stage, the multi-agent system generated a set of forecasting hypotheses, each corresponding to one possible future development.

It is important to emphasize that, at this stage, **none of these hypotheses was regarded as a forecast**. Instead, they were considered solely as a set of candidate future states generated by the Large Language Model on the basis of the current Retrieval-Augmented Generation context.

Subsequently, each forecasting hypothesis independently underwent the complete semantic system analysis cycle. The multi-agent system constructed a **Bidirectional Semantic Digital Twin**, performed scenario extraction, carried out multi-criteria scenario evaluation, and ranked the resulting scenarios.

Consequently, forecasting was transformed into a collection of independent system analysis problems, whose results were compared only after completion of the entire analytical process.

The forecasting prompt, the corresponding **AgentFlow** agent workflow, and representative fragments of the generated responses are presented in **Appendix B**.

The experiments demonstrated that the proposed methodology substantially improves the explainability of forecasting results. Instead of receiving a single prediction, the user is provided with a ranked set of forecasting scenarios accompanied by the corresponding causal chains, evidential support for each transition, and an integral assessment of the credibility of each forecast.

14.3. Comparative Analysis of Experimental Results

The comparative analysis demonstrated that the structured agentic methodology fundamentally changes the role of the Large Language Model within the analytical process. In the conventional approach, the language model directly produces the final answer. In contrast, within the proposed methodology, the final answer represents only the last stage of a multi-step semantic system analysis process.

The principal differences between conventional prompting and the proposed semantic system analysis methodology are summarized in Table I.

Table I. Comparison of Conventional Prompting and the Proposed Methodology

Characteristic	**Conventional Prompting**	**Proposed Methodology**
Retrieval-Augmented Generation	Used as contextual	Used as the information foundation of system analysis
Semantic Digital Twin	Not constructed	Automatically constructed
Bidirectional Semantic Analysis	Not performed	Performed
Agentic Functional Decomposition	Not employed	Fully implemented
Alternative Scenarios	Typically a single response	A ranked set of alternative scenarios is generated
Scenario Ranking	Not performed	Automatically generated
Result Explainability	Limited	Automatically constructed
Reproducibility of the Analysis	Partial	Ensured through preservation of all intermediate analytical stages

The obtained results indicate that the proposed methodology does not alter the capabilities of the Large Language Model itself; rather, it transforms the organization of the system analysis process. Specifically, the construction of the semantic digital twin, the bidirectional construction of its structure, the agentic functional decomposition of analytical procedures, and the multi-stage evaluation of scenarios collectively enable the generation of results that are more structured, explainable, and reproducible.

A particularly important finding is that the same methodology can be successfully applied to both decision-support and forecasting problems. The only fundamental difference between these two classes of problems lies in the way the set of target states is formed: in decision-support tasks, it is specified by the user, whereas in forecasting tasks it is generated automatically by the Large Language Model. All subsequent stages—including construction of the semantic digital twin, agentic analysis, scenario extraction, scenario evaluation, and scenario ranking—remain unchanged, thereby confirming the generality and universality of the proposed methodology.

15. Discussion and Conclusions

The conducted study has demonstrated that modern Large Language Models can be employed not only as generators of natural-language responses but also as components of a more general methodology of semantic system analysis. Unlike conventional approaches, in which Large Language Models directly generate recommendations or forecasting statements, the proposed methodology shifts the primary focus to the construction and analysis of a semantic digital twin of the problem domain. The semantic digital twin becomes the central analytical model within which all subsequent system analysis procedures—including alternative scenario generation, evaluation, ranking, and the production of explainable results—are performed.

At the same time, the proposed methodology does not guarantee the absolute correctness of forecasts, since their quality depends on the completeness of the information space, the quality of the Retrieval-Augmented Generation process, the reliability of open information sources, and the characteristics of the underlying Large Language Model. Rather than guaranteeing the correctness of a forecast, the methodology provides the construction of the most well-founded scenarios supported by the available information.

The first contribution of this work is the formulation of a methodology of semantic system analysis that integrates Retrieval-Augmented Generation, Large Language Models, bidirectional semantic expansion, agentic technologies, and scenario analysis into a unified and coherent analytical framework. Unlike the conventional use of LLMs, the proposed approach considers the language model not as a direct generator of the final response but as an instrument for constructing, refining, and analyzing the semantic model of the problem domain.

The second contribution is the introduction of the semantic digital twin of the problem domain as the central informational object of system analysis. Unlike classical knowledge graphs or static semantic networks, the semantic digital twin is automatically constructed from the Retrieval-Augmented Generation context, continuously refined according to the evolving information environment, and employed as the unified analytical model for investigating alternative development scenarios of the system under consideration.

The third contribution is the substantiation of the principle of Bidirectional Semantic Digital Twin Construction. Simultaneous construction of semantic subnetworks from both the initial and the target states, followed by their integration, substantially reduces the search space, identifies the most coherent transition scenarios between system states, and improves both the validity and the explainability of the analytical results.

The fourth contribution is the formalization of the agentic execution of system analysis procedures. The study demonstrates that an agentic architecture may be interpreted not merely as a software technology for organizing computations but as a representation of the functional structure of the system analysis process itself. Specialized agents successively construct, refine, and analyze the semantic digital twin while implementing the individual operators of the proposed methodology.

The fifth contribution is the reinterpretation of the role of Retrieval-Augmented Generation. Within the proposed methodology, Retrieval-Augmented Generation serves not only as a mechanism for extending the context of a Large Language Model but also as the informational mechanism for constructing the semantic digital twin. Through its integration with the InfoStream information-analytical system, the methodology operates on an up-to-date, representative, and verifiable information environment, thereby improving the reliability of both the constructed models and the generated scenarios.

The sixth contribution is the unification of decision support and forecasting within a single methodology of semantic system analysis. It has been shown that both classes of analytical problems share the same informational and algorithmic foundation and differ only in the manner in which the set of target states is constructed. In decision-support problems, this set consists of a single target state specified by the user, whereas in forecasting problems it is generated automatically by the Large Language Model as a collection of alternative forecasting hypotheses.

Consequently, decision support and semantic forecasting are shown to be two special cases of a unified methodology of semantic system analysis based on semantic digital twins.

The experimental study confirmed that the structured agentic organization of the analytical process produces results that are more comprehensive, logically consistent, explainable, and reproducible than those obtained through conventional prompting. The observed improvement is achieved not by modifying the underlying Large Language Model but by replacing direct response generation with the construction and analysis of a semantic digital twin of the problem domain.

The proposed methodology is universal in nature and is independent of any particular implementation of Large Language Models, Retrieval-Augmented Generation, or agent frameworks. It can be applied to a broad range of applications, including information analytics, cybersecurity, situational awareness, risk management, strategic decision support, analysis of complex socio-technical systems, critical infrastructure management, and other domains involving complex and weakly formalized processes.

In summary, this work proposes a new methodology of semantic system analysis in which Large Language Models are employed not as response generators but as instruments for the automatic construction and analysis of semantic digital twins of problem domains.

Future research will focus on methods for automatic learning of semantic digital twins, adaptive real-time updating of their structures, integration with multimodal Large Language Models, incorporation of specialized multi-criteria scenario evaluation models, and the development of comprehensive intelligent semantic system analysis platforms capable of supporting real-time decision making, forecasting, and situational awareness.

References

[1] Rajashekar, N.C., Shin, Y.E., Pu, Y., Chung, S., You, K., Giuffre, M., Chan, C.E., Saarinen, T., Hsiao, A., Sekhon, J. and Wong, A.H., 2024, May. Human-algorithmic interaction using a large language model-augmented artificial intelligence clinical decision support system. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (pp. 1-20)

[2] Bonczek, Robert H., Clyde W. Holsapple, and Andrew B. Whinston. *Foundations of decision support systems*. Academic Press, 2014

[3] Xu, Y., Kibria, G. and Peeta, S., 2025, November. Agentic LLM Framework for Generating Spatial Intelligence to Support Decision-Making in Smart Cities. In *Proceedings of the 1st ACM SIGSPATIAL International Workshop on Spatial Intelligence for Smart and Connected Communities* (pp. 63-71)

[4] Ng, K.K.Y., Matsuba, I. and Zhang, P.C., 2025. RAG in health care: a novel framework for improving communication and decision-making by addressing LLM limitations. *Nejm Ai*, *2*(1), p.AIra2400380

[5] Zhang, X., Du, C., Pang, T., Liu, Q., Gao, W. and Lin, M., 2024. Chain of preference optimization: Improving chain-of-thought reasoning in llms. *Advances in Neural Information Processing Systems*, *37*, pp. 333-356

[6] Mannava, M.K., 2024, December. Causal inference in AI based decision support: Beyond correlation to causation. In *2024 4th International Conference on Ubiquitous Computing and Intelligent Information Systems (ICUIS)* (pp. 176-181). IEEE

[7] D. Lande, I. Svoboda, A. Feger, L. Strashnoy. Decision Support in the Field of Cybersecurity Through the Use of Generative Artificial Intelligence. in Artificial Intelligence: Achievements and Recent Developments , River Publishers, 2025, pp. 29-52. ISBN:9788743800972

[8] Lande, Dmitry; Strashnoy, Leonard. AgentFlow – No-Code Agent Framework Based on Logical Primitives. SSRN Preprint 5285664, DOI: 10.2139/ssrn.5285664 (Jun 22, 2025). – 12 p.

Appendix A

Prompt for Decision Support

#Parameters:

N=4.
K=10.
M=5.
Start="Deficit of funds for defense".
Goal="Reliable defense".
Ranking criterion="Effectiveness".

#Algorithm:

We have an Initial state defined by the first concept – Start.
Let us call it Start. We have a goal, Goal — this is the final concept of the network being formed.

Two semantic networks are built. The first directed network is built as follows:
The initial state is the primary concept – "Start".
From it, K = 10 concepts are found that are consequences of the first concept and contribute to achieving the Goal. Then, for each obtained concept, another K = 10 concepts are found that are consequences of the obtained concept in the direction of achieving the goal. For each of these obtained concepts, another K = 10 concepts are found that are consequences of the obtained concept in the direction of achieving the goal: Goal. The process continues similarly for N levels. Each concept consists of no more than three words.

The second network is built as follows. It starts from the goal – Goal. At the first stage, 10 concepts are found that lead to the goal, that is, those for which the goal is a direct consequence. For each of the obtained concepts, K = 10 concepts are found that are their causes. The process continues similarly for N levels.

The two formed networks are merged and presented as a list of pairs of concepts, where the first is the primary concept – "Start", and the last is the goal – "Goal".

Select M possible chains of concepts — scripts from the initial concept to the concept of achieving the goal, starting from the initial state to the final concept, denoted as "Start", including all nodes along the path. Rank these chains-scripts according to the parameter "Ranking criterion". For each chain-script, create and provide an extended explanation and ways to solve the problem.

Provide the text in English in the form of a table with the columns:

- number - chain (scenario) - extended explanation =.

Provide everything as a table with borders. Provide everything exclusively in HTML format. Implement the arrows in the chains as ==>.

Also, in the form of a table, for each chain provide in a clear extended form a clear implementation path for this chain and an assessment of its effectiveness, with the columns:
= Number - Implementation path - effectiveness assessment =.

#Text for analysis

Text for analysis: [Text]

Agentic DSS Workflow and an Example of Decision Support for the Goal "Reliable Defense"

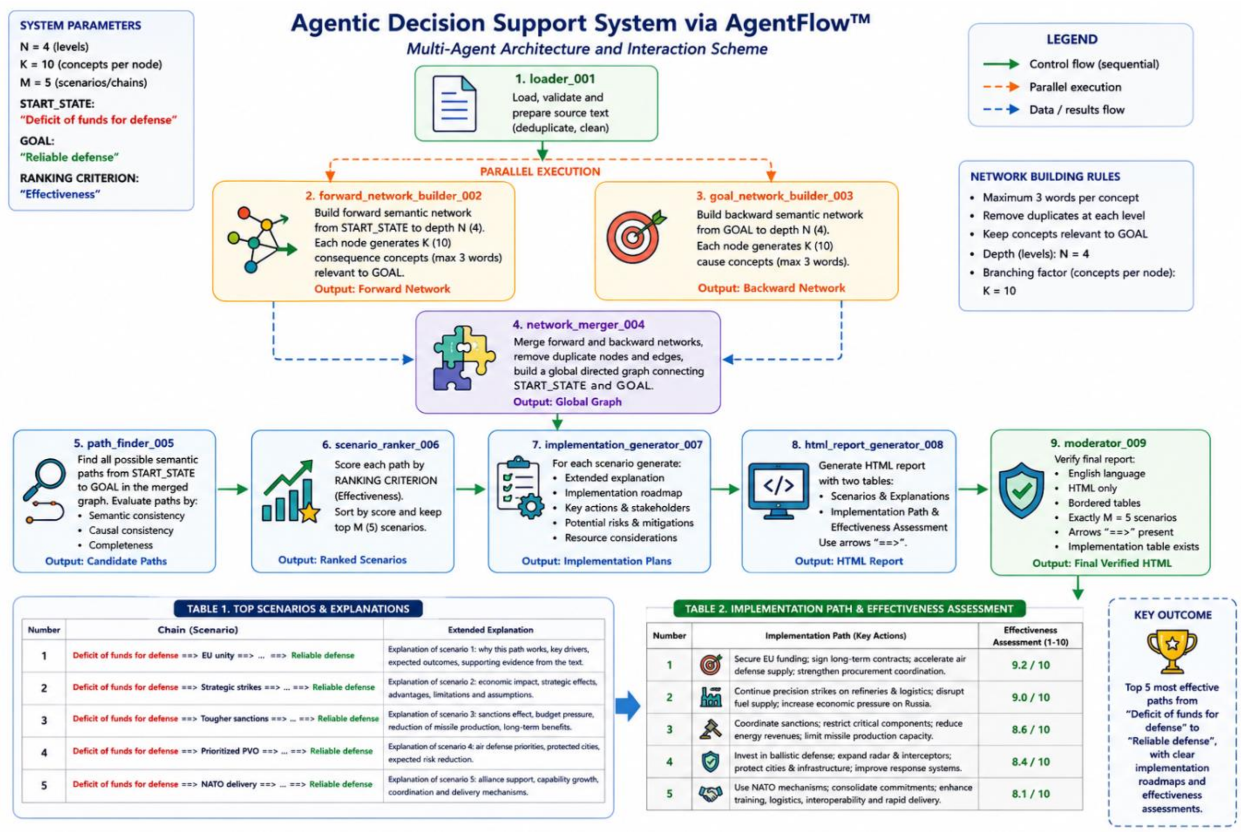

Example of the Agentic System Output

Number	Chain (Scenario)	Extended Explanation
1	Deficit of funds for defense → EU unity → Defense contracts → Air defense supply → Reliable defense	The most effective scenario is based on European consolidation and new defense contracts. The analyzed text shows that European leaders are closing ranks, NATO is preparing delivery-oriented decisions, and Ukraine urgently needs stronger air defense after massive missile and drone attacks. This path directly connects financial shortages with external procurement and practical defense reinforcement.
2	Deficit of funds for defense → Strategic strikes → Refinery disruption → Russian fuel crisis → Reliable defense	This scenario relies on asymmetric pressure. Ukrainian strikes disabling a large share of Russian refining capacity create a fuel crisis, reduce logistical flexibility, increase repair delays, and weaken Russia's ability to sustain offensive operations. It is highly effective because it uses targeted pressure to compensate for limited Ukrainian resources.

3	Deficit of funds for defense → Tougher sanctions → Russian budget deficit → Lower missile production → Reliable defense	The text links Ukraine's defense needs with sanctions and Russia's economic weakening. If sanctions deprive Russia of the ability to produce missiles and reduce energy revenues, Moscow's military-industrial capacity declines. This makes Ukrainian defense more reliable by reducing the intensity and sustainability of future attacks.
4	Deficit of funds for defense → Ballistic threat → Prioritized PVO → Protected cities → Reliable defense	The Kyiv attack demonstrates a critical gap: Ukraine intercepted most drones and cruise missiles but only a small share of ballistic missiles. Therefore, scarce funds should be prioritized toward systems capable of countering ballistic threats, protecting cities, civilians, and critical infrastructure.
5	Deficit of funds for defense → NATO delivery → Regional coordination → Capability growth → Reliable defense	This scenario follows from the emphasis on NATO delivery and implementation. Coordinated support from NATO members, including ammunition, air defense, logistics, and training, can transform limited national resources into stronger collective defense capacity.

Number	**Implementation Path**	**Effectiveness Assessment**
1	Secure EU-backed multi-year funding; connect Ukraine's urgent air defense needs with European defense contracts; prioritize delivery of interceptors, radar systems, launchers, and ammunition; create a predictable procurement schedule.	Very high: 9.2/10
2	Continue intelligence-based targeting of Russian refining, fuel storage, and logistics infrastructure; assess effects through fuel shortages, repair delays, freight disruption, and regional alerts; synchronize strikes with sanctions pressure.	Very high: 9.0/10
3	Coordinate with EU and G7 partners to tighten sanctions on missile components, energy revenue, dual-use technologies, and repair equipment; monitor Russia's budget deficit and military production constraints.	High: 8.6/10
4	Allocate scarce resources first to counter-ballistic defense for Kyiv and other high-risk cities; expand radar coverage, protected shelters, emergency response capacity, and integrated air-defense command systems.	High: 8.4/10
5	Use NATO delivery mechanisms to consolidate partner commitments; request specific capability packages rather than general support; focus on interoperability, logistics corridors, training, and rapid deployment of promised systems.	High: 8.1/10

Appendix B

Prompt for Frorecasting

#Forecast system

#Parameters:

N=10.
K=5.
M=5.
Start=”Deficit of funds for defense”.
Ranking criterion=”Effectiveness”.

#Algorithm:

I want to develop a foresight — a meaningful prediction — a forecast on the topic “Start”.

Select N possible different forecast states representing the future development of the state “Start”, expressed as N concepts.

The first network has the first node — the Initial state Start. From it, M concepts follow. Then, for each of these concepts, another M concepts that follow from it are selected. Then, for each of these concepts, another M concepts that follow from it are selected. The process continues similarly up to K levels.

The other N networks are constructed in the reverse direction. That is, for each of the N identified forecast-state concepts, determine M concepts that may lead to these states. Then perform the same procedure for each of the selected state concepts. Continue the process similarly up to K levels.

After all networks have been constructed, they are merged.

Next, find all possible paths from the first node — the current state — to all N final forecast states. Then rank all chains according to their probabilities.

Output the N=10 most probable concept chains in the form of a table with the following columns:

Number, Probability of realization (the sum of all probabilities across all scenarios must not exceed 1), Concept chain with arrow connections, Explanation.

Additionally, provide a generalized explanation covering all chains.

Provide the text in English in HTML format.

Use bordered tables as much as possible. That is, output only the table and the generalized explanation.

#Text for initial analysis

Text for initial analysis: [Text]

Agentic Forecasting System and an Excerpt from the Forecasting Results for the State “Deficit of Funds for Defense”

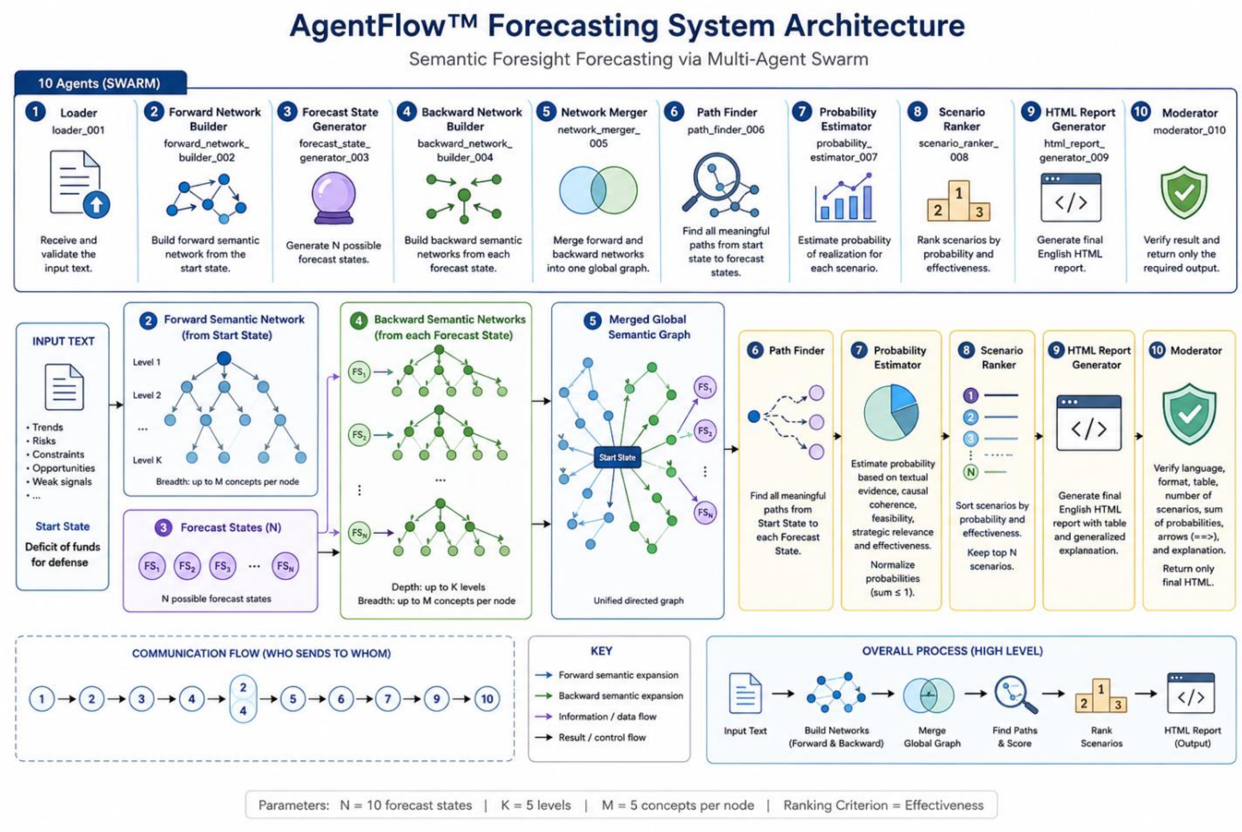

Example of the Agentic System Output

Number	Probability of realization	Concept chain with arrow connections	Explanation
1	0.17	Deficit of funds for defense → EU financial support → Joint procurement → Capability reinforcement → Sustainable defense	The most probable scenario is based on increased European financial support and multinational procurement mechanisms. The acquisition of armored vehicles for Moldova through an EU grant illustrates the growing tendency toward collective financing of defense capabilities.
2	0.14	Deficit of funds for defense → Joint production → European industry → Lower procurement costs → Defense resilience	Growing defense demand across Europe is likely to stimulate multinational production programs, allowing participating countries to reduce unit costs and improve long-term supply stability.
3	0.12	Deficit of funds for defense → Russian economic pressure → Economic crisis →	The analyzed reports suggest that prolonged economic pressure may gradually weaken Russia's ability to sustain military operations. Rather than

		Reduced military capacity → Strategic advantage	expecting rapid battlefield breakthroughs, Ukraine may increasingly rely on economic attrition.
4	0.10	Deficit of funds for defense → Drone dominance → Cost-effective warfare → Lower operational costs → Improved defense	The growing effectiveness of drone warfare significantly changes the economics of modern conflict. Low-cost unmanned systems increasingly substitute expensive conventional weapons, improving defense efficiency under financial constraints.
5	0.09	Deficit of funds for defense → NATO spending growth → Allied investment → Stronger deterrence → Regional security	Although several NATO members still miss spending targets, the alliance is moving toward significantly higher long-term defense investments. This trend is expected to strengthen collective deterrence.
6	0.08	Deficit of funds for defense → Infrastructure investment → Logistics modernization → Faster deployment → Operational efficiency	European infrastructure initiatives, including transport corridors and logistics modernization, indirectly strengthen defense capabilities by improving military mobility and supply chains.
7	0.07	Deficit of funds for defense → EU budget disputes → Funding delays → Procurement slowdown → Temporary capability gap	Disagreements regarding the future EU budget may temporarily delay funding for defense programs, creating short-term procurement challenges despite long-term strategic consensus.
8	0.06	Deficit of funds for defense → US strategic shift → European autonomy → Increased EU defense → Independent capability	The evolution of U.S. foreign policy encourages Europe to assume greater responsibility for its own security. This trend may accelerate independent European defense investments and industrial cooperation.
9	0.05	Deficit of funds for defense → Economic uncertainty → Budget constraints → Selective modernization → Prioritized capabilities	Persistent macroeconomic uncertainty is likely to force governments to prioritize the most critical military capabilities while postponing less essential modernization projects.
10	0.04	Deficit of funds for defense → Long-term adaptation → Industrial transformation → Technological superiority	The least immediate but strategically important scenario foresees a gradual transformation of defense industries toward more automated, technology-intensive production capable of

		→ Sustainable security	supporting long-term security with greater efficiency.

Generalized Forecast
The total probability of the presented scenarios equals **0.92**, leaving **0.08** for less predictable developments. The overall foresight indicates that defense funding shortages are unlikely to be solved by a single measure. Instead, the most probable trajectory combines increased European financial support, multinational procurement, joint defense production, and technological transformation driven by unmanned systems and industrial cooperation. Economic pressure on Russia is expected to complement these efforts by gradually reducing its military sustainability. At the same time, political disagreements over EU budgets, global economic uncertainty, and changing U.S. strategic priorities may slow progress, making prioritization and efficient allocation of limited defense resources increasingly important. The most resilient future scenario combines collective European investment, defense-industrial integration, cost-effective technologies, and long-term strategic planning.