A Cognitive Cyber Range Based on Semantic Networks and Large Language Models

Dmytro Lande[1], Denys Honcharov[1], Mykyta Lytvyn[1]

[1] Igor Sikorsky Kyiv Polytechnic Institute

Abstract

Modern cyber ranges are primarily designed for attack simulation, information system security assessment, and cybersecurity training. However, they provide only limited support for decision-making related to selecting optimal incident response strategies under the conditions of real network infrastructures and rapidly evolving cyber threats. This paper proposes the concept of a cognitive cyber range that integrates a digital representation of the network infrastructure, semantic knowledge networks, Retrieval-Augmented Generation (RAG), large language models, and multi-agent analysis into a unified decision support platform.

The proposed approach is based on the construction of a directed semantic network that integrates knowledge about information system assets, vulnerabilities, attack techniques, security controls, regulatory requirements, personnel competencies, and other domain-specific entities. Unlike conventional scenario discovery methods, a bidirectional subnetwork construction algorithm is proposed that independently performs forward expansion from the current system state and backward expansion from a specified security objective. After the generated subnetworks are aligned and merged, a unified causal model is constructed, from which alternative scenarios for achieving the specified objective are automatically extracted.

To evaluate the generated scenarios, a multi-agent framework of virtual experts based on large language models is employed. The framework performs multi-criteria analysis according to technical, organizational, and regulatory criteria, generates explanations for the obtained recommendations, and determines their level of justification. The proposed approach combines structured knowledge, semantic analysis, and generative artificial intelligence within a unified cognitive decision support environment suitable for both cybersecurity education and practical deployment in cyber defense systems.

Keywords: *cognitive cyber range; semantic network; large language models; Retrieval-Augmented Generation; multi-agent system; decision support; cybersecurity; response scenarios; causal analysis; directed graphs; semantic search.*

1. Introduction

The rapid development of digital technologies, the widespread adoption of cloud services, the Internet of Things (IoT), cyber-physical systems, artificial intelligence technologies, and digital twins has led to a significant increase in the complexity of modern information infrastructures. At the same time, the scale and diversity of cyber threats continue to grow, increasingly taking the form of multi-stage campaigns that employ numerous interconnected attack techniques, rapidly changing tactics, and extensive use of open information sources. Under these conditions, effective cybersecurity requires not only mechanisms for detecting individual incidents but also intelligent systems capable of analyzing large volumes of heterogeneous information, identifying causal relationships among events, predicting possible developments, and supporting the decision-making process.

One of the most important directions in the development of practical cybersecurity technologies has been the emergence of cyber ranges, which are designed for modeling network infrastructures, emulating cyber incidents, testing security mechanisms, and training cybersecurity professionals in controlled environments [1], [2]. The adoption of digital twins has significantly improved the realism of such platforms, bringing them closer to real operational

information systems [3]. Consequently, modern cyber ranges are widely used for vulnerability assessment, cyber resilience evaluation, personnel training, validation of incident response procedures, and testing of new security technologies.

However, most existing cyber ranges are primarily designed to execute predefined scenarios or emulate specific categories of cyber attacks. Even when scenarios can be partially modified during execution, their overall structure is defined in advance by experts. The analysis of alternative response options, selection of optimal response strategies, consideration of the current network state, up-to-date cyber threat intelligence, regulatory requirements, and available resources remain largely the responsibility of human analysts. Thus, contemporary cyber ranges effectively reproduce known scenarios but provide only limited capabilities for intelligent decision support.

In recent years, substantial progress has been achieved in the fields of Open Source Intelligence (OSINT) [4], Retrieval-Augmented Generation (RAG) [5], large language models (LLMs) [6], knowledge graphs, semantic networks, and multi-agent artificial intelligence systems. Retrieval-Augmented Generation constrains response generation using relevant documents, thereby reducing the risk of producing unreliable information. Large language models provide high-quality semantic processing of natural language, while multi-agent architectures enable expert analysis from multiple perspectives and the generation of consistent recommendations. Semantic networks and knowledge graphs offer structured representations of domain knowledge and support the application of graph analysis algorithms for investigating relationships among domain entities.

A review of recent research indicates that these technological directions have largely evolved independently. Studies on cyber ranges primarily focus on network environment modeling, infrastructure virtualization, traffic generation, and cyber incident emulation. Research on knowledge graphs and semantic networks mainly addresses cyber threat intelligence integration, automated analysis of cyber intelligence data, modeling dependencies between vulnerabilities and attack techniques, and domain visualization. Numerous studies investigate the application of large language models for document analysis, question answering, information summarization, and analytical support. A separate research direction concerns multi-agent systems in which different language models perform specialized analytical roles, while their outputs are integrated through automated reasoning mechanisms.

At the same time, there is a notable lack of research in which a cyber range is considered not merely as an environment for reproducing predefined scenarios but as an intelligent system capable of automatically synthesizing new decision support scenarios. In most existing studies, semantic networks [7], [8] are employed only for knowledge representation or visualization, large language models are used primarily for generating textual responses, and graph algorithms are applied to the analysis of already constructed graphs. The integration of these components into a unified cognitive framework capable of automatically generating, evaluating, and explaining alternative cybersecurity strategies for a specific network infrastructure remains insufficiently explored.

The next stage in the evolution of cyber range technology should therefore involve the transition from reproducing predefined scenarios to the automatic synthesis of response strategies adapted to the current state of the information system, current operational information, available resources, and specified security objectives. Such a transition fundamentally changes the concept of the cyber range itself. Whereas a conventional cyber range is primarily a simulation and training platform, a cognitive cyber range becomes an intelligent decision support system capable not only of reproducing known scenarios but also of autonomously generating, analyzing, explaining, and ranking alternative strategies for achieving the required level of cybersecurity.

The realization of this concept requires the integration of a digital representation of the network infrastructure, structured domain knowledge, Retrieval-Augmented Generation, large language models, and semantic network analysis algorithms. Within this framework, the central system model is a directed semantic network whose vertices represent information system assets, vulnerabilities, attack techniques, security controls, regulatory requirements, personnel competencies, resources, and objectives, while the directed edges capture causal, functional, and technological relationships among these entities.

Unlike traditional approaches, this paper proposes a bidirectional algorithm for constructing interconnected subnetworks within a directed semantic network. The algorithm independently performs forward expansion from the current system state and backward expansion from the specified security objective. The resulting subnetworks are subsequently aligned and integrated into a unified causal model, from which alternative scenarios for achieving the target objective are automatically extracted. Large language models and a multi-agent framework are employed not to construct the graph itself but to perform semantic analysis, evaluation, explanation, and multi-criteria ranking of the generated scenarios.

Accordingly, the cognitive cyber range is introduced as a new class of intelligent decision support systems that combines the capabilities of conventional cyber ranges, digital twins, semantic networks, Retrieval-Augmented Generation, and large language models. Its principal distinguishing feature is the automatic synthesis of alternative cybersecurity strategies based on structured knowledge and the current operational context rather than the execution of predefined scenarios.

The objective of this study is to develop the concept of a cognitive cyber range as a new generation of cybersecurity decision support systems capable of automatically generating, explaining, and multi-criteria ranking response scenarios based on directed semantic networks, a bidirectional interconnected subnetwork construction algorithm, and multi-agent analysis employing large language models.

To achieve this objective, the following research tasks are addressed: the development of a cognitive cyber range architecture; the construction of an integrated directed semantic network of the application domain; the development of a bidirectional algorithm for interconnected subnetwork construction and automatic scenario synthesis; the design of a multi-agent model for evaluating and ranking alternative solutions; and the experimental validation of the proposed approach using a representative network infrastructure.

The proposed approach extends the functionality of conventional cyber ranges by transforming them from platforms for rehearsing predefined scenarios into cognitive decision support platforms capable of adapting to environmental changes, integrating new knowledge, and automatically generating well-founded cybersecurity strategies. In the authors' view, this direction will define the next stage in the evolution of Cyber Range systems.

2. Semantic Cyber Range

Traditionally, a cyber range is considered a specialized software and hardware environment designed for modeling information systems, reproducing cyberattacks, testing security mechanisms, and training personnel. The main function of such platforms is to create the most realistic possible conditions for the operation of network infrastructure, in which various scenarios of cyber incident occurrence and mitigation can be practiced. In most cases, simulation scenarios are developed in advance by experts, while the cyber range ensures their execution, event logging, and evaluation of operator performance.

This approach has proven effective in professional training, cyber resilience assessment of information systems, and security technology testing. At the same time, its capabilities are limited when addressing decision support tasks under real operational conditions of information

infrastructure. The growing number of interconnected cyber threats, the continuous updating of information on new vulnerabilities, the rapid evolution of attacker tactics, and the need to simultaneously consider technical, organizational, and regulatory factors complicate the use of predefined scenarios. Under such conditions, response scenarios need to be generated directly during the analysis of the current situation.

The semantic cyber range proposed in this paper is considered the next stage in the evolution of traditional cyber ranges. Its main distinction is the transition from executing predefined scenarios to their automatic synthesis based on a structured domain model that represents knowledge about information infrastructure, cyber threats, security controls, regulatory requirements, and the relationships among them.

Unlike classical Cyber Range platforms, in which the network environment is the central object of modeling, in a semantic cyber range this role is assumed by a directed semantic network that integrates all domain entities into a unified causal model. The network vertices may represent information system assets, network services, vulnerabilities, software tools, attack techniques, security mechanisms, organizational procedures, regulatory documents, personnel competencies, resources, intermediate states, and system objectives. Directed edges represent different types of dependencies among these entities: causal, functional, technological, organizational, or logical.

This representation of knowledge makes it possible to formulate decision support as a directed graph analysis problem, in which response scenarios correspond to possible transition paths from the current system state to a specified target state. Unlike classical attack graphs or attack trees, the semantic network contains not only information about possible adversary actions but also data on available resources, constraints, regulatory requirements, personnel competencies, and other factors that affect the choice of response strategy.

An important feature of the proposed concept is that the semantic network is not static. Its structure is constructed and refined during the analysis of a specific situation by integrating information from various sources. These sources include a digital description of the real network infrastructure, results of Open Source Intelligence (OSINT) analysis, cyber threat knowledge bases, regulatory documents, corporate policies, and other sources. As a result, an up-to-date cognitive model is formed that reflects the current state of the information system and the surrounding information environment.

A fundamental distinction of the semantic cyber range from traditional decision support systems lies in the method of scenario generation. Whereas in most existing systems the analysis is performed on an already constructed graph, in the proposed approach the network is formed by a bidirectional algorithm for constructing interconnected subnetworks [9]. First, a subnetwork of possible situation evolution from the current system state and a subnetwork of the necessary conditions for achieving the specified objective are constructed independently. After these subnetworks are aligned and integrated, a unified causal model is formed that contains feasible transition paths between the initial and target states. These paths are then interpreted as alternative response scenarios.

The obtained scenarios are not final decisions. They are subsequently analyzed by a multi-agent system based on large language models, which performs their technical, organizational, and regulatory evaluation, carries out multi-criteria ranking, generates explanations for the resulting recommendations, and determines their degree of justification. Thus, the semantic network provides structured knowledge representation and alternative generation, while large language models are used for intelligent analysis and explanation of the generated scenarios rather than for direct network construction.

The proposed concept enables the integration of digital twins of information infrastructure, modern cyber intelligence tools, structured knowledge bases, semantic network analysis algorithms, and generative artificial intelligence technologies within a unified environment. As a

result, the traditional cyber range is transformed into a cognitive cyber range: an intelligent decision support platform that not only models the evolution of cyber incidents but also automatically generates, analyzes, explains, and ranks alternative cyber defense strategies according to the current state of the information system and the specified objectives.

3. Architecture of the Cognitive Cyber Range

The proposed cognitive cyber range implements an integrated decision support architecture in which the digital representation of the information infrastructure, semantic knowledge representation, graph analysis algorithms, and large language models operate as interconnected components of a unified information and analytical environment. Unlike conventional cyber ranges, where the central element is the simulation of network infrastructure and the execution of predefined scenarios, the proposed architecture is designed for the automatic synthesis of alternative response scenarios according to the current state of the system and the specified cybersecurity objectives.

The overall architecture of the cognitive cyber range is shown in Fig. 1. It is based on a modular design and provides the sequential transformation of input data into decision support recommendations. The central component of the architecture is a directed semantic network, which serves as the unified knowledge representation model and enables interaction among all functional components of the system.

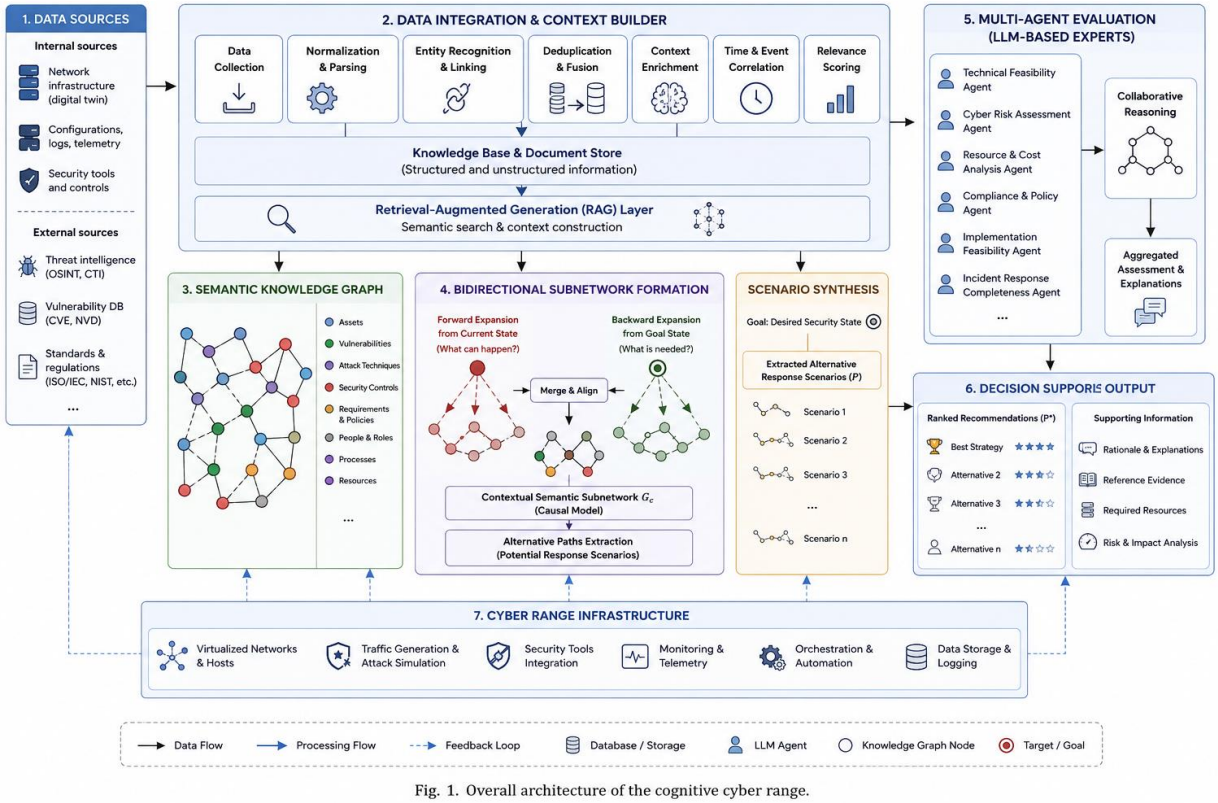

Fig. 1. Overall architecture of the cognitive cyber range.

Figure 1. General Architecture of the Cognitive Cyber Range

At the initial stage, the task context is established. The system receives a digital representation of the real information infrastructure in the form of network models, equipment configurations, or a digital twin of the target environment. At the same time, the analysis objectives, admissible constraints, evaluation criteria, and other parameters defining the problem conditions are specified. As a result, a formalized description of the current situation is created, which is subsequently used by all architectural components.

The information support of the cognitive cyber range is based on the integration of internal and external knowledge sources. These include the results of open-source information monitoring, vulnerability databases, descriptions of attack techniques, recommendations issued by cyber incident response centers, regulatory documents, corporate policies, and the accumulated operational experience of the system. A hybrid Retrieval-Augmented Generation mechanism is employed to obtain relevant information by combining full-text and semantic search, thereby constructing the information context required for subsequent analysis.

The next functional layer of the architecture is the construction of an integrated directed semantic network of the application domain. Unlike static knowledge graphs, this network is built specifically for a particular decision support task by integrating the digital model of the information system with up-to-date information obtained from external sources. The vertices of the network represent domain entities, including assets, services, vulnerabilities, attack techniques, security controls, regulatory requirements, personnel competencies, resources, and target states, whereas the directed edges describe the causal, functional, and technological relationships among them. As a result, a cognitive model is constructed that reflects the current state of the information infrastructure and the interdependencies among its components.

Based on the constructed semantic network, possible response scenarios are synthesized. Unlike existing systems, where analysis is performed on an already constructed graph, the proposed architecture employs a bidirectional algorithm for constructing interconnected subnetworks. A subnetwork representing the possible evolution of the situation from the current system state and a subnetwork representing the conditions required to achieve the specified objective are constructed independently. After these subnetworks are aligned and integrated, a unified causal model is formed, within which feasible transition paths between the initial and target states are automatically identified. These paths are interpreted as alternative cybersecurity response scenarios.

The generated scenarios are then submitted to a multi-agent intelligent analysis framework based on large language models. All agents operate on the same information context and a common semantic model but perform different expert functions related to evaluating the technical feasibility of the scenarios, the associated level of risk, compliance with regulatory requirements, required resources, and other aspects of decision-making. The individual assessments are integrated into a consolidated expert evaluation accompanied by an explanation of the rationale behind the selected alternatives.

The final stage of the architecture is the multi-criteria ranking of the generated scenarios. For each alternative, the values of criteria characterizing effectiveness, implementation cost, execution time, risk level, resource requirements, compliance with regulatory documents, and the reliability of the underlying information are determined. The system output is an ordered list of possible response strategies, each accompanied by an explanation of the logic underlying its generation and a justification of its applicability.

An important feature of the proposed architecture is the presence of a feedback loop. After the selected scenario has been implemented, the execution results, expert assessments, and accumulated practical experience are returned to the system knowledge base. This ensures the continuous refinement of the semantic network structure, the improvement of the ranking models, and the enhancement of recommendation quality when solving subsequent decision support tasks.

Thus, the proposed architecture implements a complete decision support cycle, ranging from the integration of a digital model of the information infrastructure and current cyber threat knowledge to the automatic synthesis, analysis, explanation, and ranking of alternative cybersecurity response scenarios. The use of a unified directed semantic network as the central knowledge representation model ensures the consistency of all system components and provides

the foundation for implementing the cognitive cyber range as a new generation of decision support platforms.

4. Construction of the Directed Semantic Network

The central component of the proposed cognitive cyber range is a directed semantic network that integrates information about information infrastructure objects, cyber threats, security controls, regulatory requirements, and the results of open-source intelligence analysis. Unlike conventional knowledge graphs, which are constructed once and subsequently only updated, the semantic network in the proposed approach is built dynamically according to the specific decision support task. Its structure is determined not only by the accumulated knowledge but also by the current state of the information system, the characteristics of the incident, the defined analysis objectives, and the current information context.

The network is constructed through the integration of several independent information sources. These include a digital description of the network infrastructure, the results of open-source information monitoring, databases of known vulnerabilities, cyber threat taxonomies, regulatory documents, corporate policies, accumulated incident response experience, and other structured or semi-structured data. Before integration, all data undergo normalization, deduplication, semantic alignment of terminology, and entity identification.

Formally, the directed semantic network is represented by the directed graph

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

where V is the set of vertices corresponding to domain entities, and $E \subseteq V \times V$ is the set of directed relationships between them.

Unlike conventional attack graphs, the set of vertices is not limited to information describing potential attacker actions. The network includes multiple classes of entities, including information system assets, software and network components, vulnerabilities, attack techniques, security mechanisms, organizational procedures, regulatory requirements, personnel competencies, resources, intermediate system states, and target security states.

Each directed edge represents a specific relationship between two entities. These relationships may be causal, functional, technological, organizational, or logical in nature. The direction of an edge defines the possible direction of process evolution within the domain and enables the subsequent identification of feasible transition paths between system states.

For each vertex $v_i \in V$, a set of attributes $A(v_i) = \{a_1, a_2, ..., a_m\}$ is defined, which includes its type, description, source of information, temporal characteristics, confidence level, association with a particular information system, and other relevant properties.

Similarly, each edge $e_{ij} = (v_i, v_j)$ is associated with a set of attributes $B(e_{ij}) = \{b_1, b_2, ..., b_k\}$, which characterizes the relationship type, its weight, confidence level, temporal validity, and supporting information sources.

An important feature of the proposed model is the use of weighting coefficients to evaluate the significance of relationships. The weight of an edge is determined not only by the frequency of co-occurrence of the corresponding entities in information sources but also by the results of their semantic analysis, the credibility of the information sources, and expert assessments. In general, the edge weight can be represented as $w_{ij} = f(c_{ij}, r_{ij}, t_{ij}, e_{ij})$, where c_{ij} characterizes the strength of the semantic relationship, r_{ij} denotes the credibility of the information sources, t_{ij} represents the timeliness of the information, and e_{ij} denotes the expert assessment of the relationship.

The specific form of the function $f(\cdot)$ is determined by the characteristics of the application domain and may be refined during system operation.

During information integration, different sources may contain information describing the same entity. Therefore, one of the key stages in network construction is the semantic merging of equivalent vertices. This is accomplished using methods for term normalization, named entity matching, contextual analysis, and semantic search. As a result, all information related to a particular entity is aggregated into a single vertex, while the corresponding relationships are automatically redirected to it.

The constructed semantic network is not a static structure. It is continuously expanded with new vertices and relationships as new information becomes available from open sources, the network infrastructure configuration is updated, or new cyber threats emerge. In this way, a dynamic cognitive model of the application domain is formed, reflecting both the current state of the information system and the accumulated knowledge of its operation.

A distinctive feature of the proposed approach is that the constructed directed semantic network is used not merely as a knowledge representation mechanism. It serves as the unified information model on the basis of which subnetwork construction, automatic synthesis of alternative response scenarios, multi-agent analysis, and decision support are performed. Consequently, its structure must simultaneously ensure comprehensive knowledge representation, scalability, efficient updating, and effective execution of graph search algorithms.

5. Bidirectional Algorithm for Constructing the Contextual Semantic Subnetwork

After constructing the integrated directed semantic network of the application domain, the next task is to build a model sufficient for supporting decision-making in a specific situation. The use of the entire semantic network is impractical because it may contain a large number of vertices and relationships, most of which are not directly relevant to the current incident or the specified objective. Moreover, analysis of the complete network significantly increases computational complexity and complicates the interpretation of the obtained results.

In the proposed approach, the object of further analysis is not the entire semantic network but only its contextual semantic subnetwork, which is constructed according to the specific decision support task. The context C is defined by the set of characteristics describing the current situation, including the description of the information infrastructure, the initial system state, the target state, the specified constraints, the temporal framework, the available resources, and the information corpus generated by Retrieval-Augmented Generation.

Let $C = \langle s, g, \Omega, D \rangle$, where s is the current state of the information system, g is the target state, Ω is the set of constraints, and D is the set of documents and knowledge relevant to the current task.

The directed semantic network $G = (V, E)$, constructed in the previous stage, is regarded as the global model of the application domain. Depending on the task context, a contextual semantic subnetwork $G_C = (V_C, E_C)$, is constructed from it, where $V_C \in V, E_C \in E$.

Thus, the contextual semantic subnetwork is a function of both the global semantic network and the current context: $G_C = \Phi(G, C)$, where $\Phi(\cdot)$ is the contextual subnetwork construction operator.

Unlike conventional path-search algorithms, which operate on an already constructed graph, the proposed algorithm first constructs the relevant region of the semantic network and subsequently

performs the analysis only within that region. This substantially reduces the size of the problem and enables the model to adapt to the specific situation.

The contextual semantic subnetwork is constructed using a bidirectional algorithm. In the first stage, forward expansion is performed from the vertex corresponding to the current system state. During this expansion, all vertices that can be reached along the directed edges of the semantic network and satisfy the constraints of the current context are successively included in the subnetwork. As a result, the set $V_F = \{v \in V \mid s \Rightarrow v\}$, is constructed, describing the possible directions of situation evolution.

Independently, backward expansion is performed from the vertex corresponding to the target state. Unlike conventional reverse traversal, this process does not reverse the orientation of the graph edges. Instead, the algorithm successively identifies the vertices from which the target state can be reached, thereby constructing the set $V_B = \{v \in V \mid v \Rightarrow g\}$.

Thus, two complementary regions of the semantic network are obtained: one describing the potential evolution of the situation, and the other representing the set of conditions required to achieve the specified objective.

The next stage consists of integrating both subnetworks. This involves semantic alignment of common vertices, merging of their attributes, elimination of duplicate information, and updating of the weighting characteristics of the relationships. As a result, the contextual semantic subnetwork is obtained, $G_C = (V_C, E_C)$, where $V_C = V_F \cup V_B, E_C = E_F \cup E_B$.

This subnetwork is subsequently used throughout the following stages of the cognitive cyber range.

After the contextual semantic subnetwork has been constructed, directed paths between the vertices s and g are identified. Each discovered path $P_i = (s, v_1, ..., v_k, g)$ is interpreted as an alternative cybersecurity response scenario. Unlike traditional scenarios, which are predefined by experts, the proposed scenarios are synthesized automatically based directly on the structure of the constructed contextual semantic subnetwork.

A distinctive feature of the proposed algorithm is that the contextual semantic subnetwork does not exist independently of the decision support task. Instead, it is reconstructed for each specific situation by using only the portion of the global semantic network that is relevant to the current incident and the specified objective. Consequently, the same global knowledge base gives rise to different contextual subnetworks whose structures are determined by the specific information context.

Thus, the proposed algorithm differs fundamentally from conventional graph path-search algorithms. Its result is not a path directly connecting the specified vertices but rather a contextual semantic subnetwork within which the automatic synthesis of alternative decision support scenarios is subsequently performed. The construction of such a contextual model enables the cognitive cyber range to adapt to the current situation, reduces the search space, improves the explainability of the generated recommendations, and provides the foundation for subsequent multi-agent analysis and multi-criteria ranking of the generated scenarios.

6. Multi-Agent Analysis and Scenario Ranking

As a result of applying the bidirectional contextual semantic subnetwork construction algorithm, a set of alternative scenarios describing the transition of the information system from the current state to the specified target state is generated. These scenarios represent feasible sequences of actions but do not, by themselves, determine their suitability. The same target state may be achieved through different paths that differ in required resources, risk level, execution time,

compliance with regulatory requirements, and other characteristics. Therefore, the final stage of the cognitive cyber range consists of the intelligent analysis of the generated alternatives and their multi-criteria ranking.

Unlike conventional expert systems, in which evaluation is performed according to rigidly predefined rules, the proposed architecture employs a multi-agent framework based on large language models. A fundamental feature of this approach is the separation of scenario generation from scenario interpretation. Scenario construction is determined entirely by the structure of the contextual semantic subnetwork, whereas large language models are used exclusively for analyzing the generated alternatives, explaining them, and performing their integrated evaluation. In this way, the scenario generation process remains independent of the statistical nature of large language models, while their use is confined to intelligent decision support.

Let $P = \{P_1, P_2, ..., P_n\}$ be the set of scenarios obtained in the previous stage. Each scenario represents a directed path in the contextual semantic subnetwork and contains a complete causal chain describing the transition from the initial state to the target state.

The set of agents is defined as $A = \{a_1, a_2, ..., a_m\}$, where each agent specializes in analyzing a particular aspect of a scenario. Regardless of their specialization, all agents operate on the same information context, which includes the contextual semantic subnetwork, the results of Retrieval-Augmented Generation, the digital model of the information infrastructure, and the relevant regulatory documents. This ensures the consistency of the analysis and eliminates inconsistencies arising from the use of different information sources.

The operation of each agent can be formalized by the function $F_j : P \to R$, which maps a scenario to the space of quantitative evaluations according to the agent's area of expertise. For each scenario, an evaluation vector $F(P_i) = (F_1(P_i), F_2(P_i), ..., F_m(P_i))$, is constructed, representing the results of the independent expert analysis.

The individual agent evaluations are integrated using an aggregation operator $\Psi : R^m \to R$, which produces the overall evaluation of a scenario $R(P_i) = \Psi(F_1(P_i))$.

The specific implementation of the operator Ψ is determined by the application domain and may be based on various multi-criteria decision analysis methods. In the simplest case, a linear aggregation is employed $R(P_i) = \sum_{j=1}^{m} \alpha_j F_j(P_i)$, where the coefficients α_j characterize the relative importance of the individual evaluation criteria.

In addition to a numerical score, each agent generates an explanation of its evaluation. In doing so, the large language models analyze not only the scenario itself but also the corresponding fragment of the contextual semantic subnetwork, the knowledge sources used, the relevant regulatory documents, and the logical relationships that led to the generated recommendation. This provides explainability of the obtained conclusions and enables verification of their justification.

In the proposed approach, the multi-agent analysis is considered as a separate transformation applied to the set of scenarios, $\Gamma : P \to P^*$, where $P^* = (P^{(1)}, P^{(2)}, ..., P^{(n)})$ is the ordered set of scenarios ranked according to the integrated evaluation function $R(P_i)$. The operator Γ does not generate new scenarios; instead, it performs their intelligent evaluation, explanation, and ranking according to the requirements of the current context.

Thus, the overall operation of the cognitive cyber range can be described as the composition of three interconnected operators,

$$(G,C) \xrightarrow{\Phi} G_C \xrightarrow{\Theta} P \xrightarrow{\Gamma} P^*,$$

where the operator Φ constructs the contextual semantic subnetwork, the operator Θ performs the automatic synthesis of alternative scenarios based on this subnetwork, and the operator Γ carries out the multi-agent analysis, explanation, and multi-criteria ranking of the generated alternatives.

The proposed mathematical model provides a clear functional separation between the processes of knowledge construction, scenario synthesis, and expert analysis. This makes it possible to improve individual system components independently without modifying the overall architecture of the cognitive cyber range, while also allowing different algorithms for semantic network construction, scenario discovery, and multi-agent evaluation to be employed within a unified formal decision support framework.

7. Experimental Validation

To verify the feasibility of the proposed cognitive cyber range, an experimental study was conducted using a test network infrastructure that emulates the corporate information system of a medium-sized enterprise. The use of an emulation environment made it possible to reproduce a typical architecture of a modern organization, ensure controlled experimental conditions, and repeatedly execute identical scenarios without affecting real production information resources.

To experimentally validate the proposed concept of the cognitive cyber range, a test corporate network modeling the information system of a medium-sized enterprise was employed. The objective of the experiment was to verify the correctness of the proposed algorithms for contextual semantic subnetwork construction, automatic scenario synthesis, and multi-agent ranking rather than to evaluate the scalability of the system.

The test corporate network consisted of 12 nodes and 26 network connections between them [10], [11]. It included user and administrator workstations, an external web server, a database server, an Active Directory server, a network-attached storage (NAS) server, a SIEM platform, a wireless access station, a router, a switch, and a firewall. For each node, a digital description was created containing information about the installed software, network services, operating system configuration, open ports, access privileges, deployed security controls, and known vulnerabilities. In addition, the network model was integrated with information from the CVE database, the MITRE ATT&CK knowledge base, CERT-UA recommendations, the ISO/IEC 27000 series of standards, and the organization's internal security policies. The current information context was generated using Retrieval-Augmented Generation based on a corpus of documents collected from open sources.

Based on the integrated data, a global directed semantic network of the application domain was constructed. After term normalization and semantic alignment of entities, the network contained 32 vertices and 48 directed relationships of different types describing information system assets, vulnerabilities, attack techniques, security mechanisms, and the causal relationships among them.

The experimental validation was performed for the task of responding to a simulated web server compromise incident. The initial system state was defined by the successful exploitation of a known web application vulnerability and the potential possibility of subsequent lateral movement by the attacker within the corporate network. The target state was defined as the complete containment of the incident while ensuring information confidentiality, restoring service availability, and satisfying the requirements of the organization's internal incident response procedures.

According to the mathematical model presented in the previous sections, the first stage involved constructing the contextual semantic subnetwork

$$G_C = \Phi(G, C),$$

which contained only those vertices and relationships of the global network that were directly related to the current incident, the specified objective, and the established constraints. As a result, the number of vertices was reduced to 6 and the number of relationships to 16, substantially reducing the scope of subsequent analysis without losing relevant information. The contextual semantic subnetwork contained only 6 vertices and 16 relationships, corresponding to an approximately 81% reduction in the number of vertices compared with the global model. This significantly reduced the search space for alternative response scenarios without losing the information required for decision-making.

The next stage involved applying the scenario synthesis operator $P = \Theta(G_C)$, which automatically generated a set of alternative transition paths from the current state to the target state.

After eliminating equivalent paths and scenarios that violated the established constraints, twelve feasible response alternatives were obtained. The scenarios differed in the sequence of response actions, required resources, implementation time, residual risk level, and degree of compliance with regulatory requirements. The generated scenarios were evaluated using a multi-agent framework. The framework consisted of six virtual agents that assessed the technical feasibility of the scenarios, the level of cyber risk, compliance with regulatory requirements, resource consumption, practical deployability, and the completeness of the response procedures. All agents operated on the same contextual semantic subnetwork and shared a common information corpus generated using Retrieval-Augmented Generation. The individual assessments were integrated by means of an aggregation operator, resulting in a ranked list of alternative response scenarios.

Particular attention during the experiments was devoted to the explainability of the obtained results. For each scenario, the system automatically generated a causal reasoning chain describing the logic of its construction, the semantic network nodes involved, the information sources supporting the corresponding relationships, and the regulatory documents justifying the proposed response actions. This enabled experts to trace the recommendation generation process and verify its consistency with the actual structure of the information system.

The experimental results demonstrated that the use of a contextual semantic subnetwork significantly reduces the scope of analysis compared with the use of the complete global knowledge network. At the same time, the automatic scenario synthesis mechanism generates alternative response strategies without requiring the prior programming of specific scenarios, while the multi-agent analysis ranks the generated alternatives according to the requirements of a particular decision support task.

It should be noted that the experimental evaluation was conducted on an emulated corporate network of limited scale. Therefore, the obtained results confirm the feasibility of the proposed concept and its underlying algorithms but do not claim to provide a comprehensive evaluation of system performance in large-scale corporate or critical information infrastructures. Such investigations constitute the subject of future research. Unlike conventional cyber ranges, which reproduce predefined scenarios, the proposed system automatically generates new scenarios based on the current context, up-to-date cyber threat information, and the structure of the information infrastructure, thereby creating opportunities for its application not only in cybersecurity education and training but also in operational cybersecurity management.

8. Conclusions

This paper proposes the concept of a cognitive cyber range as a new generation of decision support systems for cybersecurity. Unlike conventional cyber ranges, which are primarily

designed for simulating predefined attack scenarios and training personnel, the proposed approach enables the automatic synthesis of alternative response scenarios according to the current state of the information infrastructure, the current information context, and the specified cybersecurity objectives. This formulation makes it possible to consider the cognitive cyber range not only as a platform for cyber incident simulation but also as an intelligent decision support system.

The proposed approach is theoretically founded on the integration of a digital model of the network infrastructure, a directed semantic network, Retrieval-Augmented Generation, large language models, and multi-agent analysis into a unified information and analytical framework. The central component of the architecture is a global directed semantic network that accumulates knowledge about the information infrastructure, cyber threats, security mechanisms, regulatory requirements, and the causal relationships among them. Unlike conventional knowledge graphs, this network serves as a universal knowledge base from which a contextual semantic subnetwork relevant to a specific decision support task is dynamically constructed according to the current context.

The principal scientific contribution of this work is the development of a bidirectional algorithm for constructing a contextual semantic subnetwork. The algorithm independently generates the region of possible situation evolution and the set of conditions required to achieve the specified objective, and subsequently integrates them into a unified causal model. Unlike conventional path-search algorithms, the proposed approach does not perform direct exploration of the complete semantic network. Instead, it first constructs a contextual task model, which significantly reduces the search space, improves the explainability of the obtained results, and enables the system to adapt to changes in the information environment.

An important characteristic of the proposed architecture is the functional separation of knowledge construction, scenario synthesis, and expert evaluation. Large language models are not employed to construct the semantic network or generate response paths; instead, they perform semantic analysis, explanation, and multi-criteria ranking of the already generated alternatives. This separation combines the advantages of formal graph-based models with the capabilities of modern generative artificial intelligence, ensuring both the reproducibility of the scenario synthesis process and the high quality of scenario interpretation.

The proposed mathematical model describes the operation of the cognitive cyber range as a composition of three interconnected operators

$$(G,C) \xrightarrow{\Phi} G_C \xrightarrow{\Theta} P \xrightarrow{\Gamma} P^*,$$

where the operator (Φ) constructs a contextual semantic subnetwork according to the current task context, the operator (Θ) performs the automatic synthesis of a set of alternative response scenarios, and the operator (Γ) carries out their multi-agent evaluation, explanation, and ranking. The proposed model provides a formal separation of the individual stages of the decision support process and establishes a foundation for the independent evolution of each functional component of the system.

The results of the experimental validation on an emulated corporate network confirmed the feasibility of the proposed approach and its ability to automatically generate relevant cyber incident response scenarios without requiring programming. The use of a contextual semantic subnetwork significantly reduced the scope of analysis compared with operating on the complete knowledge network, while the application of a multi-agent framework produced an ordered set of alternative solutions together with explanations of their ranking and references to the corresponding elements of the semantic model.

The practical significance of this work lies in establishing a methodological foundation for the development of next-generation cognitive cyber ranges capable of operating both in educational

environments and in operational decision support systems for cyber incident response. The proposed architecture can be integrated with digital twins of information infrastructures, SIEM and SOAR platforms, configuration management systems, and cyber threat intelligence tools, thereby enabling the development of intelligent frameworks for the automated analysis and planning of cybersecurity measures.

Experimental validation on a test corporate network confirmed the feasibility of the proposed approach. The results demonstrated the capability for the automatic construction of a contextual semantic subnetwork, the synthesis of alternative response scenarios, and their multi-agent ranking using a unified information context generated through Retrieval-Augmented Generation (RAG).

Future work will focus on scaling the proposed approach to large-scale digital twins of corporate networks, integrating it with operational security monitoring technologies (SIEM, SOAR, and EDR), and evaluating system performance under real-world cyber incident conditions.

Further research should be directed toward developing methods for the real-time dynamic updating of the global semantic network, forecasting the evolution of contextual subnetworks, adaptively tuning the multi-agent evaluation framework, and integrating the cognitive cyber range with automated cyber incident response mechanisms. This will enable the transition from conventional decision support to cognitive cyber defense systems capable of active cybersecurity management in continuously evolving information environments.

References

[1] Kovačević, I., Katulić, F., Mamut, M. and Groš, S., 2025, June. Survey of Open Source Technologies for Building Cyber Ranges for Exercises and Trainings. In *2025 MIPRO 48th ICT and Electronics Convention* (pp. 788-793). IEEE. DOI: 10.1109/MIPRO65660.2025.11131978

[2] Lazarov, W., Schafeitel-Tähtinen, T., Squillace, J., Martinasek, Z., Coufalikova, A., Helenius, M., Gallus, P. and Fujdiak, R., 2025. Lessons learned from using cyber range to teach cybersecurity at different levels of education. *Technology, Knowledge and Learning*. DOI: 10.1007/s10758-025-09840-y

[3] Hosamo, H.H., Nielsen, H.K., Alnmr, A.N., Svennevig, P.R. and Svidt, K., 2022. A review of the Digital Twin technology for fault detection in buildings. *Frontiers in Built Environment*, *8*, p.1013196. DOI: 10.3389/fbuil.2022.1013196

[4] Lande D.V. *OSINT in Cybersecurity: A Study Guide*. – Kyiv: Engineering LLC, 2024. – 522 p. ISBN 978-966-2344-97-4

[5] Zhao, P., Zhang, H., Yu, Q., Wang, Z., Geng, Y., Fu, F., Yang, L., Zhang, W., Jiang, J. and Cui, B., 2026. Retrieval-augmented generation for ai-generated content: A survey. *Data Science and Engineering*, pp.1-29. DOI: 10.1007/s41019-025-00335-5

[6] Dmytro Lande, Leonard Strashnoy. *GPT Semantic Networking: A Dream of the Semantic Web - The Time is Now*. – Kyiv: Engineering, 2023. - 168 p. ISBN 978-966-2344-94-3, DOI: 10.5281/zenodo.14278893

[7] Sun, G., Wang, Y., Niyato, D., Wang, J., Wang, X., Poor, H.V. and Letaief, K.B., 2024. Large language model (llm)-enabled graphs in dynamic networking. *IEEE Network*, 39(4), pp.290-301. DOI: 10.1109/MNET.2024.3511662

[8] Gana, B., Palma, W., Lucay, F.A., Missana, C., Abarza, C. and Allende-Cid, H., 2025. Measuring Semantic Coherence of RAG-Generated Abstracts Through Complex Network Metrics. *Mathematics*, 13(21), p.3472. DOI: 10.3390/math13213472

[9] 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. DOI: 10.1201/9788743800989

[10] Oleksii Novikov, Dmytro Lande, and Lesia Alekseichuk. Assessment of Corporate Network Penetration Scenarios Based on Iterative Link Weight Determination Using the MRRW-PageRank Method. *Proceedings of the Workshop on Cryptology and Data Security (WCDS 2025) co-located with SMICS* 2025. Lviv, Ukraine, October 16.18, 2025. https://ceur-ws.org/Vol-4191/paper2.pdf

[11] Lande D., Novikov O., Alekseichuk L. Application of Large Language Models for Assessing Parameters and Possible Scenarios of Cyberattacks on Information and Communication Systems. *Theoretical and Applied Cyber Security, 2024*. Vol. 6 No. 1. DOI: 10.20535/tacs.2664-29132024.1.315242.