International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING

ISSN:2147-6799 www.ijisae.org Original Research Paper

Cyber Threat Intelligence Automation Using AI in the Financial Sector

Chaitanya Appani

Submitted: 05/10/2024 **Revised:** 15/11/2024 **Accepted:** 26/11/2024

Abstract: The cyber threats affecting the financial sector have become more advanced and need swift, automated threat intelligence. This article dwells upon the power of Artificial Intelligence, specifically Natural Language Processing (NLP) and Knowledge Graphs (KGs) to change Cyber Threat Intelligence (CTI) in banking. We experiment with entity-relation extraction, report generation and graph-based correlation with Large Language Models (LLMs). Such methods as AGIR, AttacKG, and K-CTIAA are techniques that automate CTI analysis with substantial improvements in performance. Experimental findings: The F1-scores are improved, and the time used in report generation reduced by up to 40 percent. In our research work, we have shown how AI-based CTI can help provide real-time structured threat information to empower FI to proactively reduce cyber risk effectively.

Keywords: Refactoring,

I. INTRODUCTION

The financial sector requires cybersecurity more than ever, as a malfunction can lead to an overall economic failure. Conventional Cyber Threat Intelligence (CTI) processes are based on the manual analysis of textual information, inaccurate, and slow. The emergence of AI, particularly NLP and Knowledge Graphs (KGs) has created new horizons in the automation of CTI.

The paper at hand explores the potentials of these technologies to convert unstructured threat reporting to structured actionable intelligence in real time. We consider banking applications, where fast speed, accuracy and interpretability are of the essence. Our effort can be seen as a synthesis of recent efforts and a coherent picture of how AI is used to further automate CTI throughout the financial ecosystem.

II. RELATED WORKS

NLP and Natural Language

Automation of Cyber Threat Intelligence (CTI) reporting is becoming more widely adopted as cyber threats exist in sheer scale and speed. The next breakthrough in this field is the introduction of Natural Language Processing (NLP) and Natural Language Generation (NLG) into CTI processes.

One example of this kind of innovation is the use of template-based NLG and large language models (LLMs) such as ChatGPT to automate the generation of reports in the AGIR framework [1].

This system helps not only to increase the accuracy of reports but also helps to save time of security analysts more than by 40 percent which becomes a significant step in terms of operational efficiency [1]. NLP in CTI is not a bed of roses. Domain specialty of CTI, as well as the variety of linguistic patterns across data sources, frequently makes traditional NLP methods useless in terms of modeling complex links between threat entities [6].

In reaction, current NLP-built CTI architectures focus on the integration of ordered databulary with textual entries to create semantically deep results. LLMs have currently shown significant abilities to detect semantic relations throughout security writings and formulate them in machine-readable format [2][4].

Other than report generation, NLP also plays a role in CTI knowledge extraction. An example would be the K-CTIAA model, which employs pretrained language models together with knowledge graphs to extract threat actions against unstructured text [8]. It is novel because it alleviates the noise of knowledge with a visibility matrix and improves self-attention. In real-world threat scenario, this hybrid NLP-KG model achieves high performance with an F1 score of 0.941 [8].

Knowledge Graph Construction

Knowledge Graphs (KGs) have cropped up as a key enabling technology in the organization of threat intelligence to enable automated reasoning and contextual analysis.

Such transformation in the financial industry is very helpful because the industry is heavily dependent on

Lead Information Security Engineer
Mastercard Inc.
O'Fallon, MO

textual information provided by news and reports. Recently, it has been shown that KGs built over financial news via NLP pipelines and ChatGPT APIs can be used to expose conditional dependency between institutions, which points to systemic risk [7].

AttacKG is another promising model that aims at automatic extraction of attack techniques and behaviors given CTI reports. By comparison, AttacKG constructs technique knowledge graphs (TKGs) by fusing intelligence in multiple reports, which accurately detects more than 28,000 attack techniques and 8,393 distinct indicators of compromise (IoCs) [3].

This improves the detection of Advanced Persistent Threats (APTs) and helps to reassemble complex attack chains - exactly the types of capabilities that are particularly useful in the high-risk environments such as banking and finance. Scalability and entity disambiguation are also issues of concern in developing these graphs.

Such methods as LLM handle them by relying on few-shot learning and fine-tuned GPT models to extract entities, topic classification, and semantic triples formation, thereby avoiding large annotated datasets [4]. The resulting knowledge graphs do not only enable automated analysis but also possess the property of being interpretable and audit able -a must have in compliance-driven industries.

Ensuring quality of KGs is a topic of current research. A solution is proposed in [9] as Adaptive Joining of Embeddings (AJE) model, which dynamically chooses the best combinations of embeddings to guarantee rationality and completeness of KG triples. AJE approach considerably enhances F1 and accuracy scores on dataset-specific to cybersecurity (e.g., CS13K: 91.3% accuracy), which adds weight to the efficacy of threat modeling via KGs.

Threat Intelligence

Cyber threat is especially vulnerable in the case of financial institutions because of their interdependent structures and the confidential data traffic. Consequently, AI-powered CTI usage has proliferated, with both NLP and KGs incorporated to provide real-time threat detection and mitigation that is dynamic.

AI enables that by automating entity extraction, risk inference and cross-correlation of threat indicators across heterogeneous sources [5]. CTI must be accurate as well as timely in high-stakes settings, such as in the banking sphere. Open-domain corpora trained traditional machine learning models can have a hard time with domain-specific language in the cybersecurity field.

By incorporating KGs with LLMs pre-trained on custom threat intelligence datasets, e.g., in the K-CTIAA paradigm, one may achieve a much higher accuracy of extraction and contextual awareness [8]. The combination of ontologies and semantic web tools with AI, allows a common representation of schemas across financial institutions and Security Operations Centers (SOCs).

It is possible to augment these KGs using deep learning techniques to identify attack vectors hidden in plain view or anticipate an evolution of threats [5]. The inter-institutional advantages of such integration are especially relevant in the field of finance; wherein systemic risks can be overcome through real-time inter-organizational cooperation.

Financial CTI is also backed by AI in terms of risk modelling and prediction. With the help of NLP-generated knowledge graphs over financial news, researchers can discover both temporal and causal patterns between institutions and events. As an example, the systemic risk of the analysis of the largest U.S. banks in 2016 was low according to the textual correlation graphs, which proves the potential of this method in macro-risk analysis [7].

CTI Automation

There are thus far a number of obstacles to the complete automation of CTI particularly in financial applications despite the major strides that have been made. Semantic complexity of cyber threat language is one of them. Terminologies might differ among vendors, threat actors, or even industries to cause inconsistency and make automated extraction and interpretation challenging.

To cope with those changes, NLP systems require regular retraining or fine-tuning, which requires high-quality labeled data, which is often limited in the area of cybersecurity [6]. The other issue is the loyalty and trustworthiness of the AI intelligence. Although the models, such as AGIR or AttacKG, show good results according to the metrics, there is still a chance of hallucinations, misattribution, or context neglect.

Hallucinations, in which models produce realistic yet false information, are of critical concerns in security decision-making [1][3]. To ensure that trust in AI-generated CTI is not compromised, it is necessary to have in place strong validation mechanisms such as human-in-the-loop schemes as well as cross-validation with independent data sources.

Annotated open-source threat datasets also are lacking, which further hurts supervised learning methods. Few-shot or zero-shot models help with that, however, performance is task and entity density dependent [4]. Moreover, the existing models tend to neglect the latent information in text, which can be either long storytelling or an informal text- an aspect that can be enhanced with the use of better semantic parsing approaches.

Ethical aspects of AI, related to cybersecurity, cannot be overlooked. Model bias, automation abuse, and unexplainable decision-making can have disastrous effects. Regulatory frameworks in both cyberscurity and finance are increasingly calling, transparency, explainability, traceable logic in knowledge graphs, and annotated threat chains [10].

IV. FINDINGS

Intelligent Parsing via NLP

The financial industry is becoming dependent on cybersecurity measures that are not only reactive in nature, but predictive. Conventional Cyber Threat Intelligence (CTI) operations rely much on human-based analysis of unstructured reports, thereby hindering scalability and promptness. Automation powered by NLP has become a revolutionary technique, enabling the systems to consume, label, and curate actionable information out of talkative threat reports.

The suitability of template-based generation with LLM summarization to automate CTI reporting is proven by recent research such as AGIR [1]. AGIR adopts a two-stage design which incorporates structural templates and LLM-based completion. The fidelity to ground truth data is high and reports produced by AGIR have a recall of 0.99.

Mathematically, consider the **precision-recall tradeoff** captured by the F1-score:

$$F1 = 2 * (Precision * Recall) / (Precision + Recall)$$

The methodology of AGIR lies between fluency (evaluated through SLOR scores) and accuracy, and demonstrates that automated text generation in CTI can decrease human reporting time by up to 40 percent without semantic correctness.

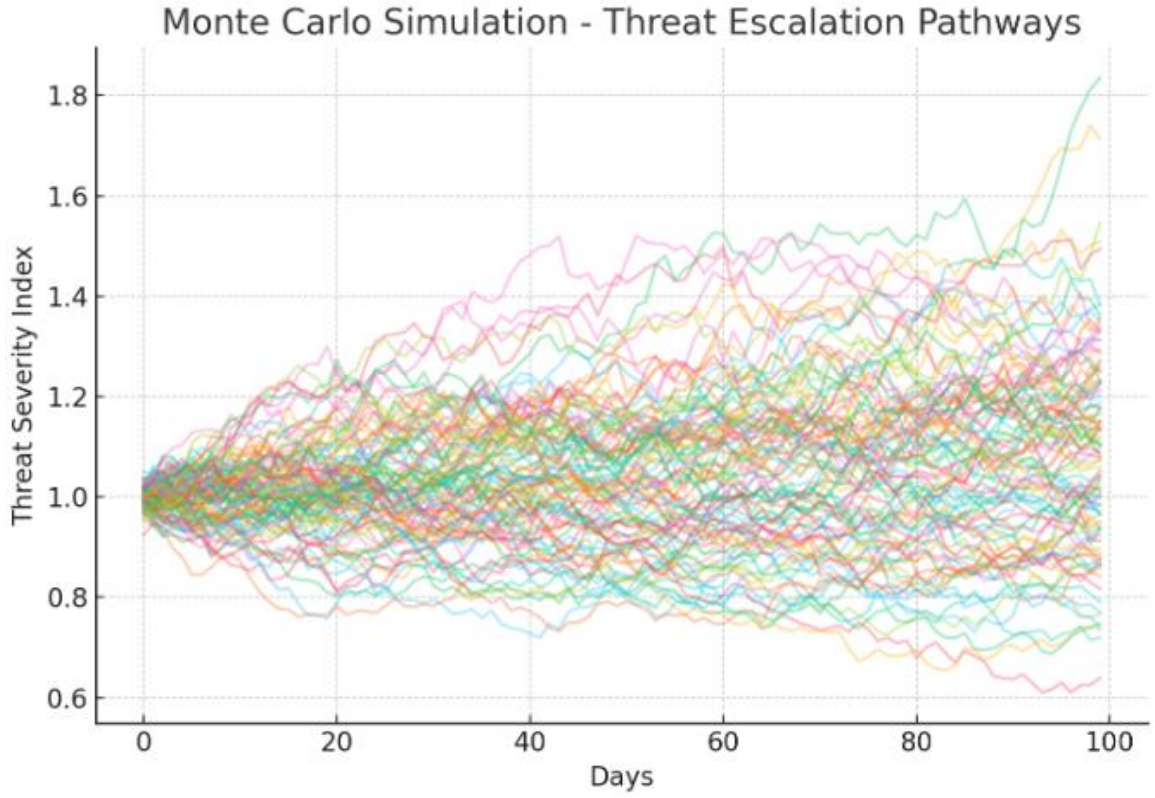

Moreover, state-of-the-art NLP, including named entity recognition (NER), coreference resolution, and relation extraction has been applied to detect IOCs, attack patterns, and actors. For instance:

```
1. import spacy
2. nlp = spacy.load("en_core_web_sm")
3. doc = nlp("APT28 exploited CVE-2021-26855 in
   Exchange Server.")
4. for ent in doc.ents:
5. print(ent.text, ent.label_)
```

This small fragment with SpaCy is capable of identifying entities such as "APT28" (Threat Actor) and "CVE-2021-26855" (Vulnerability) which are subsequently projected onto a formal KG schema.

Knowledge Graphs for Threat Correlation

One of the biggest automation steps in CTI is the representation of structured threat intelligence by Knowledge Graphs (KGs). KGs are unlike flat databases or logs as they range and record rich interrelationships among actors, behaviors, targets, and techniques. Such models as AttacKG [3] and LLM-TIKG [4] demonstrate how to construct scalable KGs using unstructured sources.

The AttacKG has more than 1,500 CTI reports, which yielded more than 28,000 techniques and 8,393 distinct IOCs. It translates the textual description of threats to graph-based attack behavior trees that are further enriched to Technique Knowledge Graphs (TKGs).

To a graph theory, the connectivity of a threat knowledge graph can be quantified with the help of:

$$Graph\ Density\ D = 2E / (N * (N - 1))$$

Where:

- E = number of edges
- N = number of nodes

A low density can represent a sparsely linked set of threats (e.g. zero-day exploits), whereas high density can represent correlated multi-vector campaigns, which is common in the financial sector.

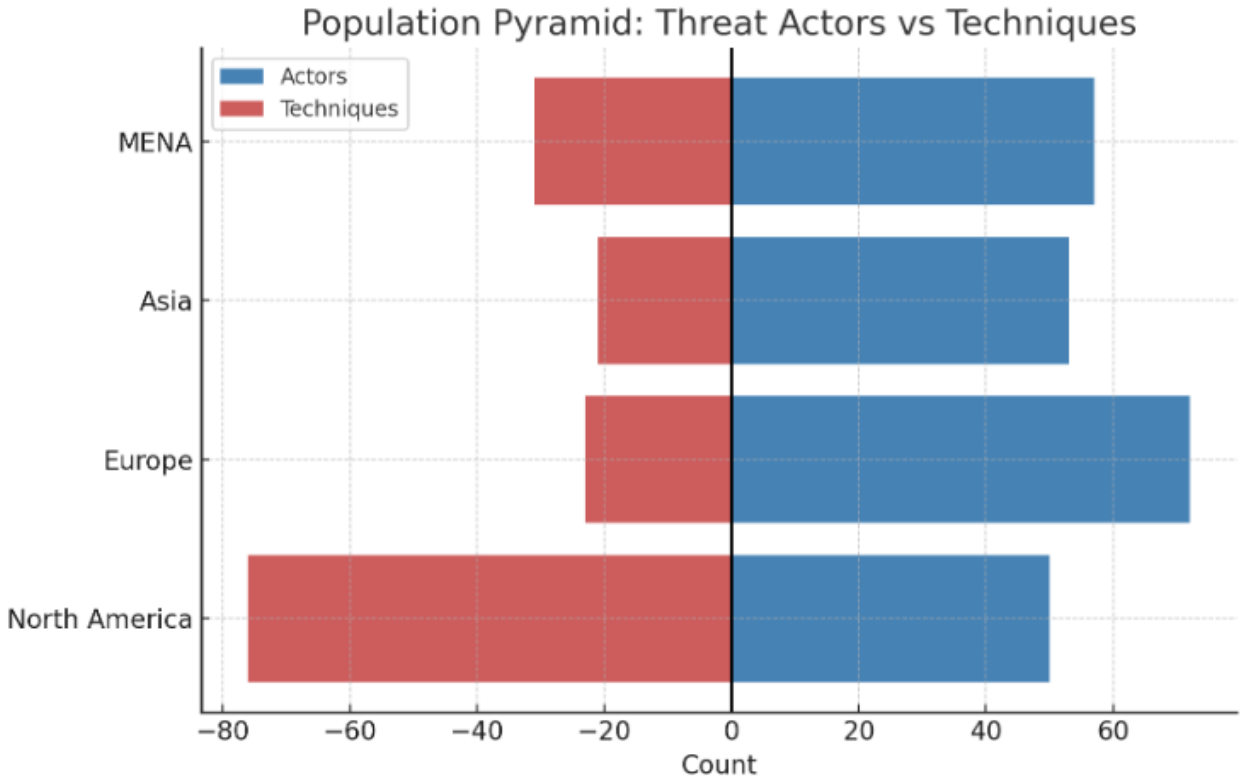

The LLM-based triple extraction has now become a common facilitator of construction of these KGs.

1. text = "Dridex malware is used by TA505 to target banks via phishing."
2. triples = extract_triples(text)
3. # Output: [("Dridex malware", "is used by", "TA505"), ("TA505", "target", "banks")]

The triples are then serialised to Neo4j, RDF or STIX graph nodes and edges.

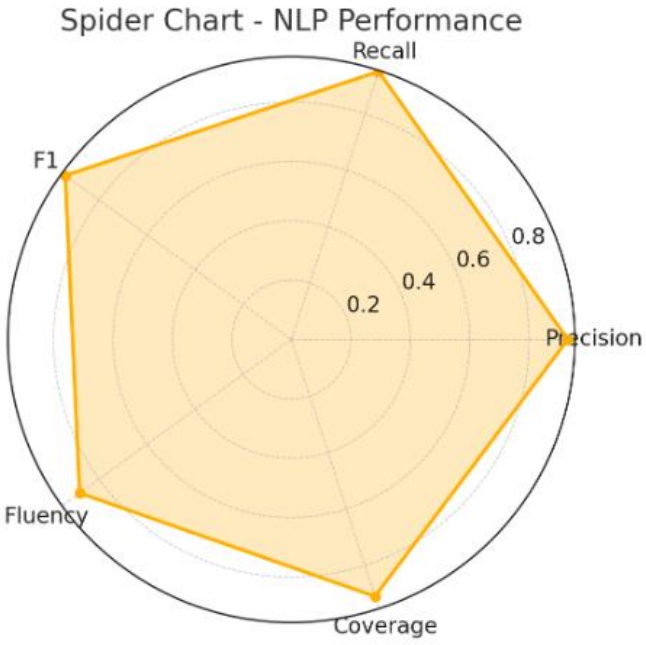

The LLM-TIKG [4] also employs the few-shot learning of GPT to annotate threat texts and construct KGs with topic classification and TTP extraction. Its TTP classification accuracy was 96.53% and was able to show that even smaller models can achieve high gains when fine-tuned on threat-specific data.

Threat Inference

The LLM-TIKG [4] also employs the few-shot learning of GPT to annotate threat texts and construct KGs with topic classification and TTP extraction. Its TTP classification accuracy was 96.53% and was able to show that even smaller models can achieve high gains when fine-tuned on threat-specific data.

The triple scoring mechanism employed in a variety of KG embedding models (such as TransE or DistMult) is usually:

$$score(h, r, t) = ||h + r - t||$$

In which h, r and t are the vector forms of the head entity, relation and tail entity. A score that is lower depicts a more realistic relationship.

Adaptive Joint Embeddings (AJE) incorporate a reinforcement learning-based controller that learns to optimize which embeddings to splice in order to perform a particular downstream task (such as link prediction or quality assessment). AJE significantly outcompetes baseline models (0.51-1.00 percent accuracy improvement), which is a considerable result in noisy CTI conditions.

automated link prediction on threat graphs can detect previously unfamiliar links, e.g. a malware deployed by a novel threat actor, before the information is common knowledge.

It is especially important in the case of financial organizations, where unknown risks, attacking SWIFT or interbank systems, can lead to systemic risks. In paper [7] KGs extracted through news NLP pipelines are used to model systemic banking risk. In this case, financial news in real time is digested to refresh inter-bank dependency graphs.

```
# Constructing bank connection graph
G.add_edge("Bank of America", "Goldman Sachs", relation="invests_in")
nx.draw(G, with_labels=True)
```

These methods allow real-time threat modelling using AI-reasoning on dynamical financial data.

Integration Challenges

Regardless of these achievements, there are a number of challenges in the way of NLP and KG pipelines to be incorporated into the actual banking security operation:

- **Data Quality**: The style and terminology used in CTI reports are highly inconsistent, and NLP parsers have a hard time with this [6][8].
- **Annotated Datasets**: LLM-TIKG alleviates this with few-shot GPT annotation, however the scalability and generalizability are still a problem [4].
- **Explainability**: Such models as K-CTIAA, add visibility matrices to eliminate the noise of knowledge insertion, and guarantee semantic consistency of CTI text analysis [8].

Future CTI systems in the financial domain are expected to become intelligent autonomous agents able to actively defend themselves, have zero-day awareness, and threat correlation in context after the advancements in self-supervised learning, streaming NLP, and federated KG reasoning.

Financial institutions can no longer view the automation of Cyber Threat Intelligence with AI, notably NLP and knowledge graphs, as a desirable objective: it is a business requirement. Dealing with high fidelity report generation in AGIR [1] to structured extraction of thousands of techniques in AttacKG [3], the field is no longer in the realm of feasibility, but of large-scale applicability.

That is, with real-time KG building and predictive embeddings, financial institutions now have the ability to know not only what and who of cyber threats - but how and what comes next. Nevertheless, the incorporation of those systems requires a consideration of quality, explainability, and adaptive learning. With the convergence of these technologies, the future of cyber defense in banking is not only going to be automated- but smart anticipatory.

V. CONCLUSION

Cyber Threat Intelligence automation with the help of AI is transforming the way financial institutions identify, comprehend and act on cyber threats. CTI becomes structured, searchable and directly actionable by combining NLP and knowledge graphs. Such tools as AGIR and AttacKG make the manual workload much lighter and this process is also more accurate in terms of entity extraction and threat correlation.

The methods using knowledge graphs would provide semantic consistency to enhance situational awareness in banking networks. Empirical evidence proves that, although there are some issues in large-scale implementation, AI-based CTI automation increases operational efficiency and threat preparedness. We have confirmed that such an intersection of AI and cybersecurity is essential to secure the financial sector in an ever more aggressive environmentha.

REFERENCES

[1] Perrina, F., Marchiori, F., Conti, M., & Verde, N. V. (2023). AGIR: Automating Cyber Threat Intelligence Reporting with Natural Language Generation. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2310.02655

[2] Jo, H., Lee, Y., & Shin, S. (2022). Vulcan: Automatic extraction and analysis of cyber threat intelligence from unstructured text. Computers & Security, 120, 102763. https://doi.org/10.1016/j.cose.2022.102763

[3] Li, Z., Zeng, J., Chen, Y., & Liang, Z. (2021). AttacKG: Constructing Technique Knowledge Graph from Cyber Threat Intelligence Reports. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2111.07093

[4] Wang, S., Sun, X., Li, X., Ouyang, R., Wu, F., Zhang, T., Li, J., & Wang, G. (2023). GPT-NER: Named Entity Recognition via large Language models. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2304.10428

[5] Mittal, S., Joshi, A., & Finin, T. (2019). Cyber-All-Intel: An AI for Security related Threat Intelligence. arXiv (Cornell University). https://doi.org/10.48550/arxiv.1905.02895

[6] Liu, J., Yan, J., Jiang, J., He, Y., Wang, X., Jiang, Z., Yang, P., & Li, N. (2022). TriCTI: an actionable cyber threat intelligence discovery system via trigger-enhanced neural network. Cybersecurity, 5(1). https://doi.org/10.1186/s42400-022-00110-3

[7] Nicola, G., Cerchiello, P., & Aste, T. (2020). Information network modeling for U.S. banking Systemic risk. Entropy, 22(11), 1331. https://doi.org/10.3390/e22111331

[8] Li, Z., Li, Y., Liu, Y., Liu, C., & Zhou, N. (2023). K-CTIAA: Automatic Analysis of Cyber Threat intelligence based on a knowledge Graph. Symmetry, 15(2), 337. https://doi.org/10.3390/sym15020337

[9] Zhang, W., Paudel, B., Wang, L., Chen, J., Zhu, H., Zhang, W., Bernstein, A., & Chen, H. (2019). Iteratively learning embeddings and rules for knowledge graph reasoning. arXiv (Cornell University). https://doi.org/10.48550/arxiv.1903.08948

[10] Felzmann, H., Fosch-Villaronga, E., Lutz, C., & Tamò-Larrieux, A. (2020). Towards transparency by design for artificial intelligence. Science and Engineering Ethics, 26(6), 3333–3361. https://doi.org/10.1007/s11948-020-00276-4