Funded PhD or Postdoctoral Research

Title: Decision-Theoretic Planning of Cyberattack and Response with Adversarially Trained Rules

Thesis advisor: Pr. Bénédicte Le Grand, Université Paris 1 Panthéon-Sorbonne

Thesis co-advisor and initial contact point: Dr. Jacques Robin, ESIEA, jacques.robin@esiea.fr

Research project: ANR-FNR funded ANCILE (AutoNomic Cybersecurity with adversarial Learning and Explanations), a cooperation between Université Paris 1 Panthéon-Sorbonne, ESIEA, ENSTAB, Gatewatcher and LIST (Luxembourg Institute of Science and Technology).

Duration: 36 months for PhD, 24 months for postdoc

Starting date: November 2024 for PhD, January 2025 for Postdoc

Monthly net compensation before income tax deduction: 2083€ for PhD, 2500€ for Postdoc

Workplace: ESIEA, Paris Campus (5th arrondissement)

Home office: 65 days per year

Working languages: English and French, B2 proficiency level required in both

Candidate profile

A PhD. position candidate should have a master's degree in computer science or computer engineering. A postdoctoral research position candidate must have a PhD degree in computer science or computer engineering. Due to the multidisciplinary nature of the research, we are looking for candidates with one of the two following profiles: (a) strong background in symbolic *Artificial Intelligence (AI)* and the intellectual curiosity to develop a focused cybersecurity expertise during the research project, or (b) vice-versa, *i.e.,* a strong background in cybersecurity and the intellectual curiosity to develop an expertise in probabilistic symbolic AI and logic programming during the research project.

Research context

To assist experts of ***Security Operations Centers (SOC)*** [1] to better detect and respond to cyberattacks, the next generation of ***Security Orchestration Automation and Response (SOAR)*** frameworks [2] will need to incorporate multiple explainable AI services to keep up with the powerful AI services that are rapidly becoming available at ever lower cost to malevolent actors. The design, prototyping and benefit evaluation of such next generation SOAR framework is the topic of the ongoing collaborative Franco-Luxembourgish research project ***ANCILE (AutoNomous Cybersecurity with Interpretable Learning)*** co-funded France's *Agence Nationale pour la Recherche (ANR)* and Luxembourg's *Fond National pour la Recherche (FNR)*. The thesis or post-doctoral research hereby described is funded by the ANR budget of the project which consortium is composed of *Luxembourg Institute of Science Technology (LIST)*, *Université Paris 1 Panthéon-Sorbonne (UP1PS*), *l'Ecole Nationale des Sciences et Techniques Avancées de Bretagne (ENSTAB)*, *l'Ecole Supérieure d'Informatique-Electronique-Automatique (ESIEA)* and the cybersecurity startup Gatewatcher.

Among the various innovative AI services to be designed, prototyped, integrated and evaluated in ANCILE, two constitutes the scope of the hereby proposed research: (a) ***cyberattack planning*** the action sequence of a multistep persistent, stealth attack[1] and (b) ***runtime response planning*** a sequence of network and service reconfiguration actions mitigating the impact of such an attack. They are both blue-sky research topics and the two faces of the same game-theoretic coin that can be investigated in synergy.

Both these tasks pertain to the most challenging variety of the AI planning task family [3] since their environment is simultaneously:

- ***Adversarial*** in a zero-sum game between an attacker and a responder that use lateral, confounding moves to hide their intentions, rather than attempting to directly achieve their objectives.
- ***Real-time*** with state changes occurring during both the elaboration and execution of the plan triggered both by asynchronous adversary's actions and the normal network activities.
- ***Non-deterministic*** for the effects of attack or response actions due to both the tasks' adversarial nature and the non-attack related, background transient connection failures occurring in any service provided over a network
- ***Partially observable***, again due to the tasks' adversarial nature, as well as the sheer size of a real network state which model cannot be maintained and reasoned upon only at the lowest level of granularity inside which attacks often hide themselves (*e.g.,* bits in packets or binary code)
- ***Relational,*** for many attack and response plans are based on multiple spatiotemporal relations among classes of services and network infrastructure elements, rather than the mere properties of a single network element at a given time point.
- ***Multi-objective,*** for mitigating an attack involves trading-off multiple conflicting objectives, such as maximizing service availability, data confidentiality and integrity (plus physical safety for cyber-physical systems), while minimizing the operational disturbances and cost overhead of the attack response network reconfiguration, that cannot all be fully satisfied by any single possible plan; similarly an attack plan must also trade-off between reaching the attack objective as soon as possible while remaining undetected as long as possible.

Cyberattack planning and cyberattack response planning are best framed as a ***Decision-Theoretic Online Hierarchical Planning (DTOHP)*** problem. *Decision-theoretic* planning [3] [4] [5] is required for multi-objective planning in a non-deterministic environment. *Online* planning [3]. with limited look-ahead, but plan execution monitoring and contingency replanning from updated sensor data is

[1] Often called *Adverse Persistent Threats* [10] in the cybersecurity literature.

required in real-time, non-deterministic and partially observable environments. *Hierarchical* planning [3] is required for environments, such as computer networks, that are too complex to be reasoned upon at a single level of abstraction, but in which knowledge must be decomposed into hierarchies of representation granularities.

The *adversarial* nature of cybersecurity requires knowledge encapsulated in all SOAR AI services, and thus in particular, in attack response planners, to continuously evolve while in operations through adversarial continual dev-ops learning [6]. This is needed to keep up with to new attack plans devised daily by malevolent actors based on social engineering, hitherto unknown vulnerabilities (so called zero-day attacks) or known vulnerabilities for which a patch has not yet been released. While the goal of a SOAR is to help orchestrate and automate the *response* to attacks, in ANCILE, we also include an *attacking* service in it, precisely to adversarially co-train [7] the response services by pitting them against attacking services in series of zero-sum attack-defense simulations. This co-training is the way we propose to leverage AI to stealthily and automatically anticipate part of human attackers' creativity.

Finally, due to multifaceted impact and ramifications that a reconfiguration, responding to an attack, can have on a critical digital infrastructure, the response plans proposed by a SOAR AI need to be clearly *explainable* to the three levels of SoC human cybersecurity experts. They must fully understand the pros and cons of the AI proposed responses, and their possible alternatives, to decide which one to execute either as is, manually altered or not all.

State of the art

Scientific research in decision-theoretic, online and hierarchical varieties of AI planning has a long and storied history [3]. However, integrating the three together with continual, dev-ops, adversarial learning and co-training and applying the whole to cyberattack and cyberattack response planning seems to be a blue-sky research direction that has never been investigated so far.

A preliminary, Google scholar search with the queries:

- *"decision-theoretic" AND ("attack planning" OR "intrusion planning" OR "Advanced Persistent Threat planning" OR "APT planning" OR ((planner OR "planning system") AND (cyberattack OR "cyber-attack" OR "cyber attack" OR "red teaming" OR "red-teaming")))*
- *"decision-theoretic" AND ("mitigation planning" OR "planning mitigation" OR "response planning" OR "planning response" OR ((planner OR "planning system") AND (cyberdefense OR "cyber-defense" OR "cyber defense" OR "blue teaming" OR "blue-teaming"))),*

only yielded single relevant paper related to the research topic hereby proposed: *"Response Planning"*, a book chapter [8] by Musman and Booker. It discusses the effectiveness on midsize networks of a decision-theoretic approach that models response planning as an adversarial *Partially Observable Markov Decision Process (POMDP)* [3]. They present and evaluate the *Automated Reasoning Cyber Response (ARCR)* planner that they implemented on top of the *Approximate POMDP Planning (APPL)* toolkit[2]. ARCR assumes that the attacker's policy is fixed and models it as one aspect of the state transition function of the POMDP. This approach shows the promise of decision-theoretic planning for cyberattack response planning. In the research hereby proposed, we will investigate its main two limitations. The first is the assumption that the strategy of the attacker is known and fixed. We will investigate how adversarial learning of planning rules can overcome it. The second limitation is the propositional nature of the network state representation as bit strings. It does not scale up for realistic operational networks and does not support neither hierarchical planning nor cognitively efficient explanations for SoC operators under the stress of devising an attack response in real time. In ANCILE we will investigate how to overcome it by using a relational first-order representation.

[2] http://bigbird.comp.nsu.edu.sg/pmwiki/farm/appl/

Research questions and hypotheses

To advance this state-of-the-art, the thesis will investigate the following open research questions:

- Q1: What ***Knowledge Representation Language (KRL)*** [15] can be reused as is or extended to uniformly represent models of cyberattack actions and plans, as well as cyberattack response actions and plans
- Q2: What ***Inference Engine (IE)*** [16] [17] for this KRL can be reused as is or extended so as to carry out and explain, in scalable fashion, attack and response DTOHP?
- Q3: What ***Machine Learning Engine (MLE***) [18] can be reused as is or extended to automatically acquire and continually revise models in this KRL for attack and response planning from new data produced either by sensors during operations or by adversarial training simulations?

Previous research suggests considering ***Decision-Theoretic Probabilistic Logic Programs with the Event Calculus (DTPLPEC)*** [11] [12] [13] [14], as the most promising starting hypothesis H1 for Q1 and the *Causal Probabilistic Logic Interpreter (CPLInt)* [19] as the most promising hypothesis H2 for both Q2 and Q3. H1 is based on the demonstrated ability (a) of an LPEC to declaratively and formally represent executable models for hierarchical planning [12], (b) of a DTPLP to do the same for decision-theoretic reasoning [20] and planning [4], (c) of a PLPEC to do the same for uncertain reasoning in complex spatiotemporal relational domains [14] and (d) of such PLPEC to be continually improved by machine learning from new data as they become available during operations [21]. H2 is based on CPLInt's [19] integration, in a single AI toolkit, of IEs for both exact and approximate first-order probabilistic and decision-theoretic inference, as well as MLE for learning both the parameters (*i.e.,* probabilities) and the structure (*i.e.,* logical rules) of a PLP, from training data, under optional background knowledge constraints themselves expressed as a partial PLP.

Proof-of-concept implementation and validation experiments

Evidence for the answers to these questions given in the dissertation will be provided through:

- The implementation of proof-of-concept prototypes of (a) a cyber-attack planning DTPLPEC and (b) a cyber-attack response planning DTPLPEC. Each one will be encapsulated in a containerized web service to be readily interoperable with the other services of ANCILE's SOAR framework.
- The results of the red team cyber-attack *vs.* blue team cyber-defense simulations on ESIEA's SOC, in which:
  - The red team will be assisted successively first by (a) open-source tools currently available as baseline, then by (b) the manually written attack planning DTPLPEC, and finally by (c) its improvement through adversarial ML.
  - The blue team will be assisted successively first by (a) The baseline integration of ESIEA's SOC individual attack clues detector with Gatewatcher's [22], then by (b) the extension of this integration with the manually implemented pipeline of (b1) an individual attack clues detection PLP, feeding (b2) an attack plan recognition PLP, feeding (b3) an attack response planning DTPLPEC, and finally by (c) the improvement of this pipeline through adversarial ML.

References

[1] H. Debar, "Security Operations & Incident Management Knowledge Area, Version1.0.2," 2021.

[2] Cortex, "What is SOAR?," [Online]. Available: https://www.paloaltonetworks.com/cyberpedia/what-is-soar. [Accessed 2023].

[3] S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 4th ed., Pearson, 2020.

[4] D. Poole and K. Kanazawa, "A Decision-Theoretic Abductive Basis for Planning," in *Proceedings of the Symposium on Decision-Theoretic Planning*, Menlo Park, CA, USA, 1994.

[5] R. Mendes, R. Weingartner, G. Geronimo, G. Bräscher, A. Flores, C. Westphall and C. Westphall, "Decision-Theoretic Planning for Cloud Computing," in *Proceedings of the 13th International Conference on Networks (ICN'14)* , Nice, France, 2014.

[6] Z. Chen and B. Liu, Lifelong Machine Learning, 2nd ed., Morgan and Claypool, 2018.

[7] Y. Vorobeychik and M. Kantarcioglu, Adversarial Machine Learning, Morgan & Claypool, 2018.

[8] S. Musman and L. Booker, "Response Planning," in *Autonomous Intelligence Cyber Defense Agent: A Comprehensive Guide*, Springer, 2023, pp. 133-158.

[9] A. a. M. S. Tesmin, "A language for capturing cyber impact effects," The MITRE Corporation, 2010.

[10] The MITRE Corporation, "MITRE ATT&CK," 2023. [Online]. Available: https://attack.mitre.org/.

[11] F. Riguzzi, Foundations of probabilistic logic programming: languages, semantics, inference and learning, Rivers Publishers, 2018.

[12] K. Eshghi, "Abductive Planning with Event Calculus," in *Proceedings of the 5th International Conference on Logic Programming*, Seattle, WA, USA, 1988.

[13] E. Mueller, Commosense Reasoning: An Event Calculus Based Approach, 2 ed., Morgan Kaufmann, 2005.

[14] P. Mantenoglou, A. Artikis and G. Paliouras, "Online Probabilistic Interval-based Event Calculus," in *24th European Conference on Artificial Intelligence (ECAI'20)*, Santiago de Compostela, Spain, 2020.

[15] F. Van Harmelen, V. Lifschitz and B. Porter, Eds., Handbook of Knowledge Representation, Elsevier, 2008.

[16] D. Gabbay, C. Hogger and J. Robinson, Handbook of logic in artificial intelligence and logic programming, vol. 5, O. U. Press, Ed., 1998.

[17] J. Robinson and A. (. Voronkov, The Handbook of Automated Reasoning, MIT Press, 2001.

[18] E. Alpaydin, Introduction to Machine Learning, 4th ed., MIT Press, 2020.

[19] F. Riguzzi, "cplint on SWISH Manual - SWI-Prolog Version," [Online]. Available: http://friguzzi.github.io/cplint/_build/html/index.html. [Accessed 2021].

[20] G. Van den Broeck, I. Thon, M. Van Otterlo and L. De Raedt, "DTProblog: A decision-theoretic probabilistic Prolog.," in *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, Atlanta, USA, 2010.

[21] N. Katzouris, A. Artikis and G. Paliouras, "Online learning probabilistic event calculus theories in answer set programming," *Theory and Practice of Logic Programming,* vol. 2023, no. 2, 2023.

[22] Gatewatcher, "AIONIQ: Behavioral and mapping analusis for augmented detection," 2022. [Online]. Available: https://www.gatewatcher.com/wp-content/uploads/2022/06/Datasheet_Aioniq_2022.pdf. [Accessed 2022].

[23] Dive Research, "Global Advanced Persistent Threat Protection Market Analysis," [Online]. Available: https://www.prnewswire.com/news-releases/global-advanced-persistent-threat-apt-protection-market-predicted-to-generate-a-revenue-of-20-290-7-million-at-a-cagr-of-20-9-from-2020-to-2027---exclusive-report-245-pages-by-research-dive-301364279.html. [Accessed 2023].