LLM-Enhanced Cyber Threat Intelligence Analysis

Assistant Prof. MA Yunshan
SCIS, Singapore Management University
06 Sep, 2025

Outline

- Cyber Threat Intelligence
- MITRE ATT&CK Knowledge Base
- (M-)LLM for Attack Graph Construction
- LLM for Attack Sequence Prediction
- Conclusion and Future Works

SMU SINGAPORE MANAGEMENT UNIVERSITY

What is Cyber Threat Intelligence (CTI)?

- **Cyber threat intelligence (CTI)[1]:** CTI is the process of collecting, analyzing, and applying data on cyber threats, adversaries, and attack methodologies to enhance an organization's security posture.

- **CTI Report:** a document that provides actionable information about potential or existing cyber threats, enabling organizations to proactively defend against attacks and minimize their impact.

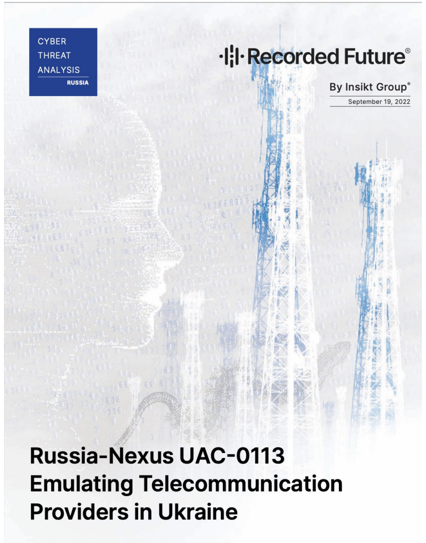

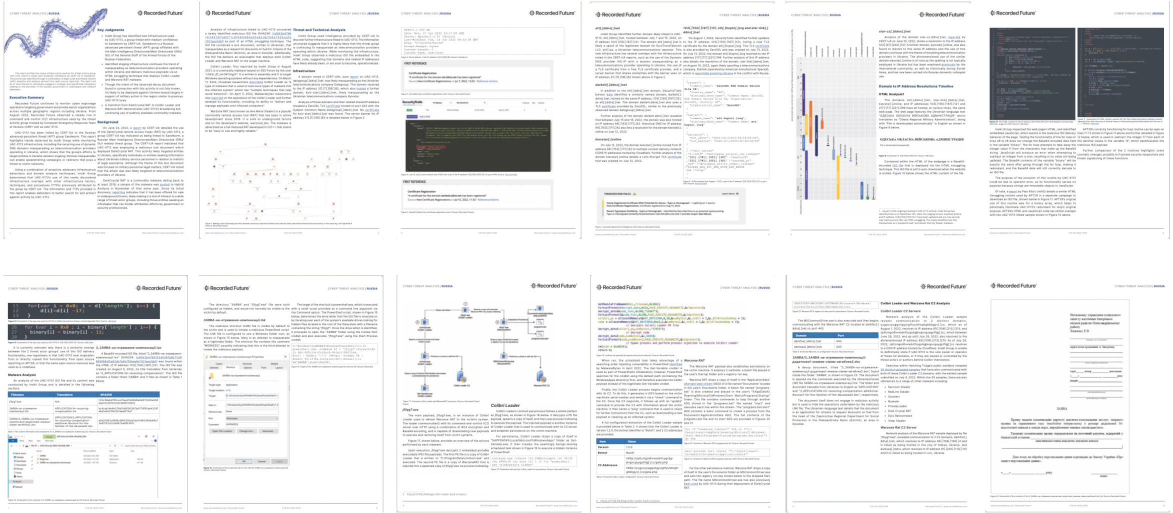

[1] https://www.paloaltonetworks.com/cyberpedia/what-is-cyberthreat-intelligence-cti
[2] https://assets.recordedfuture.com/insikt-report-pdfs/2022/cta-2022-0919.pdf

What is Cyber Threat Intelligence (CTI)?

SMU SINGAPORE MANAGEMENT UNIVERSITY

CTI lifecycle[3]:

- ➢ CTI requirements
- ➢ CTI collection
- ➢ CTI processing
- ➢ CTI analysis
- ➢ CTI dissemination
- ➢ CTI feedback

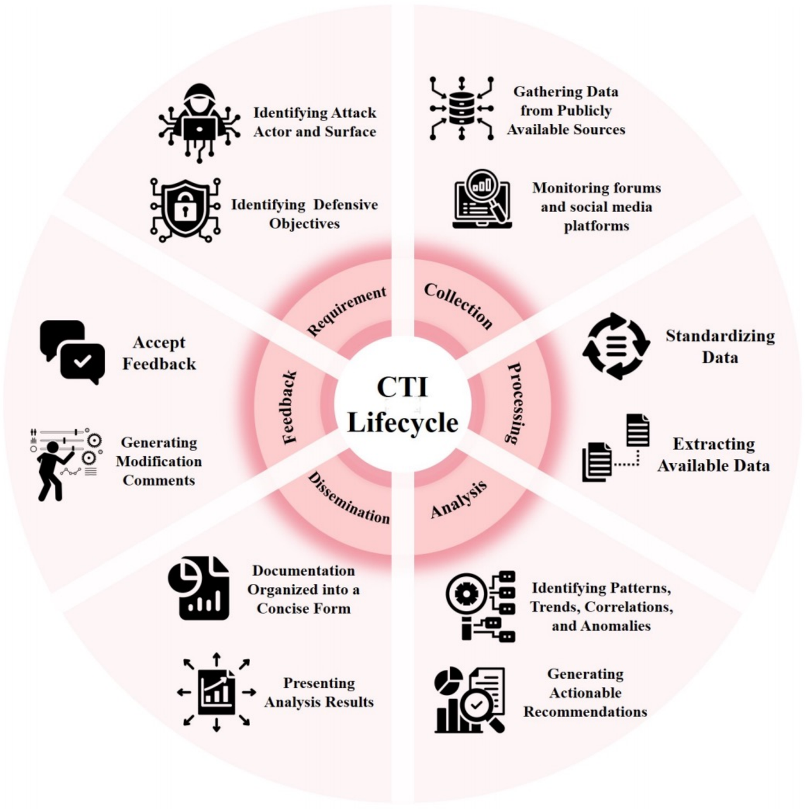

[3] Security and resilience in sustainable smart cities through cyber threat intelligence. K. Nova et al. International Journal of Information and Cybersecurity.

Why is CTI important?

Significance:

- Cyber threat intelligence is an essential component of an organization's cyber resiliency, which includes “the ability to anticipate, withstand, recover from, and adapt” to threats, attacks, or compromises on systems[4].

Benefits:

- **Establishing proactive defense**
  - ➢ Anticipates potential attackers and attacks rather than reacting to known threats.
- **Improving risk management**
  - ➢ Provides insights into adversaries’ motives, methods, and means for better resource allocation.
- **Enhancing incident response**
  - ➢ Equips organizations to respond faster and recover more effectively from breaches.
- **Increasing employee awareness**
  - ➢ Educates staff on threats and reinforces security-focused practices.

[4] https://csrc.nist.gov/glossary/term/cyber_resiliency

MITRE ATT&CK Knowledge Base

- MITRE ATT&CK® is a globally-accessible knowledge base of adversary tactics and techniques based on real-world observations. The ATT&CK knowledge base is used as a foundation for the development of specific threat models and methodologies in the private sector, in government, and in the cybersecurity product and service community.

- With the creation of ATT&CK, MITRE is fulfilling its mission to solve problems for a safer world — by bringing communities together to develop more effective cybersecurity. ATT&CK is open and available to any person or organization for use at no charge.

- The MITRE Corporation is an American not-for-profit organization, which supports various U.S. government agencies in the aviation, defense, healthcare, homeland security, and **cybersecurity** fields, among others.

[5] https://attack.mitre.org/

MITRE ATT&CK Knowledge Base

SMU SINGAPORE MANAGEMENT UNIVERSITY

TTP: Tactics, Techniques, and Procedures:

- Tactics: The high-level goals of an attacker, such as gaining initial access.
- Techniques: The specific methods used to achieve those tactical goals, like phishing or exploiting vulnerabilities.
- Procedures: The detailed steps and actions taken to execute the techniques.

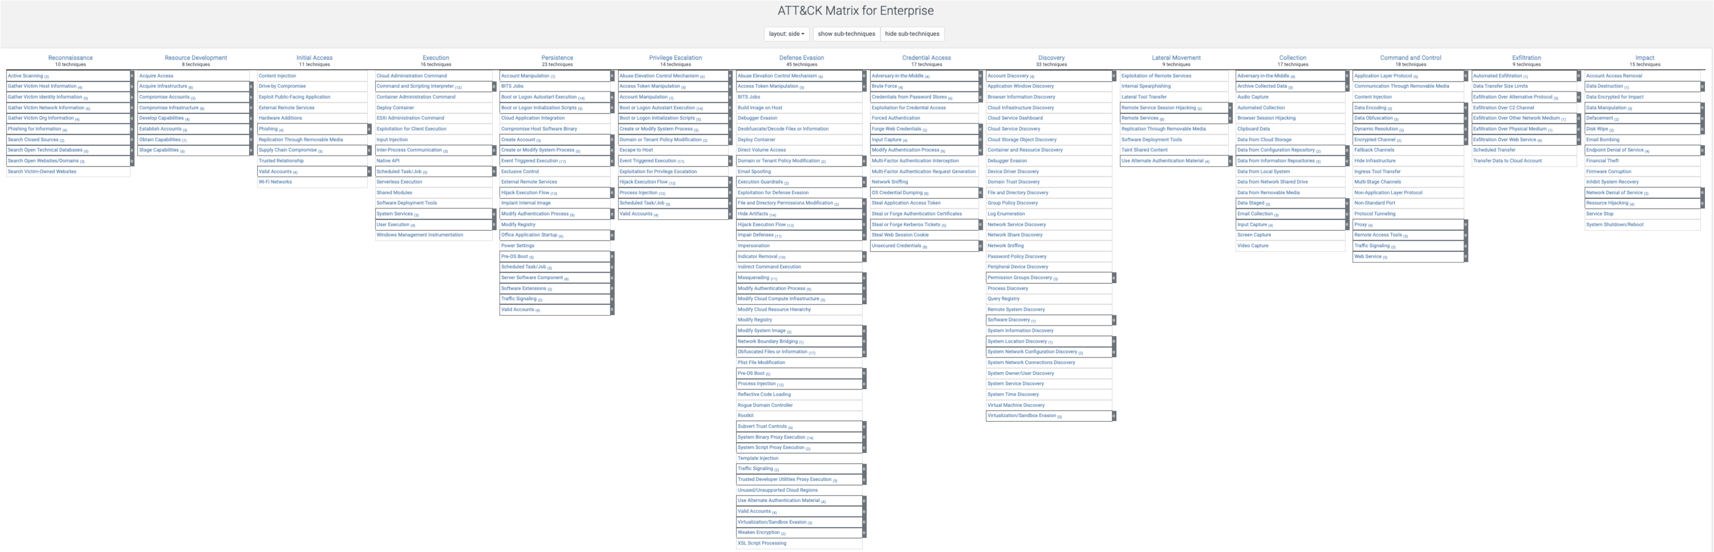

[5] https://attack.mitre.org/

MITRE ATT&CK Knowledge Base

SMU SINGAPORE MANAGEMENT UNIVERSITY

TTP: Tactics, Techniques, and Procedures:

- Tactics: The high-level goals of an attacker, such as gaining initial access.
- Techniques: The specific methods used to achieve those tactical goals, like phishing or exploiting vulnerabilities.
- Procedures: The detailed steps and actions taken to execute the techniques.

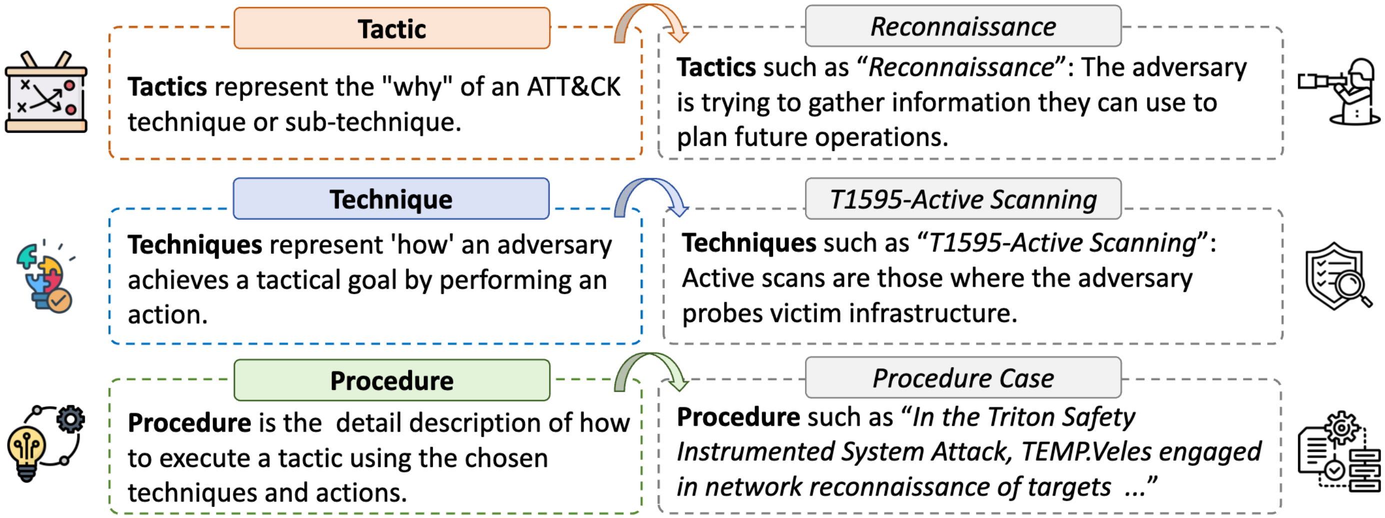

[5] https://attack.mitre.org/

MITRE ATT&CK Knowledge Base

SMU SINGAPORE MANAGEMENT UNIVERSITY

- Tactics (14):
  - Reconnaissance
  - Resource Development
  - Initial Access
  - Execution
  - Persistence
  - Privilege Escalation
  - Defense Evasion
  - Credential Access
  - Discovery
  - Lateral Movement
  - Collection
  - Command and Control
  - Exfiltration
  - Impact
- Techniques:
  - Active Scanning (3)
  - Gather Victim Host Information (4)
  - Gather Victim Network Information (6)
  - Gather Victim Org Information (4)
  - Phishing for Information (4)
  - Search Closed Sources (2)
  - Search Open Technical Databases (5)
  - Search Open Websites/Domains (3)
- Sub-techniques:
  - Scanning IP Blocks
  - Vulnerability Scanning
  - Wordlist Scanning

Tactics → Techniques → Sub-techniques

[5] https://attack.mitre.org/

Attack Graph (AttacKG) Construction

- Task formulation: CTI report pdf → an attack knowledge graph

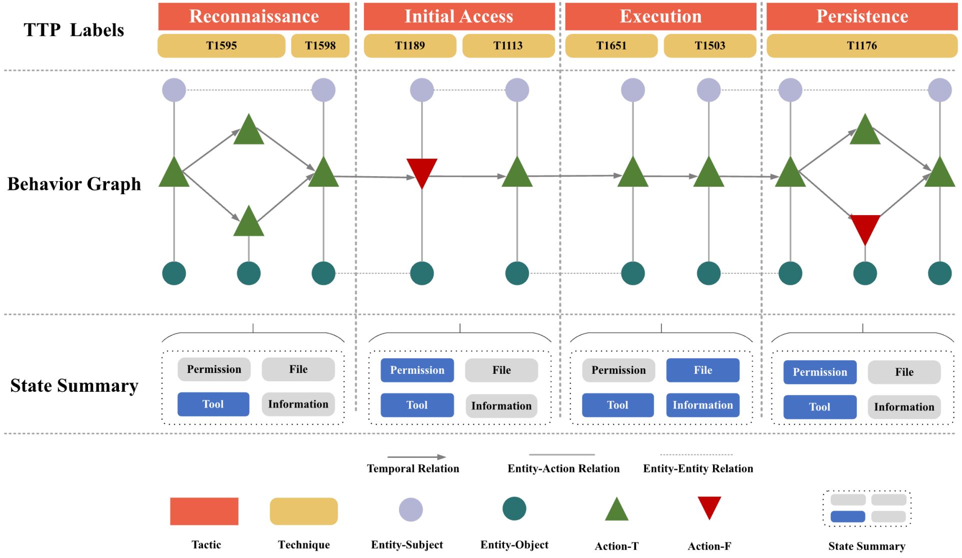

Attack Graph (AttacKG) Construction

- Conventional approaches:
  - Non-learning methods
    - ➢ Regular expression
  - Learning-based methods:
    - ➢ Information extraction

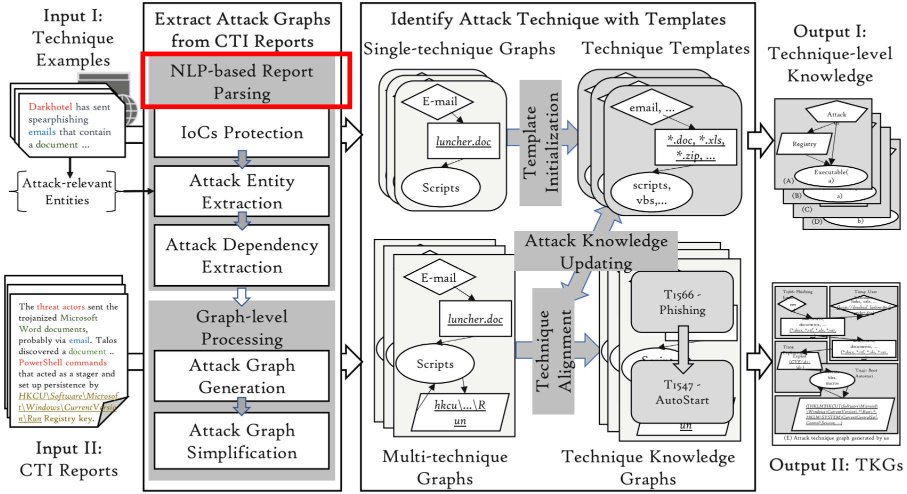

- Limitations:
  - Poor performance due to limited semantic understanding capabilities
  - Need large-scale annotated dataset, which is expensive, time-consuming, infeasible
  - Hard to generalize to new knowledge (ATT&CK regularly updates new types)

[6] Zhenyuan Li, et al., AttacKG: Constructing Technique Knowledge Graph from Cyber Threat Intelligence Reports. In Computer Security – ESORICS 2022

Attack Graph (AttacKG) Construction

Our proposal – AttacKG+

- Using Large Language Models (LLMs) for attack graph construction

Key motivations:

- Leverage LLMs' strong semantic understanding capabilities
- No need to annotate datasets, we can directly do extraction using instruction following (zero-shot) and in-context-learning (few-shot)

Key modules:

- Rewriter
- Parser
- Identifier
- Summarizer

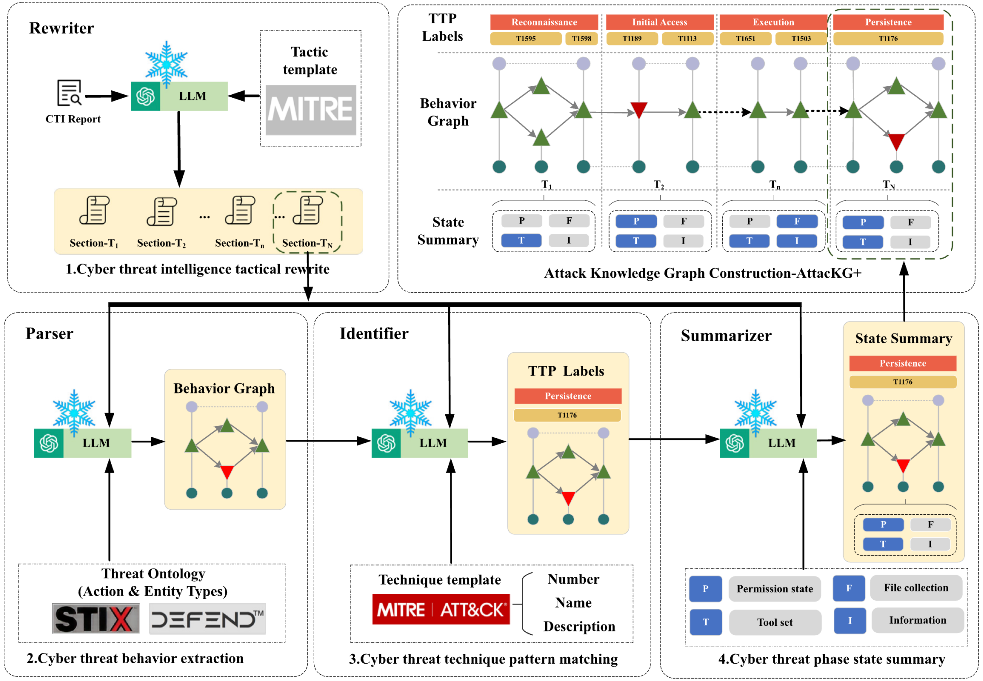

[7] Yongheng Zhang et al. AttacKG+: Boosting attack graph construction with Large Language Models. Computers & Security. 2025

Attack Graph (AttacKG) Construction

SMU
SINGAPORE MANAGEMENT UNIVERSITY

Results

Table 2
Accuracy of AttacKG+ construction and technique identification.

CTI reports	Entities			Relations			Techniques		
	Manual	Extractor	AttacKG+	Manual	Extractor	AttacKG+	Manual	AttacKG	AttacKG+
BRONZE	13	−13 (+10)	−2 (+9)	8	−5 (+18)	−2 (+9)	4	−1 (+18)	−3 (+4)
Chat_Mimi	15	−15 (+9)	−5 (+8)	10	−7 (+15)	−5 (+4)	4	−1 (+7)	−2 (+1)
North_Korea	22	−19 (+15)	−4 (+5)	9	−4 (+22)	−2 (+4)	7	−3 (+23)	−2 (+2)
Nitro_Attacks	28	−28 (+8)	−8 (+5)	19	−6 (+22)	−7 (+5)	8	−5 (+14)	−3 (+6)
Moon_Bounce	12	−12 (+5)	−1 (+10)	10	−6 (+22)	−5 (+10)	5	−2 (+12)	−3 (+4)
Stuxnet_Under	24	−22 (+21)	−8 (+3)	18	−6 (+31)	7 (+5)	11	−8 (+19)	−5 (+6)
Stellar_Particle	33	−32 (+12)	−6 (+5)	13	−5 (+18)	−5 (+7)	10	−10 (+10)	−1 (+3)
Prime_Minister	19	−19 (+10)	−5 (+9)	12	−4 (+12)	−4 (+3)	12	−8 (+11)	−1 (+1)
Mustang_Panda	37	−37 (+10)	−9 (+3)	19	−13 (+28)	−10 (+7)	12	−7 (+22)	−3 (+9)
Shuckworm_APT	17	−16 (+24)	−2 (+11)	9	−5 (+18)	−1 (+8)	7	−3 (+9)	−2 (+4)
C5_APT_SKHack	13	−11 (+4)	−4 (+4)	9	−5 (+18)	−3 (+1)	5	−3 (+17)	−3 (+4)
Cisco_Talos_Bitter	17	−17 (+10)	−9 (+3)	8	−5 (+18)	−3 (+1)	3	−2 (+21)	−1 (+1)
Log4Shell_Rootkits	38	−36 (+8)	−14 (+7)	22	−13 (+17)	−10 (+7)	16	−12 (+8)	−9 (+5)
Cisco_Talos_Iranian	14	−14 (+8)	−3 (+7)	6	−3 (+19)	−3 (+2)	4	−2 (+9)	−3 (+1)
Asylum_Ambuscade	21	−21 (+10)	−9 (+3)	11	−6 (+24)	−4 (+3)	4	−1 (+16)	−1 (+3)
Overall precision	**1.000**	**0.046**	**0.668**	**1.000**	**0.221**	**0.601**	**1.000**	**0.179**	**0.545**
Overall recall	**1.000**	**0.034**	**0.732**	**1.000**	**0.472**	**0.647**	**1.000**	**0.458**	**0.588**
Overall F-1 score	**1.000**	**0.039**	**0.698**	**1.000**	**0.301**	**0.623**	**1.000**	**0.258**	**0.566**

[1] Accuracy of threat behavior graph construction and technique identification in 15 CTI reports.
[2] Columns 2–10 present the ground-truth and **false negative/positive** in extracting entities, relations, and techniques.
[3] Rows 18–20 present the overall Precision, Recall, and F-1 Score.

Green: FN
Red: FP
Manual: GT

- LLMs (GPT-4) exhibit much better performance observed in **more recent** CTI reports (unseen data), compared with baselines (Extractor).

Attack Graph (AttacKG) Construction

SMU SINGAPORE MANAGEMENT UNIVERSITY

Case study

- Taking the attack on SK Communications (C5 APT SHack) as an example, AttacKG+ extracts structured knowledge of threat event scenarios from this event.
- The multi-level attack graph representation shows the development process of threat events more clearly and intuitively.

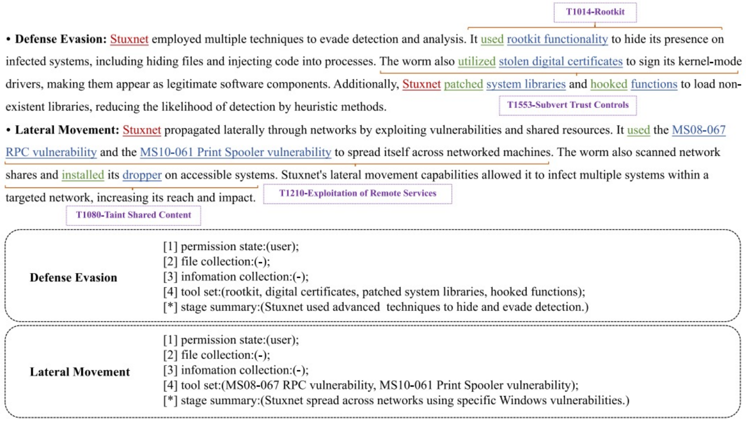

Example of AttacKG+ extraction (Stuxnet)

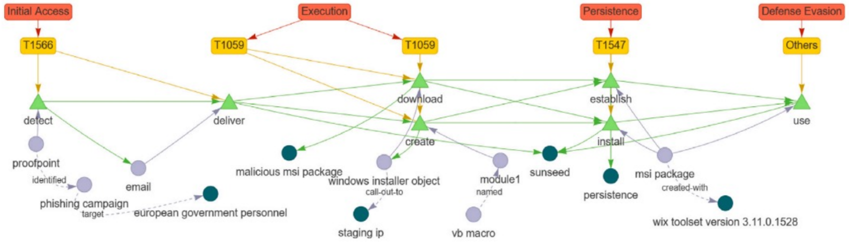

An attack case against SK Communications

Multimodal Attack Graph Construction

MM-AttacKG: A Multimodal Approach to Attack Graph Construction with Large Language Models

- Motivation: leverage the images (visual information) in CTI report.

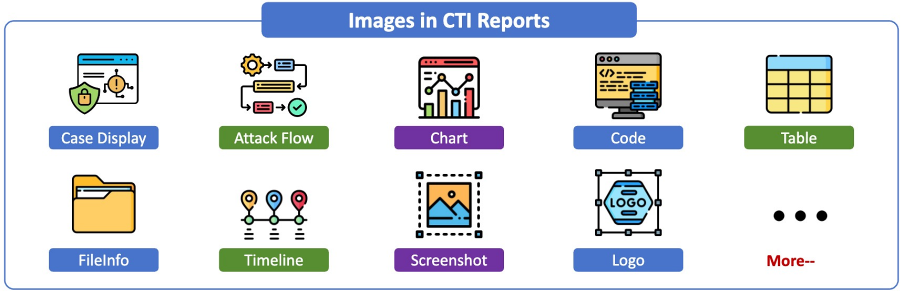


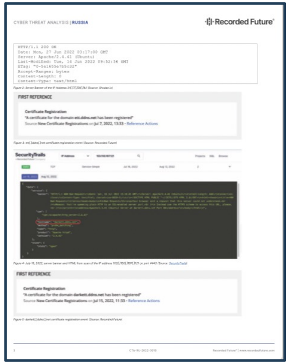
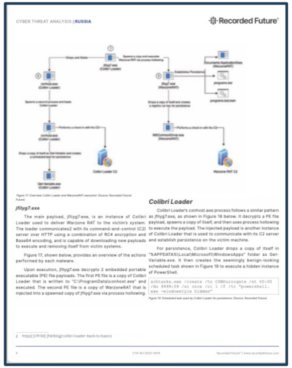
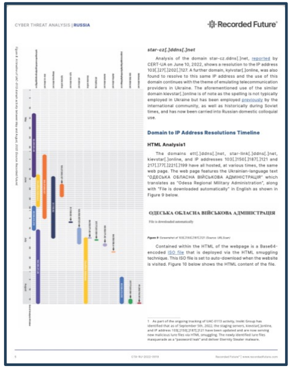
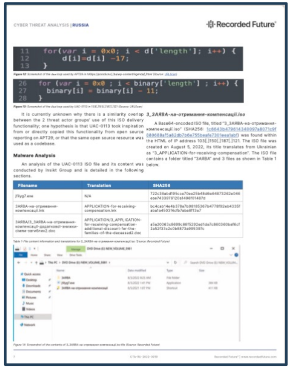

[8] MM-AttacKG: A Multimodal Approach to Attack Graph Construction with Large Language Models. Yongheng Zhang et al. arXiv. 2025

Multimodal Attack Graph Construction

SMU
SINGAPORE MANAGEMENT UNIVERSITY

MM-AttacKG: A Multimodal Approach to Attack Graph Construction with Large Language Models

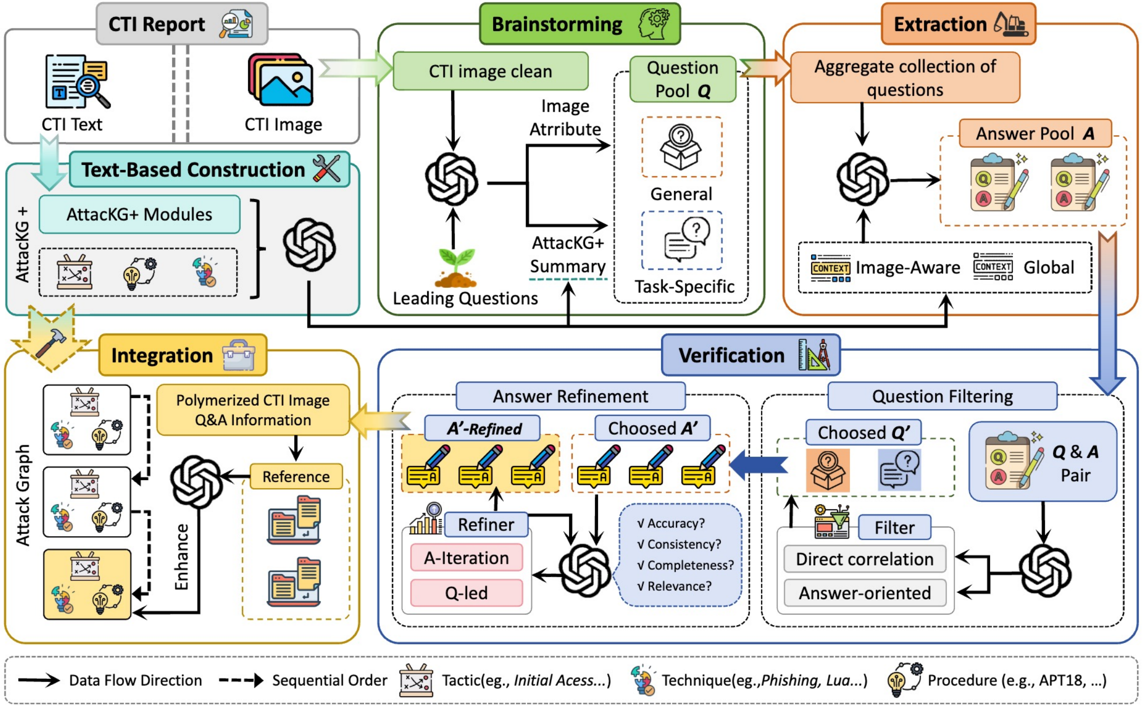

Key modules:
- Brainstorming
- Extraction
- Verification
- Integration

[8] MM-AttacKG: A Multimodal Approach to Attack Graph Construction with Large Language Models. Yongheng Zhang et al. arXiv. 2025

Multimodal Attack Graph Construction

Overall performance

Method	Entity			Relation			Technique		
	Precision	Recall	F-1	Precision	Recall	F-1	Precision	Recall	F-1
Text-based Method									
Extractor	0.6568	0.5387	0.5919	0.2158	0.1026	0.1391	-	-	-
AttacKG	0.5580	0.2612	0.3559	-	-	-	0.2060	0.3399	0.2565
AttacKG+	**0.7701**	0.5294	0.6274	**0.7693**	0.6806	0.7222	0.4502	0.4481	0.4491
Human Anotation-Text	1.0000	0.4559	0.6263	1.0000	0.6820	0.8109	1.0000	0.6547	0.7913
Image-enhanced Method									
ICL	0.6901	0.7326	0.7107	0.7106	0.8261	0.7640	0.4948	0.5383	0.5156
CoT	0.6805	0.7432	0.7105	0.6949	0.8383	0.7599	0.5063	0.5508	0.5277
MM-AttacKG	0.7224	**0.8280**	**0.7716**	0.7460	**0.8973**	**0.8147**	**0.5256**	**0.6232**	**0.5703**

- Image-enhanced methods achieve much higher recall and F-1, showing that leveraging images within CTI reports provides significant more information to enrich the attack graph.
- MM-AttacKG outperforms both ICL and CoT, showing our framework well caters to the CTI report characteristics, thereby extracting more valuable attack information.

Multimodal Attack Graph Construction

SMU SINGAPORE MANAGEMENT UNIVERSITY

Different LLM backbones and prompting strategies

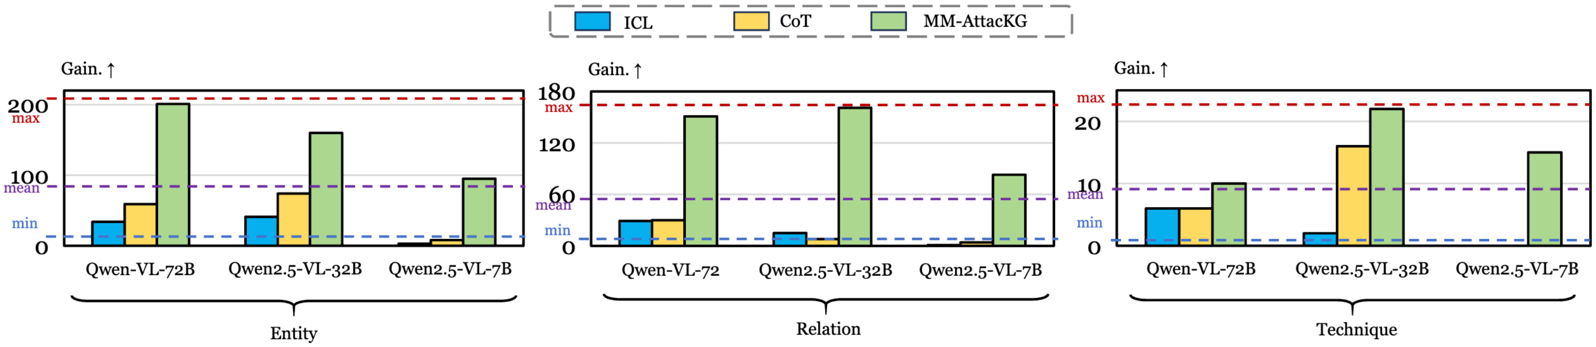

Iterative QA improves the quality with increasing iterations

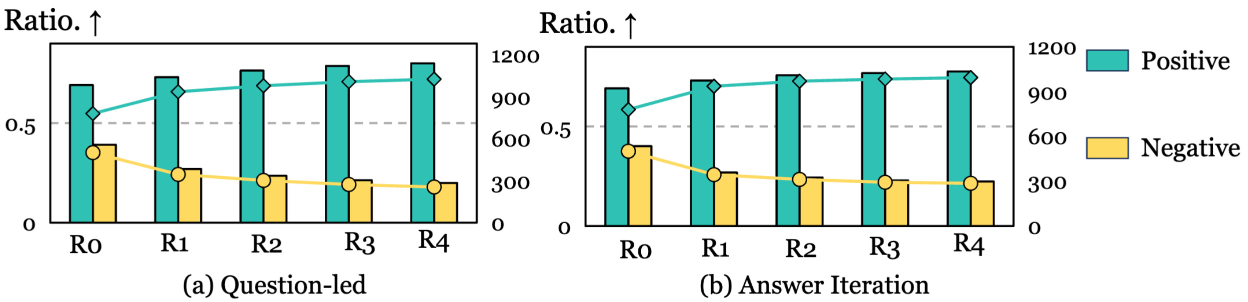

Multimodal Attack Graph Construction

Case study: pure-text based attack graph

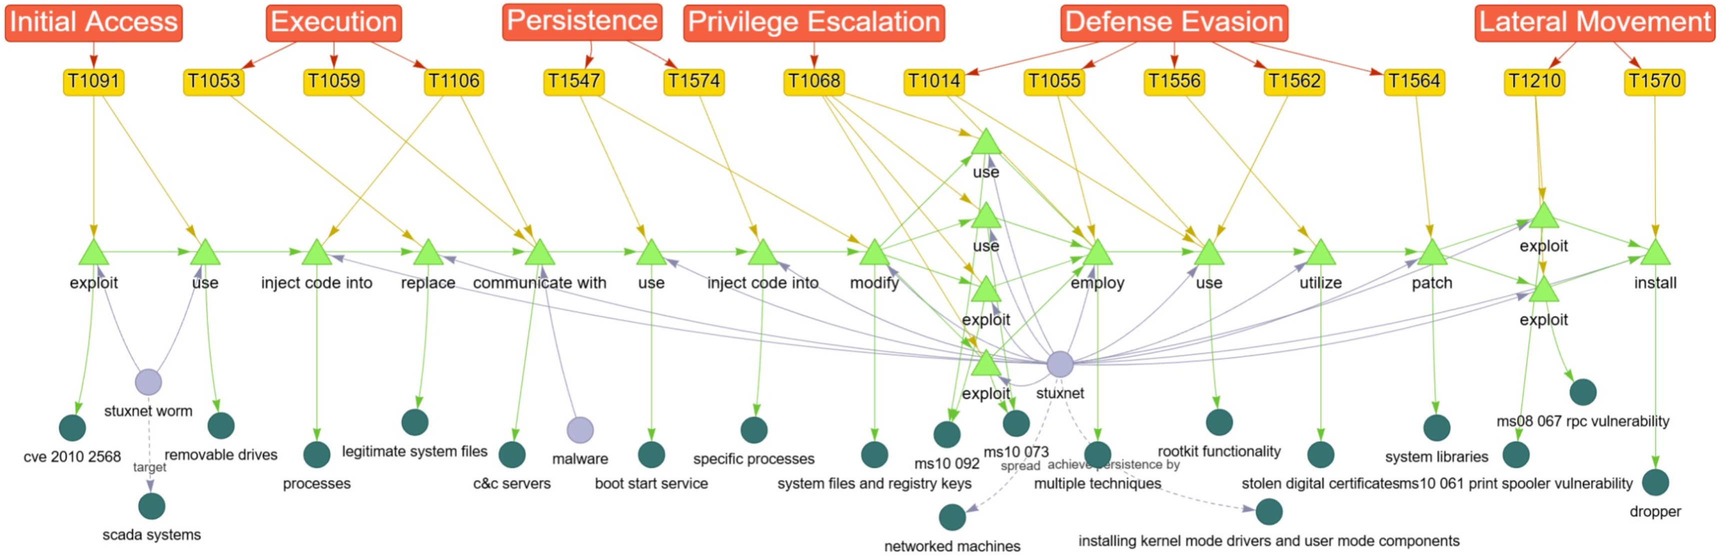

Multimodal Attack Graph Construction

SMU SINGAPORE MANAGEMENT UNIVERSITY

Case study: incorporating images in CTI

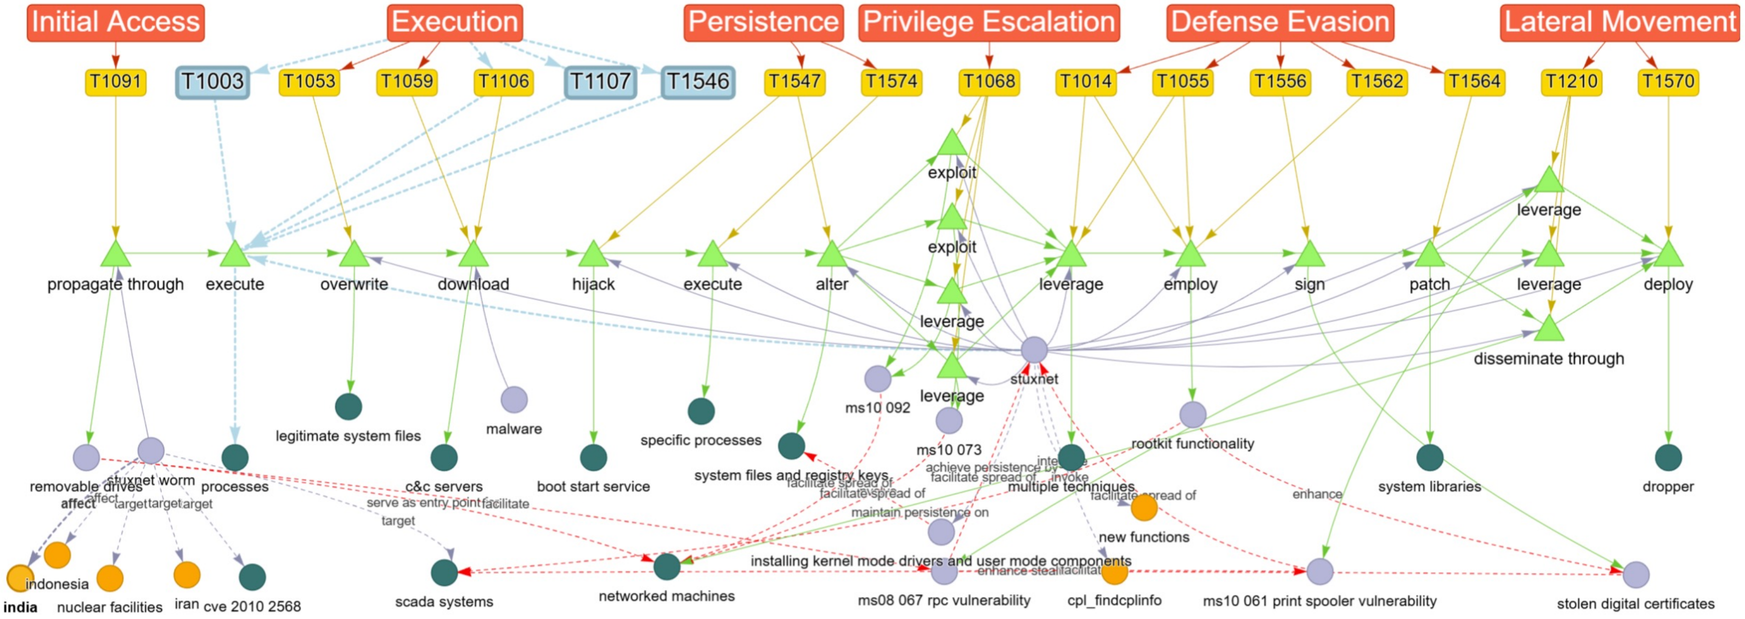

Note: red dotted lines indicate the newly extracted knowledge from images.

Attack Sequence Prediction

Motivation:

- Cyber attacks often involves multiple consecutive steps, forming an attack sequence (attack flow).
- Understanding the sequential patterns and making accurate prediction are essential for cyber attack analysis.
- We aim to extend pure textual or multimodal from **understanding** to **prediction.**

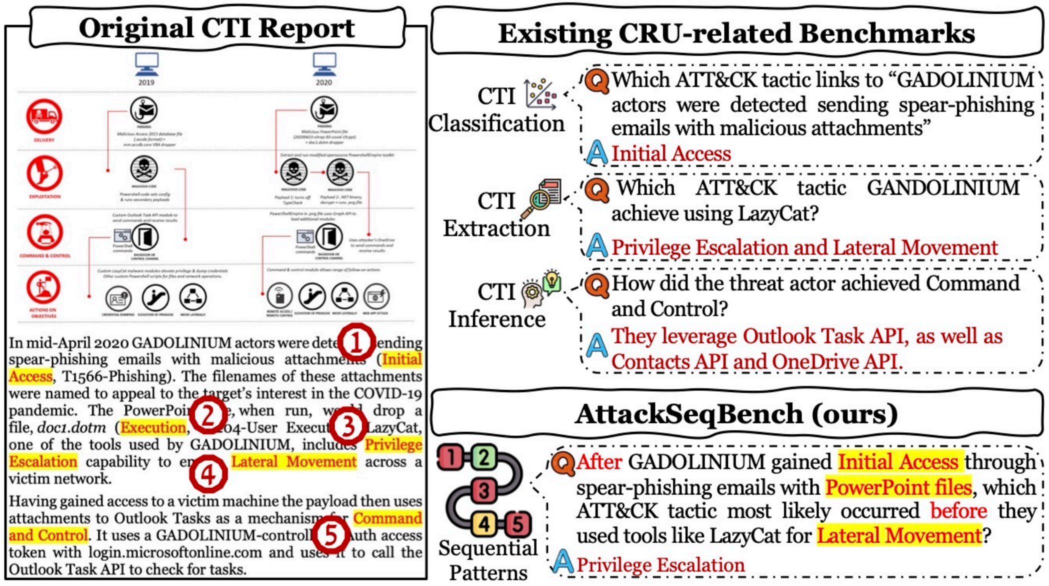

Attack Sequence Prediction

AttackSeqBench:

1. Model attack sequences within CTI reports.
2. Design an automated QA dataset construction pipeline based on 3 tasks (*i.e.,* TTP).
3. Perform benchmark on a diverse set of LLMs.

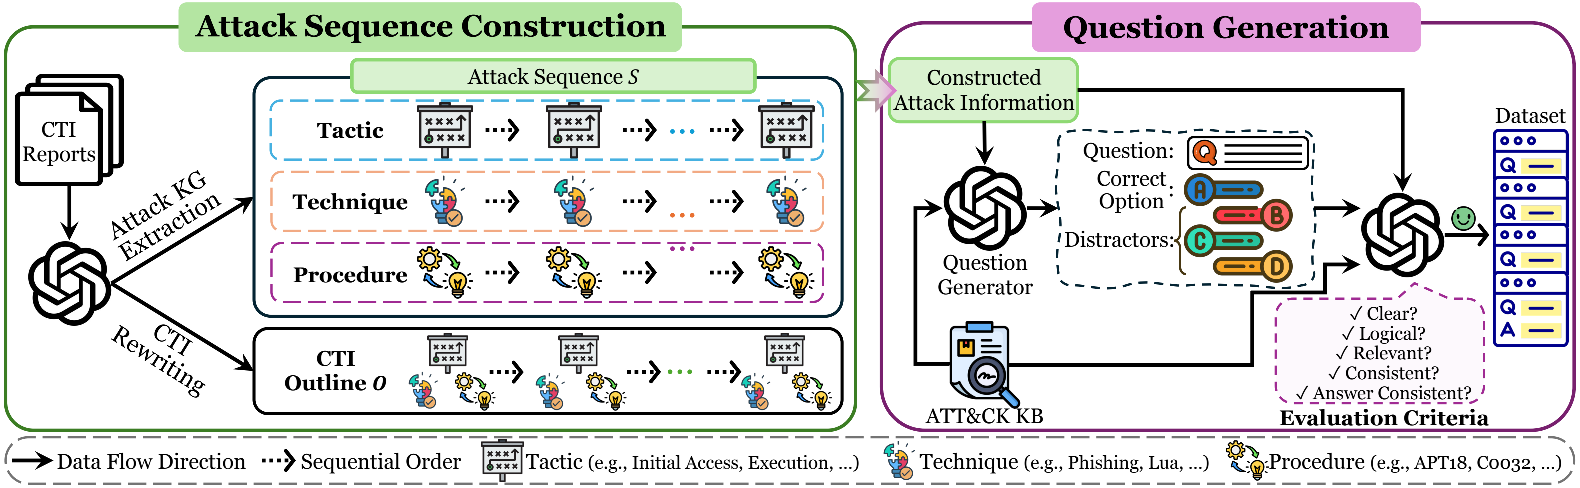

Q&A dataset construction pipeline

[9] AttackSeqBench: Benchmarking Large Language Models' Understanding of Sequential Patterns in Cyber Attacks. Javier et al. arXiv 2025.

Attack Sequence Prediction

SMU SINGAPORE MANAGEMENT UNIVERSITY

Attack Sequence Formulation:

- S is a 4-tuple (T, E, P, O) where:
  - Tactic Sequence $T = (t_1, \dots, t_n)$.
  - Technique Mapping $\forall t_k \in T, E(t_k) = \{e_{1,k}, \dots, e_{i_k,k}\}$.
  - Procedure Mapping $\forall e_{j,k} \in E, P(e_{j,k}) = \{p_{1,j,k}, \dots, p_{m_{j,k},j,k}\}$.
  - CTI Outline $O = \{o_1, \dots, o_n\}$

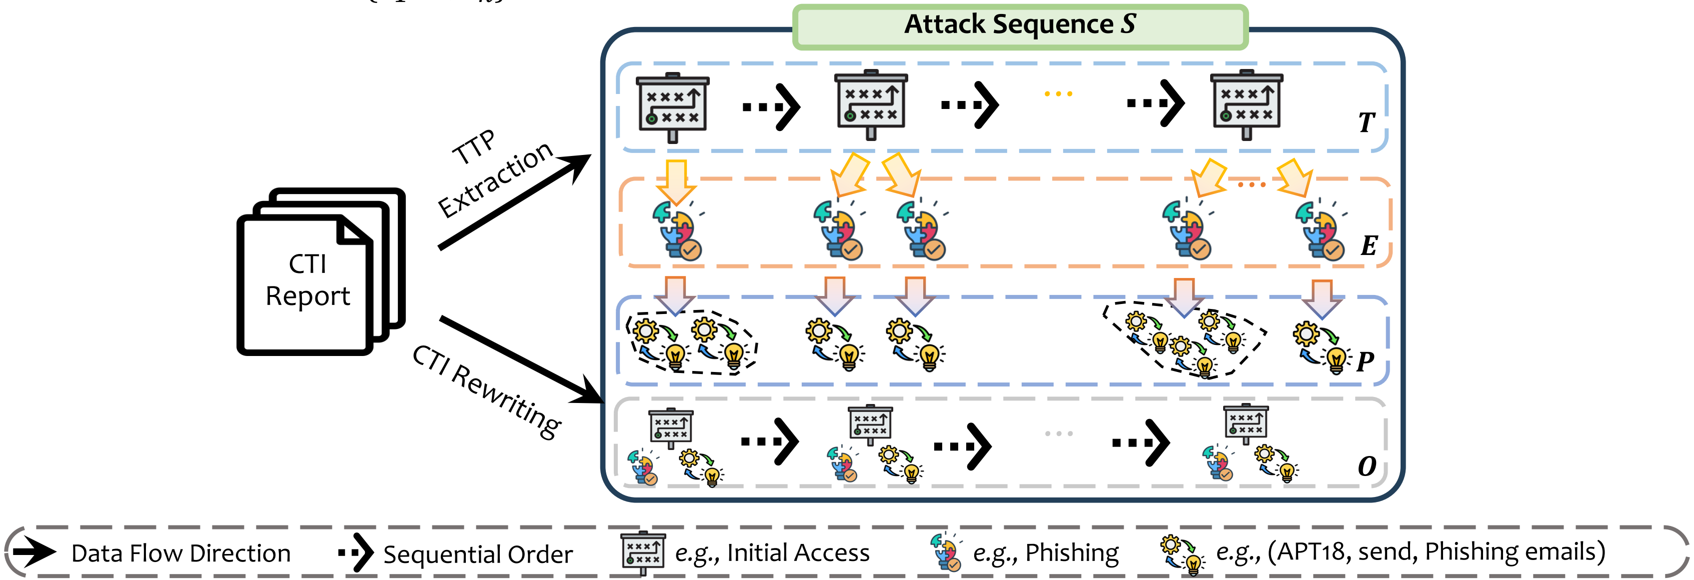

Attack Sequence Prediction

SMU
SINGAPORE MANAGEMENT UNIVERSITY

Attack Sequence Prediction task:

- Three Question Answering (QA) tasks based on TTP.
- Evaluate abductive reasoning abilities in attack sequences.
    - i.e., Infer most plausible TTP in the sequence given remaining TTPs.

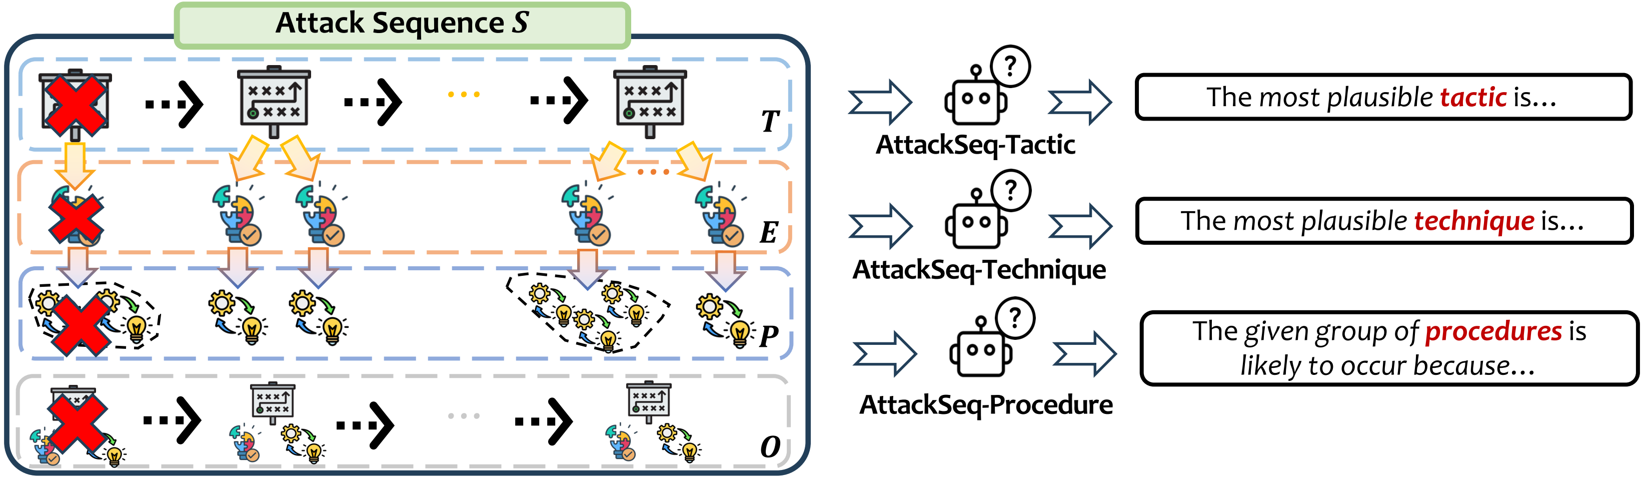

Attack Sequence Prediction

Question Generation:

- Construct attack sequences from 500 real-world CTI reports.
- Answer-aware QG approach using LLM.
  - A given tactic, technique, or group of procedures in S.
  - Distractors are randomly selected from ATT&CK KB.

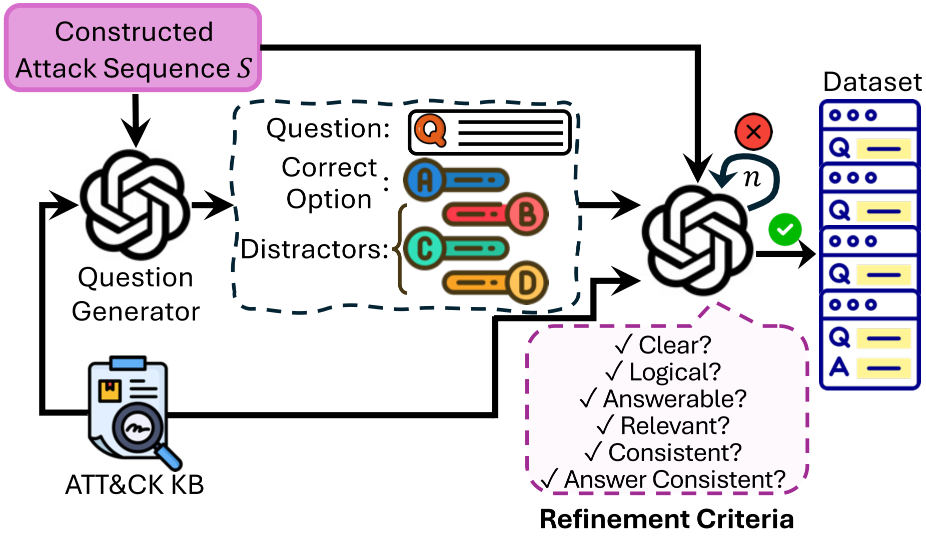

[10] https://github.com/blackorbird/APT_REPORT

Attack Sequence Prediction

SMU SINGAPORE MANAGEMENT UNIVERSITY

Benchmark Methods

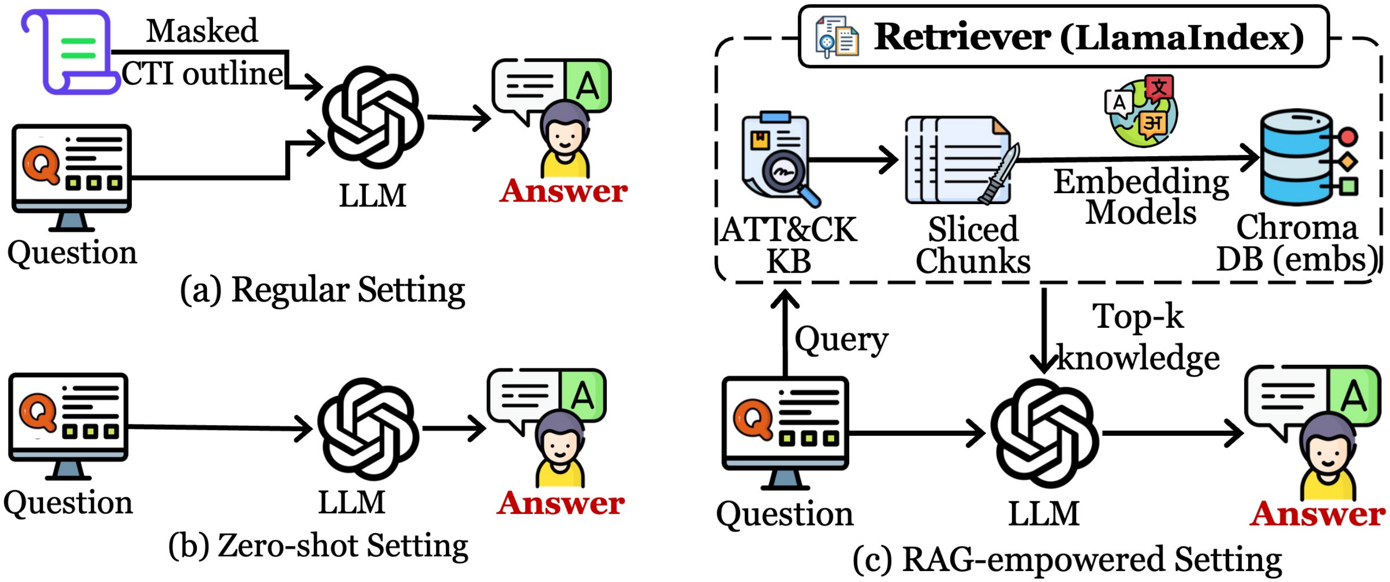

Attack Sequence Prediction

Dataset Evaluation:

- Utilize 5-point Likert scales using the same criteria.

Human Evaluation: 3 domain experts on a random sample of questions.

Task	Num.	Hum. Perf.	Answerability	Clarity	Logical	Relevance	Consistency	Answer Consistency
			Scores (out of 5)					
AttackSeq-Tactic	35	0.5143	4.2952	4.3619	4.4476	4.5619	4.4571	4.4381
AttackSeq-Technique	35	0.7143	4.0857	4.2095	4.4000	4.4476	4.4381	4.4095
AttackSeq-Procedure-Yes	35	0.7429	4.8762	4.6952	4.8762	5.0000	4.8095	4.9429
AttackSeq-Procedure-No	35	0.5619	4.5524	4.838	-	-	4.8190	4.6571
Total	140	0.6333	4.4524	4.5262	4.5746	4.6698	4.6310	4.6119

Automatic Evaluation: G-eval (LLM-based) evaluation on entire dataset.

Task	Answerability	Clarity	Logical	Relevance	Consistency	Answer Consistency
	Scores (out of 5)					
AttackSeq-Tactic	4.5200	4.6510	4.7901	4.8360	4.6530	4.7590
AttackSeq-Technique	4.1040	4.3960	4.6200	4.6300	4.3870	4.5910
AttackSeq-Procedure-Yes	4.0170	4.0640	4.6110	4.4650	3.7760	3.8940
AttackSeq-Procedure-No	3.2930	3.6600	-	-	2.7650	3.2490
Average	3.9835	4.1928	4.6737	4.6437	3.8953	4.1233

Attack Sequence Prediction

SMU SINGAPORE MANAGEMENT UNIVERSITY

Findings:

- No LLM dominates in all benchmark tasks.
- LLMs performed worst in Tactic-level task.
- Contextual information is critical in Procedure-level task (i.e., Regular vs. Zero-Shot).

LLMs	AttackSeq-Tactic			AttackSeq-Technique			AttackSeq-Procedure		
	Regular	Zero-Shot	RAG	Regular	Zero-Shot	RAG	Regular	Zero-Shot	RAG
Fast-thinking LLMs									
Mistral-7B-Instruct-v0.3	0.2823	0.3371	0.2752	0.3432	0.3975	0.2999	0.5359	0.5795	0.5484
Qwen-2.5-7B-Instruct	0.5121	0.4903	0.4761	0.6693	0.6568	0.6067	0.6584	0.5184	0.4941
Llama-3.1-8B-Instruct	0.4744	0.5085	0.4926	0.6260	0.6333	0.5827	0.6577	0.5328	0.5230
ChatGLM-4-9B-Chat	0.4885	0.4979	0.5009	0.6275	0.6109	0.6030	0.641	0.5408	0.5131
Llama-3.3-70B-Instruct	**0.6588**	0.5551	0.5681	<u>0.7058</u>	0.6797	<u>0.7037</u>	0.6903	0.5469	0.5279
Qwen-2.5-72B-Instruct	0.5793	<u>0.5863</u>	0.5657	0.5430	0.7162	0.6959	<u>0.7188</u>	<u>0.6285</u>	0.6030
GPT-4o-mini	<u>0.6517</u>	**0.6005**	<u>0.5692</u>	**0.7387**	<u>0.7058</u>	0.7021	0.6968	0.5491	0.5340
GPT-4o	0.6093	0.5740	**0.5787**	0.6755	0.6995	**0.7188**	**0.7359**	0.6755	<u>0.6353</u>
Slow-thinking Reasoning LLMs									
DeepSeek-R1-Distill-Llama-8B	0.4178	0.4467	0.4532	0.5389	0.5519	0.5138	0.6194	0.5044	0.4968
DeepSeek-R1-Distill-Qwen-32B	0.5421	0.5698	0.5504	0.5879	0.5816	0.5816	0.6945	0.6258	0.5852
QWQ-32B-Preview	0.5345	0.3377	0.4638	0.5112	0.3918	0.5342	0.7036	0.5696	0.5457
GPT-o3-mini	0.4643	0.5445	0.5215	0.5373	0.5915	0.5822	0.6854	**0.6877**	**0.6459**

Table: Performance (Accuracy) comparisons in our benchmark. In each column, **bold** values refers to best performance, while <u>underline</u> values refers to second best.

Conclusion

- Cyber Threat Intelligence – background
  - What is CTI (report)? Why CTI matters?
- MITRE ATT&CK Knowledge Base – background
  - TTP: Tactics, Techniques, Procedure; Hierarchical knowledge
- (M-)LLM for Attack Graph Construction – LLM for CTI – task 1
  - AttacKG+, MM-AttacKG
- LLM for Attack sequence Prediction – LLM for CTI – task 2
  - AttackSeqBench

SMU SINGAPORE MANAGEMENT UNIVERSITY

Future works

1. Integrate more modalities

Natural Language  Image 

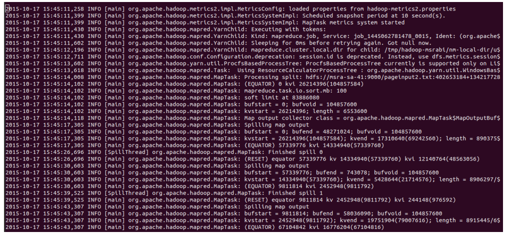
System Log

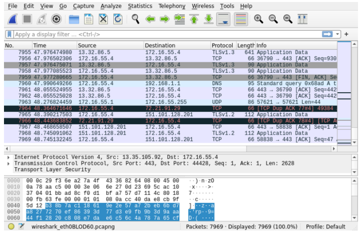
Network Traffic

2. Cross-source verification

- Different sources provide complementary information
- Different sources of data can cross-verify the facts

3. From CTI analysis to CTI generation.

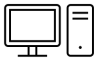
Local file (code, log)

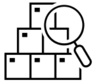
Open-source Intelligence


Domain Knowledge


Human Instruction

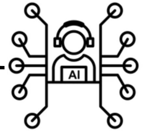
Multi-agent System


CTI Report

SMU
SINGAPORE MANAGEMENT UNIVERSITY

Thank You & QA