

Technical Disclosure Commons

Defensive Publications Series

February 04, 2019

VIRTUALIZED INTELLIGENT HONEYPOT AGENT

Plamen Nedeltchev

Mani Kesavan

Hugo Latapie

Enzo Fenoglio

Follow this and additional works at: https://www.tdcommons.org/dpubs_series

Recommended Citation

Nedeltchev, Plamen; Kesavan, Mani; Latapie, Hugo; and Fenoglio, Enzo, "VIRTUALIZED INTELLIGENT HONEYPOT AGENT", Technical Disclosure Commons, (February 04, 2019)
https://www.tdcommons.org/dpubs_series/1928



This work is licensed under a [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).

This Article is brought to you for free and open access by Technical Disclosure Commons. It has been accepted for inclusion in Defensive Publications Series by an authorized administrator of Technical Disclosure Commons.

VIRTUALIZED INTELLIGENT HONEYPOT AGENT

AUTHORS:

Plamen Nedeltchev

Mani Kesavan

Hugo Latapie

Enzo Fenoglio

ABSTRACT

A honeypot system is described that can expand to any attack surface as it learns and grows with the changing device landscape. The system also takes into account the human elements that originated the attack. By using adversarial training mechanisms, the system may be quickly trained to become a doppelganger and attract attacks. Moreover, a unique quantum cognitive framework provides a robust adaptivity to ever-changing attacker strategies. Virtualized intelligent honeypot agents may be introduced into the network, device, or server, to connect and share knowledge to facilitate federated learning for similar type of agents. The agents may also be operated in multitasking for many similar types of devices, users, applications, and the like.

DETAILED DESCRIPTION

Honeypot is a well-known security mechanism for luring attackers, as well as studying and countering attack attempts [6]. However, with the advent of the Internet of Things (IoT) and wearable devices, and the ever-changing technology landscape, the attack surface is expanding. It is becoming increasingly difficult to adjust and create convincing handcrafted honeypots. IoT devices are not scanned by legacy cyber defense systems in-depth, and remain prime targets for attackers within the network. Also, honeypots that are set up at a demilitarized zone or periphery of the network do not cover vulnerabilities that occur inside the network through wearables and other devices that enter the network. Recent issues with ransomware highlight the fact that the attack surface is expanding with devices that can enter and exit a network. It would be useful to have an intelligent system that can adapt to this expanding attack surface, cater to different areas of that surface, and make proper counterattack decisions even under a given level of uncertainty.

Cyber attackers are improving at detecting conventional honeypots. Combining the ability to open attachments, interact with websites and other online services, and interact

with IoT devices, provides a rich contextual environment which makes client honeypot detection that much more difficult.

A smart attacker knows that a simple or direct attack is unlikely to be effective. Hence, s/he may try to deceive the system by mixing up actions in a rather sophisticated plot through machine learning techniques and opportunistic decisions. For example, s/he can act stealthily by waiting for an opportunity (passive attacker) or can launch an attack immediately (active attacker), as described in [1] and illustrated in Figure 1 below. Moreover, the defender faces an unknown attacker with unknown intentions where decisions can be made dynamically on observable events that may be non-separable (entangled) and non-commutative, since they may depend on the order in which they are presented.




Figure 1 - Attacker and defender in the network

Accordingly, provided is a mechanism to identify potential attackers laterally entering the network using IoT and wearables. The system injects intelligent doppelganger agents running in a virtualized sandbox for any IoT device, application, or user that would normally be allowed into a network. This may include wearable devices, mobiles, computers, user accounts, etc.

In one embodiment, the system locates a subject device or user and provides an associated agent. For a new device like a user laptop, an agent associates and acts as a honeypot agent for that user account. While many Mobile Device Management (MDM) and Information Technology (IT) systems already inject support agents, this system supports agents dedicated for the honeypot attack surface. The same agent then acts as an inciting agent in an adversarial reinforcement learning system, where the agent is trained

to operate in the presence of a destabilizing adversary that applies disturbance forces to the system [5]. The purpose of this function is to imitate and fake device functionality while appearing more attractive for being probed, attacked, or compromised than production systems that have value. This agent constantly learns how to better mimic a user's interaction with the system. The agent may employ generative adversarial networking principles of deep learning to imitate and adapt to any device, user, or application. The agent changes its personality dynamically, rather than based on a fixed set of known vulnerabilities (e.g., manually coded for the fixed set of known vulnerabilities) as is typically performed on existing morphing honeypots systems.

In a second embodiment, for more sophisticated attackers with unknown intentions and limited information, the classical Bayesian theory or expected utility maximization could be applied, despite leading to overlooking and oversimplify the essence of the attack. Actually, within the classical rational model of decision-making, it is assumed that decision makers (honeypot agents) comprehensively define a problem, understand all possible alternatives and their consequences, and select the very best action after evaluating all available options. On the contrary, the present system considers some paradoxical aspects of real systems in which (human) attackers operate, which are characterized by limited cognitive resources, uncertainties, and complexity [2]. In particular, environments may be characterized by indeterminacy effects (e.g., the state of the system does not determine a unique collection of values for all its measurable properties), complementary effects, interference effects (e.g., probing a honeypot virtual agent changes its state), violation of total probability effects (e.g., prior probability distributions may be zero while posterior distributions exist), order effects (e.g., probing two agents may lead to different measures depending on the order), entanglement effects (e.g., agent states for the same IoT device class may not be separable), and the like. To model an attacker ingenuity, the system may use heuristics and intuition. However, it is very difficult to find the heuristic that works in every case, especially when a (human) attacker is involved. In this case, a quantum probability framework is preferable for building a corresponding quantum cognitive model for decision-making that resolves the irreconcilability between observed features and classical ideal models.

This formalism is already used in social science and behavioral economics as a probabilistic tool [2]. The use of quantum techniques requires that neither the brain nor consciousness would have anything to do with genuinely quantum systems. The techniques of quantum theory are used in this context solely as a convenient and efficient mathematical tool and language to capture the complicated properties associated with the decision-making process [4]. Actually, the honeypot virtual agents use quantum probability as a framework to model uncertainty for a compound entangled system where agents and environment are inseparable, and to model how agents reflect upon their strategies for adapting to changing environments where defender decisions are made dynamically on observable events.

The present system targets IoT devices in the network by associating each device with a virtual attractive (i.e., easily attacked) smart agent to form a meta-network that continuously locates and quarantines new threats by exchanging information and improving future counterattack strategies through deep adversarial learning and a reasoning engine. The system may build knowledge from learned strategies observed on attackers shared among all agents by making small changes to data samples. This model is predictive. Prediction is based on the generative adversarial networking principle of deep learning to generate adversarial behavior that imitates a device and learns new counter-strategies. For more sophisticated protection, agents may not maximize their utility, but instead adapt their behavior to the incomplete information acquired from other agents and the environment within the limited resources they have. Non-Bayesian quantum probability together with a reasoning engine may be used to describe this process of adaptivity [3].

Three examples are provided as follows. First, if an Active Directory (AD) management system comes online, an agent associated with that system may begin faking group or user creation vulnerabilities. Second, for a personal device, the agent may select links and attachments, fill in forms with its (fake) “personal information,” and otherwise mimic a human who is completely unaware of good security practices. Third, for any device, an agent may add time warp and execution environment analytics (e.g., change in code signatures, memory, processes, startup code, sensitive storage areas, etc.). This addresses dormant ransomware and other hacks, as the sandbox agent may intelligently

proceed in time in a manner that would be difficult to determine because the ransomware could try to detect such time warps.

Since attackers are always searching for different vulnerabilities, they can be lured into such an agent pretending to be a device. Since the agent is constantly learning and updating its behavior, the attacker may be fooled into believing that the hard-coded fake honeypot device is a transaction system. The quantum cognitive model provides an augmented framework to estimate probabilities for correct decision-making strategies and human probabilistic inferences.

A general decision-making formalism is provided which includes machine (e.g., honeypot agents) and human (e.g., cyber attacker) elements. Artificial quantum intelligence may enhance the quality of the decision-making capabilities under environmental uncertainty and incomplete information. This demands adaptive behavior by virtual agents to respond to the changes in complex dynamic environments.

Every attempt to understand the environment is an interaction (i.e., a measure) that in turn affects the environment. The interaction of a cyber attacker between the environment and the other virtual agents is a generative process to acquire knowledge. The generative process takes place as a continuous interaction with the environment. Due to the uniqueness of every virtual agent, each generative process is unique, but decisional conflicts may emerge. Current approaches based on classical Bayesian theory oversimplify the nature of the decisional conflicts, and as a result, solutions are also oversimplified or, more likely, prone to being wrong or ineffective (i.e., need to be constantly adapted by proper heuristics) to defeat attackers.

As opposed to classical Bayesian approaches [1], quantum cognition [2],[3],[4] provides axiomatically coherent answers to decisional conflicts among virtual agents under incomplete information and uncertainties.

Failure to recognize that defender/attacker interactions occurred in an uncertain environment may lead to serious consequences, including: (1) lack of interoperability; (2) information vulnerabilities, which can impede the operation environment because of the dependencies on the decision aid tools; and (3) impeding of adaptive behavior due to inadequately available heuristics.

A non-Bayesian cyber-attack game (quantum cognitive model) of incomplete information reflects the defender (virtual agent)'s imperfect knowledge of incoming attacks (e.g., malicious or not). This one-shot game model addresses the basic questions as to how the defender should dynamically predict new scenarios/observations for the cyber attacker, and whether deception is optimal for both the attacker and defender without using knowledge-based techniques. This may be performed with the generative adversarial model.

These techniques use generative adversarial learning and imitating devices, users on production network to newer device personalities, and newer user account baits for offering up both Personally Identifiable Information (PII) and network system vulnerabilities.

In summary, a honeypot system is provided that can expand to any attack surface as it learns and grows with the changing device landscape. The system also takes into account the human elements that originated the attack. By using adversarial training mechanisms, the system may quickly train to become a doppelganger that can attract attacks. Moreover, a unique quantum cognitive framework provides a robust adaptivity to ever-changing attacker strategies that adequately model human causal reasoning. Virtualized intelligent honeypot agents may be introduced into the network, device, or server, to connect and share knowledge to facilitate federated learning for similar type of agents. The agents may also be operated in multitasking for many similar types of devices, users, applications, etc.

REFERENCES

- [1] Quang Duy La Tony Q. S. Quek Jemin Lee "A game theoretic model for enabling honeypots in IoT networks" 2016 IEEE International Conference on Communications (ICC).
- [2] White, L. C., Pothos, E. M., & Busemeyer, J. R.(2015). Insights from quantum cognitive models for organizational decision making. *Journal of Applied Research in Memory and Cognition*.
- [3] A quantum probability framework for human probabilistic inference. Trueblood JS, Yearsley JM, Pothos EM. *Journal of Experimental Psychology Gen.* 2017.

[4] Vyacheslav I. Yukalov, Didier Sornette: Quantum Probabilities as Behavioral Probabilities. *Entropy* 19(3): 112 (2017).

[5] L. Pinto, J. Davidson, R. Sukthankar, and A. Gupta, “Robust adversarial reinforcement learning,” *CoRR*, vol. abs/1703.02702, 2017.

[6] F. Cohen. The use of deception techniques: Honeypots and decoys. In H. Bidgoli, editor, *Handbook of Information Security*, volume 3, pages 646–655. Wiley and Sons, 2006.