

Swarm of Virtual Experts in the Implementation of Semantic Networking

Dmytro Lande¹, Leonard Strashnoy²

¹*National Technical University of Ukraine “Igor Sikorsky Kyiv Polytechnic Institute”*

²*The University of California, Los Angeles (UCLA)*

Abstract. This article explores the concept of the "swarm of virtual experts" in the context of semantic networking, representing a methodology that combines human interaction and language models for the creation and modification of semantic networks. The article highlights the significant advantages of this approach, including flexibility, multifaceted analysis, and economic efficiency, while also discussing existing challenges, including the need for coherence and interpretability of data. An important component of the work is the mathematical formalization of the processes for creating and modifying semantic networks, which allows for a clear description of the interactions between nodes and their connections. In addition to theoretical aspects, the article presents a real-world example of applying the semantic networking methodology. Based on the analysis of data from a chosen field (in the area of cybersecurity), the processes of entity extraction and connection identification are demonstrated, illustrating the practical value of the swarm of virtual experts in generating new knowledge. The article underscores the significance and potential of using the swarm of virtual experts as an effective tool for analyzing complex relationships and forming deep semantic networks, opening new horizons in scientific research and practical activities.

Keywords: swarm of virtual experts, semantic networking, large language models (LLMs), cybersecurity, semantic network creation, data validation

Introduction

The age of digital information and artificial intelligence has introduced a new approach to organizing and processing data: semantic networking. This process involves the creation, expansion, and modification of semantic networks that capture the meaningful connections between concepts, allowing for the analysis and visualization of their structure and relationships. These networks, seemingly simple graph models at first glance, are intelligent constructs formed through the interaction between humans and large language models (LLMs). On one hand, they rely on the knowledge available in LLMs, and on the other, they draw upon the intuition and experience of the human who formulates queries and interprets the results. Semantic networking is a domain where human experience, intuition, and knowledge merge with the linguistic capabilities of modern AI systems.

In an era of rapid technological advancement, LLMs have become tools for extracting entities (concepts) and the connections between them from texts. This process underlies the creation of semantic networks that can be applied across various

fields, from big data analysis to the development of knowledge systems. However, several significant problems arise when using these models, including hallucinations, excessive or irrelevant connections, and the omission of important concepts.

To partially address these issues, an approach has been proposed that relies on multiple calls to LLMs with various prompts, roles, and models, along with the aggregation of results. This approach is based on repeated interactions with LLMs through different requests, utilizing various roles and models. It increases the likelihood of identifying relevant concepts and connections, yet raises questions about the optimal number of queries and roles, as well as the need for human verification of results.

A key role in the process of semantic networking is played by the virtual expert—an entity represented by the LLM in response to each user query. Each query can be seen as the opinion of a virtual expert, who provides network nodes (concepts), their properties, as well as the connections between these nodes and their attributes. Virtual experts operate based on the knowledge embedded in the models and the texts or documents provided by the user, transforming the LLM into a kind

of linguistic processor capable of analyzing, classifying, and expanding semantic networks.

However, the true breakthrough in semantic networking is achieved through the swarm of virtual experts—a multitude of responses to queries made at different times, on various LLM models, using different approaches to query formulation. This swarm combines the diversity and variability of opinions, providing the semantic network with additional layers of accuracy and completeness. This swarm forms a network of concepts, interactions, and relationships, which a person can then filter, aggregate, and organize, thereby enhancing the credibility and completeness of knowledge. Each query, each “expert” in the swarm, is a particle of collective knowledge that, like drops in the ocean, come together to achieve a common goal: building a more accurate and complete picture of the explored domain.

An additional innovation proposed in this article is the introduction of a validation stage involving human experts. This helps ensure that the resulting semantic networks are not only statistically sound but also practically meaningful. Such a hybrid approach—a combination of machine and human experts—can significantly enhance the accuracy and reliability of the results.

The role of the human in this process, despite the power and erudition of the LLM, remains unique and irreplaceable. It is the human, as a mentor, who determines the goal, sets the directions, and controls the quality of the created network. Their experience and contextual knowledge allow them to distinguish the significant from the trivial, identifying key nodes and their connections. In this unique symbiosis, the human not only utilizes the LLM to form the network but also directs the swarm of virtual experts toward a deeper understanding of the topic.

The “swarms of virtual experts” in semantic networking can be used for analyzing texts, reports, and data when addressing tasks such as uncovering hidden meanings and interconnections, monitoring and assessing trends, creating recommendations and making decisions, identifying and combating misinformation, generating educational materials and content, and analyzing large volumes of documents.

The swarm of virtual experts helps identify semantic connections that may not be immediately obvious. They analyze texts at

different levels of abstraction and can suggest new nodes and connections that reflect hidden dependencies within the data.

By utilizing a body of texts, such as reports, news articles, or academic publications, the swarm of virtual experts assists in highlighting key trends and dynamics of change, representing them in the form of a semantic network. This allows tracking the evolution of topics and forecasting future changes.

The swarm of experts can be employed in the analysis of numerous data sources, revealing patterns and connections that are useful for decision-making. For instance, in business or politics, strategies and recommendations can be substantiated based on the analysis of reports and forecasts prepared by AI using historical data.

Semantic networking based on a swarm of virtual experts allows for the analysis of text fragments to identify contradictory or potentially false information. This is particularly important when working with social media, where the swarm of virtual experts can assess the credibility of posts and suggest the most likely true connections.

Swarms of virtual experts can process and simplify complex materials, highlighting key concepts and relationships that are easier for the audience to grasp. This can be useful for preparing training courses or self-study materials.

When processing and analyzing large volumes of information, such as reports or research publications, the swarm of virtual experts enables the systematization of knowledge by identifying main themes and concepts, creating coherent connections and structures for further analysis.

In semantic networking, a swarm of experts can incorporate new information into the network and rapidly restructure the conceptual framework, improving the network as fresh data becomes available. This makes it flexible and relevant, capable of adapting to changing conditions.

The use of a swarm of virtual experts in semantic networking addresses these tasks, extracting maximum value from complex information structures and aiding in making informed conclusions across various fields.

In recent years, there has been a rapid increase in the use of artificial intelligence and machine learning models in cybersecurity for threat detection and data analysis. The

implementation of large language models (LLMs), such as GPT, has revolutionized natural language processing tasks, including the extraction of concepts and relationships from unstructured data. Several studies have demonstrated the potential of LLMs in detecting phishing attacks, analyzing social engineering tactics, and identifying patterns of cyberattacks [1, 2].

However, a significant issue remains—the tendency of LLMs to "hallucinate," where the model generates irrelevant concepts or fails to identify key relationships between important entities in the field of cybersecurity [3].

The method of the swarm of virtual experts aims to address these issues by employing multiple iterations of prompts to one or several LLMs to extract concepts and relationships. There are several studies dedicated to the extraction of concepts and relationships using LLMs and other artificial intelligence methods. For instance, the study [4] examined methods for the automatic construction of semantic networks from text data using LLMs.

Recent research indicates that such methods can be applied to improve accuracy in other fields, such as medical diagnostics and financial analysis [5-7]. An important topic is the methodology for using role models to expand the diversity of concepts. The research [8] suggests methodologies similar to the "swarm of experts," but without the application of repeated prompts with different roles. Another crucial direction is the inclusion of the human factor during the verification stage.

The approach based on the "swarm of virtual experts" resonates with the concepts of a swarm of intelligent agents discussed in the work [9]. While previous studies focused on creating specialized agents for specific tasks, the swarm approach offers extensive opportunities for flexible adaptation to complex scenarios.

Issues related to the generation of accurate and relevant entities during the repeated execution of prompts are also raised in several articles, such as [10], which examines the impact of cognitive biases and prompt reuse on the quality of extracted data.

The Concept of the "Swarm of Virtual Experts"

The term "swarm of virtual experts" refers to a group of responses obtained from a large language model (LLM) based on a series of prompts directed at exploring a particular domain of knowledge. Each individual response from the LLM serves as the opinion of a "virtual expert," reflecting the model's interpretation of the data for a specific query. However, a single response may only contain a part of the possible information, especially when the prompt requires subjective interpretation or complex analysis.

To enhance the accuracy and completeness of the semantic network, specialists employ the "swarm" approach by repeatedly querying the LLM with various prompts and even across different models to accumulate a wide array of perspectives and interpretations. These expert responses collectively form a "swarm," which can be aggregated to identify key nodes (concepts) and the relationships between them, as well as their properties that most frequently align across all responses. Statistical processing of numerous responses allows a human mentor to refine the semantic network by highlighting the most significant concepts and connections. Thus, the "swarm of virtual experts" serves as a powerful tool in constructing comprehensive and accurate semantic networks, offering a symbiosis of human analysis and AI capabilities.

In the context of the "swarm of virtual experts," the consistency of the actions of virtual experts is a key aspect that allows for the creation of coherent and reliable semantic networks. Although each "expert" represents a separate prompt execution, several mechanisms contribute to the consistency of their output and help create cohesive results:

The frequency of recurring nodes and connections in responses to similar prompts allows for the identification of more significant and stable concepts. The more often a specific node or connection appears, the higher its probability of significance for building the network.

Structuring results based on a hierarchy of concepts helps to prioritize key ideas within the network. This allows for the identification of higher-level nodes that encompass related but less significant elements.

Modern LLMs often retain information about previous queries within a session, helping the models maintain a certain consistency in responses. Additionally, using cumulative prompts that refine or develop previous ones enables clarifying connections and adjusting network nodes.

A human serves as the mentor of the swarm of virtual experts, posing a sequence of prompts and tracking the responses. This enables adjustments to questions and the addition of new ones while verifying hypotheses if necessary. Human analysis helps to avoid random responses by discarding irrelevant connections and reinforcing the most significant ones.

Weights can be assigned to different nodes and connections, reflecting their significance based on frequency or uniqueness. Connections with high weights are considered more reliable, contributing to the creation of a more accurate and coherent semantic network.

These mechanisms allow the "swarm of virtual experts" to be more than just a random collection of opinions; they facilitate a purposeful process of generating reliable knowledge. Through the consistency of actions, the swarm ultimately constructs a complete and stable semantic network where humans and AI operate as a unified system.

To increase the diversity and accuracy of extracted connections, several methodologies have been proposed, including entity and relationship extraction using LLMs, the use of roles, the application of various LLM models, result aggregation, and human verification.

The method is based on the sequential extraction of concepts and relationships between them from the text. This employs the same prompt, which is reformulated and executed multiple times to minimize the influence of repetitive responses. Repeated prompt execution helps "reset" the system, increasing the likelihood of finding more relevant connections. Frequent resets enable the generation of more diverse results, avoiding repetitions and enhancing the chances of extracting unique and relevant relationships.

To diversify responses and minimize errors, a role-switching methodology has been introduced. Each query to the LLM is formulated with a specific role in mind, for instance, from the perspective of an expert in a

certain field, which aids in extracting new contextually significant connections.

The variety of models allows for a reduction in the influence of specific model characteristics on the outcome, thereby enhancing the reliability of extracted data. Each model generates its unique entities and connections, improving the completeness of the resulting semantic networks.

After collecting data from various models and roles, aggregation occurs. Connections are assigned weights corresponding to their frequency of appearance across different networks. Connections with low frequencies may be excluded to eliminate false or insignificant data.

In the final stage, the inclusion of a human expert for network verification is proposed. This step is essential for minimizing errors made by the LLM, especially in cases where important connections have been missed or non-existent ones added.

Role definition is an important aspect of the "swarm of virtual experts" methodology, which enables the extraction of entities and connections in various contexts and from different perspectives. Role differentiation provides a multi-faceted understanding of the data and helps create a semantic network capable of integrating diverse interpretations and approaches to the same information.

Roles Defined by Humans

A human, acting as a mentor, can determine key roles based on the goals of the analysis and the specifics of the subject area. These roles can represent various perspectives, approaches, and contexts, such as expert, analytical, and evaluative roles. The expert role focuses on extracting entities and relationships with an emphasis on scientific or technical terms, approaches, and concepts.

The analytical role involves analyzing relationships in terms of causal dependencies, sequences of events, etc. The evaluative role is concerned with identifying entities and relationships with a focus on priorities, significance, and weights. By assigning specific roles, the human shapes the direction of the analysis, indicating which concepts and relationships are most important to highlight. This helps the swarm focus on

extracting information within the given perspectives and interpretations.

Roles Suggested by LLMs

Large language models can also propose their own roles and contexts that a human might not consider. LLMs can recognize and suggest roles based on the structure of data or text, the frequency analysis of concepts, or even based on embedded semantic categories. For instance, a model may suggest historical, social, and cross-disciplinary roles. The historical role interprets information with an emphasis on historical context and the evolution of concepts. The social role focuses on the connections between the social aspects of data, analyzing information within the framework of interactions and their significance in society. The cross-disciplinary role provides data extraction from the perspective of the interrelationship of different disciplines or approaches.

LLMs can automatically identify such roles using their built-in knowledge and analytical capabilities, opening up the swarm of virtual experts to see data from unexpected and innovative perspectives.

Role Structure in the "Swarm of Virtual Experts"

The application of various roles allows virtual experts to form a network of entities and relationships that encompass diverse approaches and viewpoints. This multifaceted analysis makes the semantic network more coherent and flexible. The human, as a conductor, defines the overall direction and adjusts roles as necessary, while LLMs, acting as virtual experts, complement the process by initiating and proposing new, potentially significant roles. Thus, the distribution of roles between humans and LLMs not only ensures depth in the analysis but also allows for the creation of semantic networks capable of more accurately and comprehensively reflecting complex interrelationships in the studied context.

One of the key questions is determining the optimal number of roles for virtual experts. An excessive number of roles may lead to a decrease in the quality of the network due to

informational noise, while too few roles may fail to identify all possible connections. A statistical analysis approach for the growth of new connections can be employed for this purpose. In the proposed model of the "swarm of virtual experts," evaluating the number of roles and their confirmation by a human are critical elements for enhancing the accuracy and reliability of the semantic networks created with the help of LLMs.

Assessing the Number of Roles for Virtual Experts

The roles of experts in the model serve to increase the diversity of responses. Each role sets a different context or perspective from which the LLM can view concepts and their connections. This helps minimize the impact of possible errors or omissions that may arise from queries from a single role. Finding the optimal number of roles is crucial. Excessive increases in the number of roles may lead to the generation of redundant, irrelevant connections, while too few will not provide sufficient diversity in responses. The optimum can be sought through statistical analysis: after several iterations of prompts with different roles, one can observe at what number of roles the most stable and relevant network of connections emerges. The role-counting method may be based on the concept of convergence. If the frequency of new concepts and connections begins to decrease with the addition of a new role, this may signal that most relevant concepts have already been accounted for and that further role additions are ineffective. This can be formalized by monitoring the growth of new concepts with each additional role and establishing a threshold at which adding new roles becomes statistically insignificant.

Validation of Results by Humans

The role of the human in the model serves as the final check of the aggregated semantic network. After the LLM has processed numerous prompts from different roles, the human expert performs validation by analyzing the connections between concepts, their weight, and relevance. The human expert helps eliminate LLM hallucinations, irrelevant or minor connections, and omissions of

important links. For convenience in validating results, interactive interfaces can be implemented to allow the human to observe the weights of connections (obtained based on their frequency identified by virtual experts) and adjust them as necessary. Thus, the system gains an additional degree of reliability through the inclusion of the human expert, who confirms or corrects the results provided by the virtual experts.

Mathematical Formalization

Let:

- $S = \{s_1, s_2, \dots, s_n\}$ represent the set of entities extracted from the text;
- $R = \{r_{ij}\}$ represent the set of relationships between entities, where r_{ij} denotes the relationship between entities s_i and s_j ;
- P represent the set of prompts used for relationship extraction;
- $M = \{M_1, M_2, \dots, M_k\}$ represent the set of different LLM models;
- $W(r_{ij})$ represent the weight of the relationship between entities s_i and s_j , which depends on the frequency of the relationship's occurrence.

For each prompt $p \in P$ and each model

$M_i \in M$:

- The sets of entities S_p^l and relationships R_p^l are extracted.
- The relationship weight is updated: $W(r_{ij}) = W(r_{ij}) + 1$, if $r_{ij} \in R_p^l$.

Threshold θ

A relationship r_{ij} is included in the final network if its weight $W(r_{ij}) \geq \theta$, where θ is defined as a statistical threshold based on the number of virtual expert iterations and the frequency of relationships.

Determining Optimal Threshold θ

The following methods are proposed for determining the optimal threshold θ :

- Let $f(r_{ij})$ denote the frequency of the relationship r_{ij} in the results across different prompts and models.
- Define a frequency distribution $P(f)$ for relationships, representing the probability of relationships appearing with different frequencies.

To select an optimal threshold θ , a significance criterion based on the number of iterations and the variance of relationship frequencies can be used:

$$\theta = \mu f + \alpha \sigma_f,$$

where:

- μf is the mean frequency of relationship occurrences;
- σ_f is the standard deviation of frequencies;
- α is a significance coefficient (e.g., $\alpha = 1.5$ to filter out less significant relationships).

This threshold helps exclude relationships that are random and lack substantial weight.

Estimating Optimal Number of Virtual Expert Invocations N

The number of virtual expert invocations N directly affects the reliability of the network construction. The method of stabilization of relationship frequencies can be used to assess the optimal number of invocations. Let:

- N represent the number of iterations;
- $f(r_{ij})$ represent the frequency of relationship r_{ij} after N iterations.

If the increase in new relationships $\Delta f(r_{ij})$ becomes negligible (e.g., $\Delta f(r_{ij}) < \varepsilon$, where ε is a small value), stability is achieved. Thus, the optimal N can be determined based on the analysis of relationship growth:

$$N = \min \left\{ N' : \frac{1}{|R|} \sum_{r_{ij} \in R} \Delta f(r_{ij}) < \varepsilon \right\}.$$

Formalization for Role Quantity and Verification

1. Let R denote the number of roles from which data is requested from LLM. For each role i , we obtain a set of concepts C_i and a set of relationships between them S_i .

2. Aggregated set of relationships: $S = \bigcup_{i=1}^R S_i$, assigning each relationship a weight $w(s)$, equal to its frequency across roles.
3. Introduce a threshold value τ to determine which relationships are considered significant. For example, if a relationship appears in less than τ proportion of roles, it is discarded: $S^* = \{s \in S \mid w(s) \geq \tau\}$.

The necessary number of roles R^* can be determined by observing the growth in new relationships with additional roles. If the growth rate of new concepts and relationships becomes negligibly small (e.g., $\frac{\Delta S_i}{\Delta R} \rightarrow 0$), increasing the number of roles can be halted.

4. After network construction S^* , a human expert conducts a final verification, confirming or discarding suggested relationships, which ensures the final adjustment of the network.

Determining Optimal Frequency Threshold

To determine the optimal frequency threshold θ , the following approach can be applied:

- Conduct a series of experiments with varying values of θ , starting from minimal values (considering all relationships) and increasing the threshold.
- A relationship with a weight of 1 may appear most frequently with a small number of experts. To avoid this, an experimental analysis with varying roles and thresholds θ can identify the point where the addition of new relationships becomes negligible.
- This value will be considered optimal.

Example from the Field of Cybersecurity

Suppose we need to construct a causal network corresponding to the concept of a "phishing attack." First, we will build this network without applying the "swarm of virtual experts," and then we will apply the

proposed concept, relying on the role-based approach and multiple prompt iterations.

In both cases, we will use hierarchical refinement of concepts that lead to a phishing attack. As the basic query, let's consider

Prompt: List the main reasons (indicated by short keywords) that lead to a phishing attack. Present the results as a list in the format "reason; phishing attack", for example, "human error; phishing attack"

LLM system response:

- Human error; phishing attack
- Weak passwords; phishing attack
- Lack of awareness; phishing attack
- Unprotected email systems; phishing attack
- No multi-factor authentication; phishing attack
- Outdated software; phishing attack
- Social engineering; phishing attack
- Poor security policies; phishing attack
- Suspicious links; phishing attack
- Insufficient spam filters; phishing attack

At the second level, similar prompts are generated but directed toward the concepts already obtained through LLM, for example:

Prompt: Identify the main reasons (indicated by short keywords) that lead to Human Errors leading to a phishing attack . Present the results as a list in the format "reason; Human error"

The resulting simple network is visualized using Gephi software (Figure 1).

To implement a small swarm of virtual experts, we first use an LLM to identify the primary expert roles.

Prompt: Specify the roles of experts involved in analyzing a phishing attack, such as "organization manager" or "system administrator"

We obtain the following set of roles:

- Organization Manager
- System Administrator
- Security Analyst
- Incident Response Team Member
- Network Administrator
- IT Support Specialist
- Forensic Analyst
- Compliance Officer

- User Awareness Trainer
- Legal Advisor
- Threat Intelligence Analyst
- Data Protection Officer
- Software Developer
- Chief Information Security Officer (CISO)

From which, through expert judgment, we retain only three roles:

- System Administrator
- Security Analyst
- Incident Response Team Member

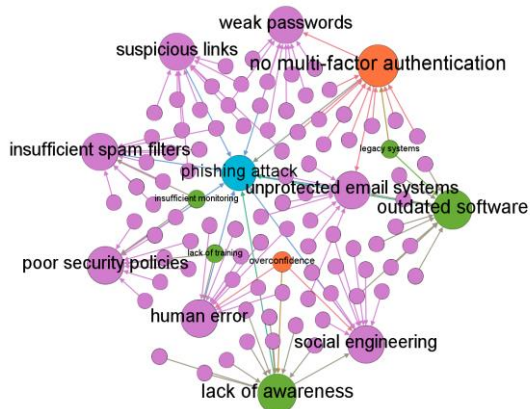


Figure 1 – Network Obtained through Simple Hierarchical Query Input

On behalf of these roles, each of the prompts mentioned above is entered twice. Example of a prompt:

Prompt: List the main reasons (indicated by short keywords) that lead to a Phishing Attack from the perspective of a System Administrator. Present the results as a list in the format 'reason; Phishing Attack', for example, 'human error; phishing attack'

In the process of aggregating results, we do not consider relationships whose frequency in the responses is less than 2. Additionally, we involve a human expert in obtaining the final result. Figure 2 presents the network obtained using a swarm of virtual experts.

Comparing Networks Obtained with and without the Participation of a Swarm of Virtual Experts Using ChatGPT

We will conduct this evaluation by also involving both real and virtual experts. To engage a virtual expert, we will write a prompt

that we will provide to another LLM—DeepSeek:

Prompt: I am interested in the questions of statistical reliability and robustness of the networks. Which one is more reliable and robust for decision-making.

Next, we introduce the set of relationships in the network, obtained through simple hierarchical query input (Network 1) and the network generated using the swarm technology of virtual experts (Network 2).

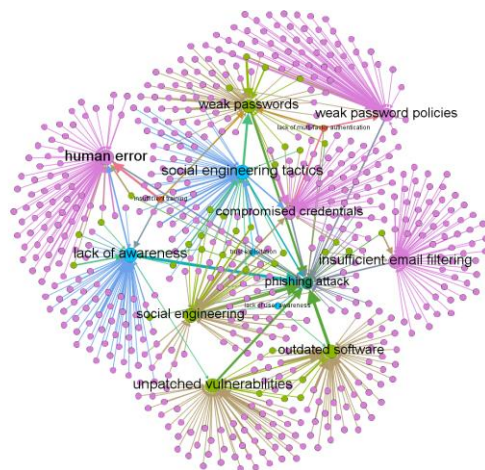


Figure 2 – Network Obtained Using a Swarm of Virtual Experts

Analyzing the structure and content of the networks, LLM - DeepSeek provides the following response:

Conclusion:

Network 2 appears to be more reliable and robust for decision-making as it is more focused on specific aspects such as email security and multi-factor authentication. This reduces the risk of redundancy and repetition in links, increasing statistical reliability and robustness of the data.

If you need a more detailed and focused network, **Network 2** will be the better choice.

Problems and Limitations

The "swarm of virtual experts" methodology, despite its usefulness, faces a number of problems and limitations that are important to consider for the correct application of this approach. Below are the main challenges and limitations:

Limited Coherence and Contradictory Responses

Since each response from the LLM is generated independently, contradictions may arise between different "virtual experts." This complicates the construction of a coherent network and requires additional filtering or aggregation. Often, human intervention is necessary to resolve contradictions, determine relevant connections, and exclude less accurate responses.

Dependence on Prompt Quality

Prompts play a key role in the "swarm of virtual experts," and the quality of responses directly depends on their clarity and focus. Poorly formulated prompts can lead to the acquisition of irrelevant information or the loss of important connections. This demands a high level of skill and expertise from the individual formulating the questions.

Long-Term Consistency Issues

When using different LLM models or querying at different times, responses may vary significantly, especially if updates occur in the data or algorithms of the models. This can lead to difficulties in maintaining the consistency of results, reducing the reliability and predictability of the information obtained.

Complexity of Data Interpretation and Control

In a scenario where multiple virtual experts offer their viewpoints and interpretations, it becomes challenging for a human mentor to process and interpret all available data. Without well-developed methods for aggregating information, it may be difficult to identify meaningful connections among a large volume of suggestions and build a coherent semantic network.

Issues with Network Robustness and Completeness

Although the swarm method can enhance the completeness of the semantic network, it may still be incomplete, particularly when the swarm encounters rare or difficult-to-define concepts. Additionally, there may be omissions in critically important connections, necessitating further cycles of queries and improvements in construction methods.

Potential Information Distortion

The LLMs on which the virtual experts are based are trained on limited and historically

conditioned data, which can lead to the dissemination of biases or distorted information within the semantic network. This limitation is especially concerning if the network's results are used for decision-making.

Difficulty in Updating and Scaling the Network

As new data and concepts are added, the network can become cumbersome and challenging to update. This requires the swarm of virtual experts to exhibit high adaptability and the ability to adjust the existing network without sacrificing quality. However, the adaptability of LLMs regarding new data does not always meet the requirements for rapid and accurate network updates.

Advantages

The "swarm of virtual experts" methodology offers a range of advantages that make it promising for the creation and analysis of semantic networks. Here are the main benefits:

Flexibility and Adaptability

The swarm of virtual experts can swiftly adjust to changes in context and requests. By modifying prompts and interacting with different LLMs, the analysis can be quickly tailored to specific tasks, highlighting necessary data aspects or focusing on particular details.

Multifaceted Analysis

Virtual experts are capable of generating responses based on various roles and contexts, allowing the same subject to be examined from different perspectives. This promotes the construction of a deeper and more holistic semantic network that incorporates diverse interpretations and approaches. As a result, the network becomes more informative and comprehensive.

Automation and Acceleration of the Process

Using LLMs to generate nodes and relationships automates the network construction process, significantly saving time. Instead of involving numerous human experts, a single operator can leverage the swarm of virtual experts to perform complex analytical tasks, thereby increasing productivity and the speed of analysis.

Capacity to Process Large Volumes of Data

Thanks to the parallel operation of virtual experts, large amounts of information can be effectively processed, extracting the most significant relationships and concepts. This is particularly useful in situations that require analyzing extensive datasets or a large number of documents within a short timeframe.

Enhanced Completeness and Accuracy

The swarm can refine and augment semantic networks by requesting additional information on contentious or incomplete aspects. Through the statistical aggregation of frequent nodes and relationships, the accuracy and completeness of the network are improved as the swarm gradually identifies the most significant elements while eliminating random and less important ones.

Continuous Updates and Data Relevance

The swarm of virtual experts allows the network to remain current by utilizing fresh queries and processing new data. This is especially valuable in rapidly changing fields, where the relevance of information is critical. The swarm can be "reconfigured" to analyze new data and identify relevant relationships, helping to prevent the network from becoming outdated.

Economic Efficiency

Compared to involving a large number of human experts, the swarm of virtual experts significantly reduces costs. A single operator, using LLMs and the swarm methodology, can perform the same analytical work that previously required the participation of many specialists, making the method more accessible and economically viable.

Support for Human Creativity and Expertise

While the swarm automates a significant portion of the work, it can also support human creative and analytical processes. Virtual experts can generate ideas that inspire new approaches and solutions, as well as provide additional information that enables more informed decision-making.

Conclusions

The proposed "swarm of virtual experts" method offers a novel approach to building semantic networks using LLMs by combining

multiple queries, roles, different models, and aggregating results. This approach partially addresses issues of hallucinations and redundant connections through result aggregation. The introduction of a statistical threshold and the evaluation of the necessary number of runs of virtual experts enhance the accuracy and reliability of the network. Multiple queries with different roles increase the likelihood of finding relevant concepts and their connections. The significance of the role concept also manifests in the ability to improve the accuracy and diversity of the network. Statistical analysis of the increase in new connections helps determine the optimal number of roles, reducing the likelihood of information noise. It is important to note that the coherence of the swarm's actions is achieved not only through multiple prompt runs but also through human participation and the internal mechanisms of LLMs, which maintain the context of queries within a single session.

The "swarm of virtual experts" methodology represents a promising tool for building, analyzing, and modifying semantic networks by leveraging LLMs as "virtual experts." The main advantage of this methodology is its ability to automate analysis processes and accelerate the retrieval of relevant information, effectively addressing a wide range of tasks, from processing large datasets to creating complex, multifaceted models.

The strengths of the approach include high flexibility, cost-effectiveness, and the ability to adapt to rapidly changing conditions. The swarm of virtual experts enables the generation of in-depth, multifaceted views on analyzed concepts while ensuring high completeness and relevance of data. However, the method also faces several challenges, such as the difficulty of maintaining coherence, the need for careful prompt tuning, and periodic human intervention to resolve contradictions and manage the network.

It is important to note that while the method significantly reduces the need for numerous human experts, it does not eliminate it completely. Human participation remains critically important at the stages of task formulation, result filtering, and quality control of the network. This hybrid approach, where humans act as mentors and operators,

helps overcome the weaknesses of the methodology, such as potential inconsistencies in responses and limited interpretability of some data.

The proposed approach bridges the gap between automated AI-based analysis and human expertise, providing a reliable solution for processing complex data in the field of cybersecurity. Future research could include expanding the range of roles used in generating virtual experts, testing the approach on various cybersecurity report datasets, and improving the validation process with human involvement through the integration of active learning methods.

The prospects for the development of the "swarm of virtual experts" are quite broad, especially considering the current development of LLMs and their potential use in semantic analysis. The future of this methodology lies in the development of enhanced data aggregation and filtering mechanisms to increase network coherence and automatically resolve contradictions among the responses of virtual experts, creating new interaction models between humans and LLMs, where humans can more flexibly manage swarms of virtual experts by setting precise criteria for which responses should be considered meaningful and reliable; improving role and context alignment mechanisms, allowing for more accurate and comprehensive analysis without significant human intervention; and expanding the use of swarms in specific fields, such as medical analytics, financial monitoring, or social research, where the need for data processing and analysis of complex relationships demands high completeness and accuracy.

Further research could focus on optimizing the method, developing more precise aggregation criteria, and automating verification procedures with the involvement of LLMs. Based on the presented results, this method can be useful for building semantic networks in information retrieval tasks, big data analysis, and creating cognitive models. The proposed approach can be applied in big data processing, knowledge analytics, and the development of expert systems.

References

- [1] Gallagher, Sean, Ben Gelman, Salma Taoufiq, Tamás Vörös, Younghoo Lee, Adarsh Kyadige, and Sean Bergeron. Phishing and Social Engineering in the Age of LLMs. In *Large Language Models in Cybersecurity: Threats, Exposure and Mitigation*, pp. 81-86. Cham: Springer Nature Switzerland, 2024. DOI: 10.1007/978-3-031-54827-7_8
- [2] Lande, Dmitry; Strashnoy, Leonard. Hierarchical Formation of Causal Networks Based on ChatGPT. SSRN preprint (May 8, 2023). - 13 p. DOI: 10.2139/ssrn.4440629
- [3] Gabrijela Perković; Antun Drobnjak; Ivica Botički. Hallucinations in LLMs: Understanding and Addressing Challenges. In 2024 47th MIPRO ICT and Electronics Convention (MIPRO) DOI: 10.1109/MIPRO60963.2024.10569238
- [4] 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
- [5] Akshay Goel, Almog Gueta, Omry Gilon, et al. LLMs Accelerate Annotation for Medical Information Extraction. Proceedings of the 3rd Machine Learning for Health Symposium, PMLR 225:82-100, 2023. URL: <http://dblp.uni-trier.de/db/conf/ml4h/ml4h2023.html>
- [6] Rian Dolphin, Joe Dursun, Jonathan Chow, Jarrett Blankenship, Katie Adams, Quinton Pike. Extracting Structured Insights from Financial News: An Augmented LLM Driven Approach (2024). Preprint arXiv, arXiv: 2407.15788. DOI: 10.48550/arXiv.2407.15788
- [7] Wiest, Isabella C., et al. "LLM-AIx: An open source pipeline for Information Extraction from unstructured medical text based on privacy pre-serving Large Language Models. medRxiv (2024): 2024-09. DOI: 10.1101/2024.09.02.24312917
- [8] Kommineni, Vamsi Krishna, Birgitta König-Ries, and Sheeba Samuel. From human experts to machines: An LLM supported approach to ontology and knowledge graph construction. arXiv

preprint arXiv:2403.08345 (2024).DOI:
10.48550/arXiv.2403.08345

- [9] Steenstra, I., Nouraei, F., Arjmand, M., & Bickmore, T. W. (2024). Virtual agents for alcohol use counseling: exploring LLM-powered motivational interviewing. arXiv preprint arXiv: 2407.08095. DOI: 10.48550/arXiv.2407.08095
- [10] Lim, S., Kim, Y., Choi, C. H., Sohn, J. Y., & Kim, B. H. (2024). ERD: A Framework for Improving LLM Reasoning for Cognitive Distortion Classification. arXiv preprint arXiv: 2403.4255. DOI:10.48550/arXiv.2403.14255