Warning and management of cyber threats by a hybrid AI system (robot and operator)

Isaac Faber[a], **Elisabeth Paté-Cornell**[b]
[a]U.S. Army AI Integration Center, Pittsburgh PA, USA, isaacfab@gmail.com
[b]Stanford University, Stanford, CA, USA. mep@stanford.edu.

Abstract: This paper presents a warning system and risk management model, in which early signals of cyber threats are generated using machine learning and artificial intelligence. Cyber threats and attacks can be modeled as a set of discrete observable steps called a "kill chain". The data support automatic defensive responses based on decision analysis and machine learning, before losses occur. A hybrid AI system (robot and human being) guides decisions to open or close gates in a system, based on attack signals at the beginning of the kill chain.

1. INTRODUCTION TO CYBER RISK

Organizations must defend their information systems from threats of cyber-attacks because they rely on these systems for operations and the storage of confidential/sensitive data. A probabilistic analysis of cyber risks based on systems' analysis and probability allows characterizing the risk and guiding risk management decisions [1]. Threats from attackers need to be recognized and blocked. At the same time, legitimate customers need to have access to that system. The question is to distinguish between the two cases. As in classical engineering risk analysis, defended systems can be decomposed into their parts, but cyber analysis must focus primarily on the threat itself. Here, the cyber risk management domain includes three distinct categories of actions:

- Offense i.e., activities involving exploitation or attempted exploitation of unknown systems [2].
- Active defense involving a dedicated work force and tools that control and update the defensive posture.
- System Operation and Development involving defensive security practices such as password change, token management, software design, and network design.

In practice, cyber security efforts are executed after damage has occurred [3]. Enhancing these cyber security practices involves improving both network and software (host) security. General network defense includes cryptography, internet-security protocols, user authentication, firewalls, and intrusion protection [4], [5]. Software security means better software design [6], [7], [8]. To be most effective, defenders must be able to identify hostile activity at an early stage and to obtain warnings with enough lead time [7].

Defenders continuously react to new information and update their organization's security policies. What we propose here is the integration of machine learning techniques, expert analysis, and decision making. Warning systems are built using machine learning models based on threat actors' behavior, a digital fingerprint recorded security device databases [9], [10]. These behaviors are identified via forensic activities, however, network defenders are often overwhelmed by the high prevalence and ignore them [9]. However, in recent years, the increase in computing power has made it possible to integrate Artificial Intelligence (AI) and human operators in hybrid systems [10]. This integration allows for the scale needed to properly react to digital all relevant forensic evidence.

A system analysis allows identification of a 'gate-set', defining access control decisions for the defended system which balance security and effectiveness. The defender controls a series of gates, which they can either leave open or close. The goal of the system's defender is to find an optimal gate-set policy (open or close the gates inside the system) as quickly as possible. The model in this paper thus includes two decision analyses: (1) for the robot to decide to close a gate or to pass the hand to the human-in-the-loop with a risk attitude programmed in the machine algorithm, and (2) for the operator either at the beginning to train the robot, or when later called upon by the robot to decide. It is assumed here that the risk attitude is the same for the robot and the human expert.

2. BACKGROUND

In recent years, nation-states as well as private actors have been acting both as defender and attackers. The United States created a Cyber Command, and currently, over 60 countries have commissioned similar military units [11]. Active cyber attackers also include activists, criminals or state-sponsored non-government groups. For instance, the international hacker community known as Anonymous has taken credit for a broad range of offensive activities from defacing websites to directing attacks [12].

It is a challenge to secure cyber systems [13]. Uncertainties abound; however, models can help. Cyber models are abstract representations of the real world. Three common types of cybersecurity models are physical, graphical, and mathematical.

The physical model is a scaled down representation of an information network, e.g., the National Cyber Range of the United States [14]. A graphical model is commonly used to represent the general system architecture, but it is limited to physical topology. For instance, UML and SysML provide a graphical foundation and basic interface definitions but do not include the semantics (language) level and the fundamental uncertainties associated with the cyber domain. Mathematical models can represent information, inter-connections, and semantics in an abstract quantitative form. Automatically generating attack graphs is the most popular method of modeling cybersecurity systems using mathematics. Significant research on the application of quantitative methods to graphical models has been explored in the literature [15],[16].

In this context, a graph is a network of nodes and edges, rather than a visual network topology. New models are needed to include complex interactions and actors, as well as actionable parameters for risks and decisions. Network systems are deployed to maximize the defenders' visibility and situational awareness in order to do threat detection and response. Security devices are commonly placed at the boundary of a defended system that routes inbound and outbound traffic through a set of monitored devices. For example, firewalls are a common device found in a security stack.

Typically, threat detection is followed by a defensive response involving a change in access to the system. Two approaches are used for access control: blacklisting and whitelisting. Blacklisting is the practice of rejecting something if the sensor rule is met. Whitelisting means rejecting everything that does not match a rule. Firewalls can restrict networks from communicating directly with external geographical locations (or entities) that the defender considers of high risk. Such activity restrictions are commonly referred to as "blocking". Whitelisting determines access based on a set of known, non-threat signatures, thus reducing the value of a system for entities that interact with the general population.

One common approach to standardization of incident response (the process of updating access control) and study of events is to use a framework such as the 'cyber kill chain' [6], [7], developed by Lockheed Martin Corporation (see Table 1). The lead time for the first 3 signal types are weeks, days and minutes; zero for the others.

The use of machine learning facilitates the automation of signature recognition and threat detection and response by relying on historical data.

Table 1: Cyber Kill Chain Stage

Steps i,	Examples of signal
1.Recon.	Boundary access behavior
2.Weaponize	Rate of vulnerability of systems
3.Deliver	Virus detection in IDS systems
4.Exploit	Anomaly in behavior
5.Control	Improper access control us
6.Execute	Anomaly of data transmission
7.Maintain	Pattern in external communications

Behavior-based detection is one solution to the limitations of black- and white-listing. One common approach similar to behavior-detection techniques is anomaly detection, similar to statistical hypothesis testing [17]. In this paper, a broader interpretation of behavior-based techniques allows using robust machine-learning methods. A small sample of recent methods includes regression, deep learning [18], [23], clustering [24], and Bayesian networks [4], [5], [15]. The following section presents a decision support model, which leverages machine learning to identify and warn of threat actor behaviors early in the kill chain.

3. MODEL OF SYSTEM DEFENSE OPERATIONS/DECISIONS

This model is meant to support the decisions of the system defender. It is limited to a single organization, with a specified risk preference, conducting active defense through a "super-agent" combining a machine and a human being in which the robot can either make an automatic decision or pass the hand to the human in the loop (Fig. 1).

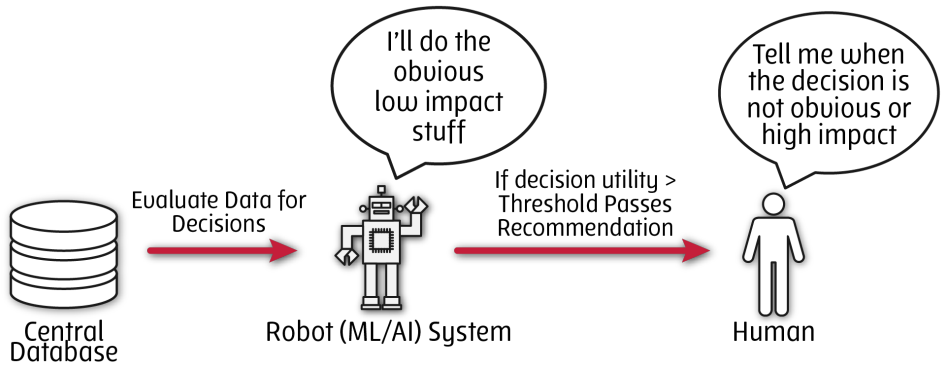

Figure 1 The "Super-Agent": Robot and Human Cyber-Defense Policy Making

The objective is to identify and deploy an optimal defensive policy consisting of a set of decisions that maximize the defender's expected utility. The goal in practice is to allow legitimate actors to use the system and to block threat actors. The defender must thus develop a defensive posture that balances uncertainties, allowing entrance to as many customers as possible while blocking likely threats.

The decision space for the system's defender consists of a series of blocking actions (closing gates). The uncertainty is that closing selected gates may block legitimate network activity, resulting in a cost to the defender's organization Figure 2 shows a simple view of the defender's decision space over a single discrete confidence. Signals (or warnings) of possible attacks at each gate are generated by machine learning from time interval governed by the kill chain, Timing matters. Optimal policies involve acting at early stages, when signals can be observed with sufficient historical data updated as new information is acquired, real-time context, and threat intelligence.

Figure 2: Cyber security decisions as opening or closing gates

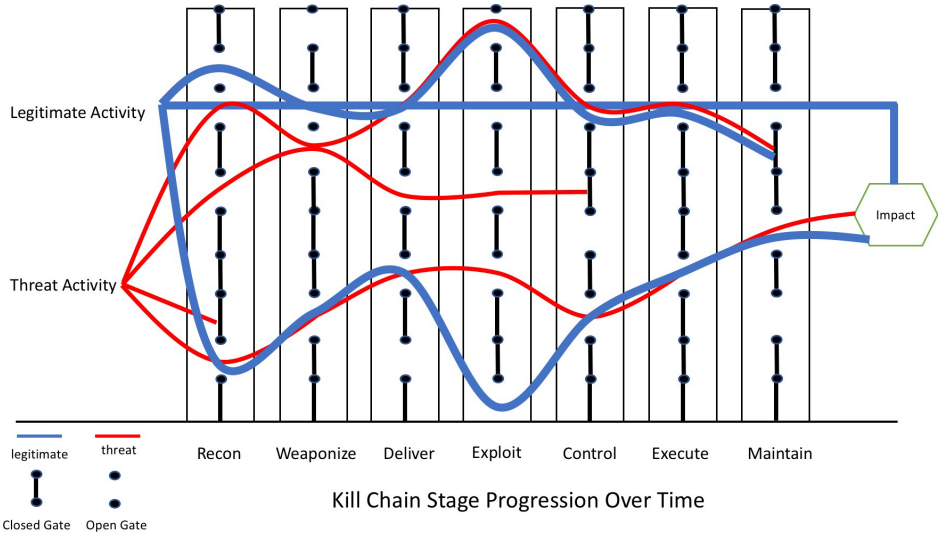

Three key elements of the model are thus: the Bayesian updating that allows the machine to learn as new observations are added to the system; the encoding of the prior probabilities that each possible type of actor may attempt a cyber-attack; and the choice of a type of distribution for the probability (prior and posterior) of each key factor affect the decision. The Bayesian updating adopted here consists in choosing a prior for each factor and a probability distribution for the likelihood of the signals.

3.1 Defender's decision model: Influence diagram

The purpose of an influence diagram is to represent the uncertainties of the systems components' performance and their dependencies in the decision to be made (here, the defensive response). Figure 3 represents the decision model of the defenders as an influence diagram.

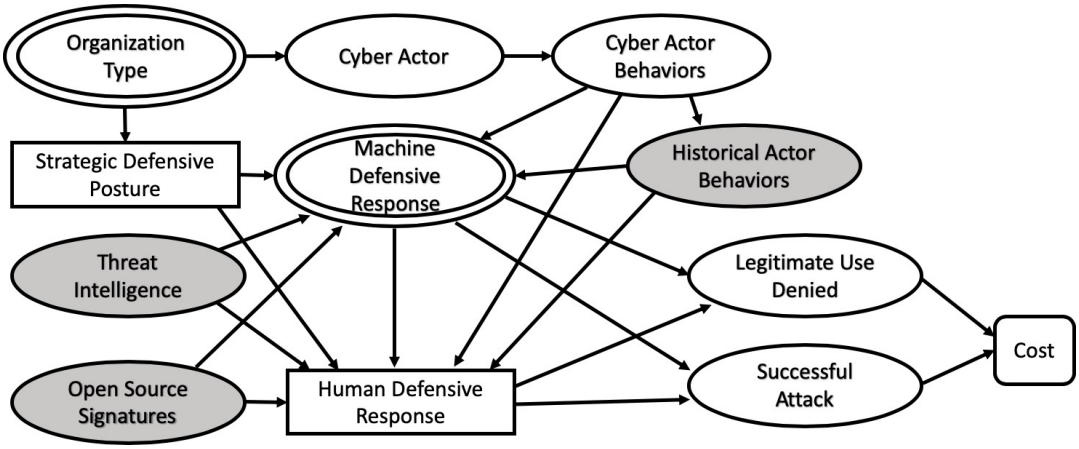

Fig. 3: Influence Diagram for a Cybersecurity Defender

.

Figure 3 yields the risk of a cyber-attack at a given time and the related uncertainties from the defender's point of view. It includes decisions, random variables, events, and their dependencies. The nodes are deterministic double-ovals (DN), rectangular decision nodes (DCN), or oval for random variables and events (RV). The arrows represent dependencies among nodes

The nodes of the influence diagram are:

- Organization Type: Type of Organization, i.e., Government, Small business. (DN)
- Strategic Defensive Posture: The decision on technical defensive capabilities.(DCN).
- Cyber Actor: Actor (good or bad) that interacts with the defended system. (RV).
- Cyber Actor Behavior: A set of the behaviors observed from the security stack. (RV).
- Threat Intelligence: Near-term information gathered outside the defended system, about threat actors. (RV).
- Open-Source Signatures: Near-term open-source information collected about threat actors. (RV).
- Historical Behavior: Long-term observed behavior associated with specific actors in the defended system. (RV).
- Specific Behavior: current combination of attributes for behaviors of actors in real time. (RV).
- Cyber Threat: The short-term probability that the current actor is a threat (RV)
- Machine Defensive Response: The decision of the system defender. DN given the rules encoded in the software.
- Human Defensive Response: The decision of the system defender (DCN).

- Successful Attack: The uncertainty about the success of an attack if the actor is a threat given the signals. (RV).
- Legitimate Use Denied: The uncertainty about access denial to a legitimate actor (RV).
- Cost: Incurred by the defender (RV)

The uncertainty about cyber actors is influenced by the defender's type of organization. An observed behavior, along with external and historical information, determines the success of an attack, as well as a denial of legitimate use, determine the cost, or the effect of errors on the defended system.

3.2 Strategic Defensive Posture

The mathematical model is based on an expected utility optimization over time through the tuning of a set of decisions. Circumstantial conditions can support policy development governed by long-term strategic choices made when designing the information system, including:

- Central database design
- Defensive security-stack technology
- Risk-tolerance thresholds encoded in the system.

Strategic defensive postures are based on thresholds of acceptable risk. Decisions include the purchase of risk management technologies (the security stack) such as Intrusion-Detection-Systems (IDS), firewalls, and proxies. The amount and the type of data collected from the stack are determined by the choice of a database technology.

In general, the more significant the investment in the security stack and central repository, the greater the decision space and amount of data available to the system defender.

The strategic posture of a defender defines the decision space for responding to events. The more resilient the defensive system, the greater the number of possible actions after or during an attack. At the most basic level, a defensive system typically involves three key features:

- Sense: Identify threats through signatures
- Block: Stop threat activity by breaking the kill chain
- Mitigation: Respond to damage done after a successful attack

Given these three capabilities, a defender may be able to identify an attack, know where it is coming from, stop the threat from accessing a system, and respond when attacks take place. Blocking of components represents the gate-set defined earlier in Figure 2

Internal policies and guidance complement the technical capabilities of a security stack and most organizations can modify systems under threat [21]. Decisions to change the infrastructure are often made late in the kill chain when severe steps have been identified because the cost of these changes can be significant and often require considerable manual effort.

3.3 Defensive Response

The rules of cyber engagement vary significantly depending on the timing and mission of the attacked organization. Legal considerations in the US restrict private entities from engaging in offensive cyber operations [14], but nation-states can leverage advanced techniques to conduct improved defense, including intelligence gathering and deterrence. Private entities can purchase intelligence from third parties if this information is collected legally but this is exceptional and defensive response is often limited to internal mechanisms.

With an optimal defensive posture, certain types of system access can be granted or denied. In a common approach, session-level data are used to make this determination. Source/destination, IP addresses or time of the day can be used as decision support. With more advanced systems, session-level data can be enriched by external data. These can include geospatial information or behavior information such as multi-session temporal patterns. Access control decisions can be made in near-real time when allowed by defensive strategic posture decisions, but the posture changes at the pace of the organization's requirements and purchasing systems [17]. Defensive response decisions can be considered in the context of kill-chain lead-time. As warnings generated by the system defender are observed, a gate policy is formulated to break the kill chain in its earliest stage. In practice, the kill chain is broken by closing a previously open gate. The earlier the chain is broken, the smaller the damage.

3.4 Cyber Actors and behavior types

As described earlier, cyber actors can be classified as follows:

- Legitimate: Administrators, developers, users/ customers, incidental
- Threat: Competitors, Data collectors, Criminal, Hackers, Foreign nations, Insider Threats

Each of the actors described earlier shows a set of behavior types associated to their interaction with the defended system. Behavior types are discrete categories, each of which allows a quantifiable observation. Gates correspond to specific behaviors. When a gate is open that behavior is permitted, when the gate is closed it is not. Unlike traditional signatures, behaviors are general characteristics, which are not automatically identifiable as belonging to a specific actor type. These include:

- Signature: IP address, MAC address, geo-location
- Temporal: Inter-arrival time, Time of day, Pattern of life
- Frequency: Top talkers, Infrequent visitors
- Other: Amount of data sent, Type of interactions

In this perspective, an actor produces a digital, behavior-based signature. The behavior types are inputs to a machine learning model and signatures are generated as outputs which become alerts in the warning system. In the AI construct, the machine will learn to categorize behaviors through the use of historical data and Bayesian updating. This learning is then used to update the gate-set policy and improve the super-agents performance.

3.5 Open Source Signatures; Threat Intelligence

The condition of general (global) cyber-threat activity at a given time is fluid and based on several factors: impacts on the defender of known and unknown software vulnerabilities, organizational positioning, and hacker capabilities all contribute to trends that impact the effectiveness of a defense. In most organizations, information is generated internally (threat intelligence) and collected externally (open source or shared signatures). Open-source security communities or custom-developed threat intelligence thus affect directly the defender's posture. These sources of data are quantifiable and actionable and these quantities, if they are observed, can support security decisions.

3.6 Decision Formulation

Each gate g has an unknown probability distribution of a visit (f_g) by an actor i with utility $u(g, i)$ based on the risk attitude γ of the actorand his reward for penetrating gate g. In the mathematical formulation of the decision analysis for the system manager, the notation used is as follows:

Table 2. Model Notation

Notation	Description
G_t	Gate-set policy
g	Gate-set policy for a specific gate
y_t	Observed actor type
r_t	Realized reward from actor i
γ	Risk attitude of the decision system
$u_t = u(g, y_t, r_t, \gamma)$	Utility of the decision system
e_t	External Information
y_{t-1}, r_{t-1}	Historical data used to train the machine
i	Index of actor type
$\hat{u}(g,i)$	Statistical estimate of the utility of the decision system for a specific gate and a specific actor
$\hat{p}_{g,i}$	Statistical estimate Of the probability of arrival Of a specific actor type at a specific gate

At each time t, the defender of an open gate has a single opportunity to observe an actor, whose type is described by a multinomial distribution, and a continuous reward distribution. The system defender updates the probabilities p_i of actors' intrusions over time, which allows him to make the best gate-set policy decision given his current state of knowledge. Figure 4 shows the current defensive process and the mathematical factors associated with the logic flow.

The utility of the hybrid decision system for the rewards r_t of penetrating gate g is evaluated as a function of its risk attitude γ =-u"/u' [22] . It is assumed that this risk attitude is the same for the machine software and for the human decision maker. Equation 1 shows the functional form used in this paper for both attackers and defenders, assuming a constant risk attitude with respect to the outcomes r_t (risk neutrality).

$$u(g,i,\gamma,r) = -e^{-\gamma r} \qquad (1)$$

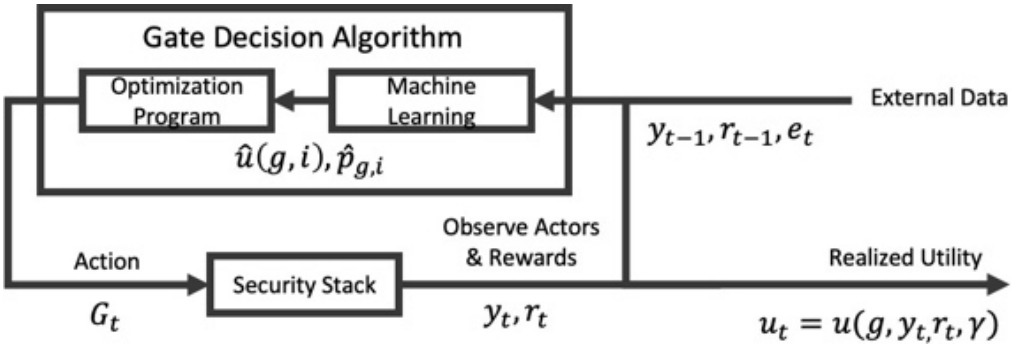

Fig. 4: Logic flow in the Mathematical Model

Both the defender's probability and current belief about utility distributions for each gate are outputs of the learning process. For the system defenders to learn, they must receive a consistent stream of new information. Therefore, their knowledge and utilities depend on the time *t*.

3.7 Machine Learning and Bayesian Updating

Each scenario is guided by the learned probability about a particular gate's state and the super agent needs to decide whether to keep it open or to close it. Figure 5 represents the decision tree.

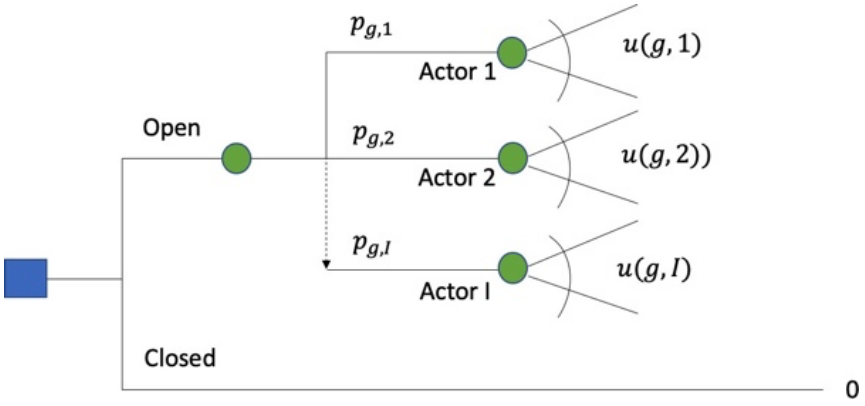

Figure 5: Decision Tree of the Super-Agent For the Management of Each Gate

There are two distribution types for each gate, one for the actor types, (f_g), and one for the rewards expressed as the realized utility $(u(g,i))$. The distributions f_g and $u(g,i)$ are uncertain and the parameters need to be learned. As shown in Figure 5, a critical element is the ability of the robot to pass the hand to the human actor.

There is a trade-off in the decision cycle between exploring to gain new information and exploiting what is already known. Therefore, in order to find a global optimal policy, the system defender must take risks and explore other alternatives.

Bayesian techniques allow both exploiting and exploring. Exploiting means taking advantage of the best alternative implied by estimates from observed data. Bayesian techniques leverage conjugate distributions such as Beta/Binomial conjugates that yield posteriors of the same for as priors.

Exploring can be done by sampling from a prior distribution at each time period, then using those realizations to build a gate policy. Over time, the priors are updated into posteriors, generally narrowing the variance of the parameter estimates.

Human intervention thresholds need to be set to direct the robot to call on the human being. The "robot" may act based on the automated system in routine situations. Beyond a certain threshold of loss magnitudes or uncertainty, the robot should not be trusted to act autonomously and needs to pass the hand to a human. Clearly, these two factors are not independent. Human analysts have contextual understanding of the system and external demands that the machine does not. Appropriate thresholds ensure that an organization is conscious of consequential actions.

3.7 Optimization of Gate Set Policy

Using information from the learning process, the optimization yields a vector representing the best security policy. The vector $\boldsymbol{G_t}$, at given point in time, represents $\boldsymbol{m}$ gates with binary values. 1 represents an open gate and 0 a closed gate.

$$\boldsymbol{G_t} = \begin{bmatrix} \mathbf{1} \\ \mathbf{0} \\ \mathbf{1} \\ \vdots \\ \boldsymbol{m} \end{bmatrix} \qquad (2)$$

The general formulation of the optimization model is the following:

$$\boldsymbol{maximize} \quad \boldsymbol{G_t U_t} \qquad \boldsymbol{subject\ to} \quad \boldsymbol{G_{g,t} = 1 \ \forall \ g \in S} \qquad (3)$$
with $\boldsymbol{g \in [0, 1]}$

$\boldsymbol{G_{g,t}}$ is a value in vector $\boldsymbol{G_t}$ and $\boldsymbol{S}$ is a security policy (decision option) for a specific gate as dictated by the exploring algorithm. The vector $\boldsymbol{U_t}$ is the defender's utility for each gate at time t. The overall goal of the model is to find the optimal defense policy, $\boldsymbol{G_t^*}$ (open or close gates at time t).

3.8 Results

As an example of the model, consider a cyber defender controlling two points of entry s1 and s2, two destination ports (d1, d2), two destination IP addresses (i1, i2), and two times of the day (work hours: w and non-work hours: nw). In this case, the defender has 16 combinations of behavior attributes (options), which constitute the gate set. Table 5 shows a summary of the gate-set with additional information about the distributions of actors and their utility at each of the gates. If the defender has perfect information about the utility/actor distributions, his optimal policy is to keep eight of the gates open and the other eight closed.

Table 5: Gate sets and threat probabilities (normal distributions (10, 2)), Expected utility (Eu(g)) and optimal decision

Gate set	Threat probability	Eu(g)	Best action
s1 d1 i1 w	0.1	8	open
s1 d2 i1 w	0.1	8	open
s2 d1 i1 w	0.1	8	open
s2 d2 i1 w	0.1	8	open
s1 d2 i1 nw	0.2	6	open
s2 d2 i1 nw	0.2	6	open
s1 d1 i1 nw	0.3	4	open
s2 d1 i1 nw	0.3	4	open
s1 d2 i2 w	0.5	0	close
s2 d2 i2 w	0.5	0	close
s1d1 i2 w	0.6	-2	close
s2 d1 i2 nw	0.6	-2	close
s1 d2 i2 w	0.6	-2	close
s2 d2 i2 nw	0.6	-2	close
s1 d1 i2 nw	0.7	-4	close
s2 d1 i2 nw	0.7	-4	close

Using the information about risk and uncertainty from the learned distributions, the decision threshold where the human takes over from the robot can be defined. Figure 6 shows an example that shows expert and robot decisions during the machine learning process.

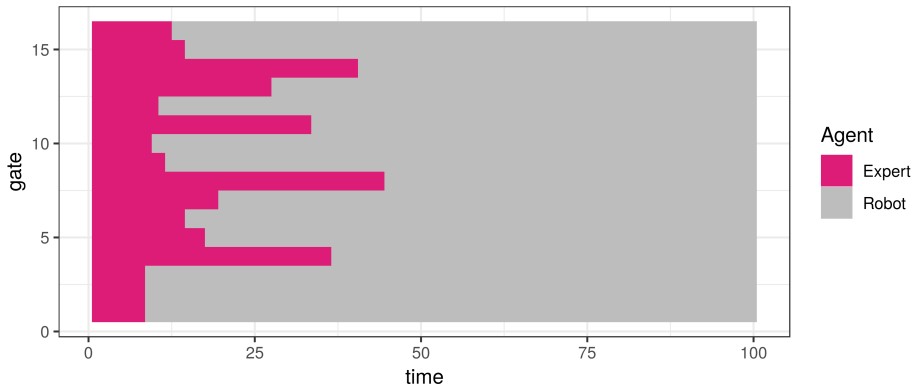

Figure 6: Example Showing When the AI system Starts Making the Decision vs Human Expert

In the beginning, as information is scarce, the human makes all decisions, and the robot learns from them. As more data are collected and uncertainty is generally reduced, the AI-based robot makes decisions more frequently. The highest and lowest expectation gates (1,2,15,16) are learned quickly enough to be automated in a few time periods; but the gates towards the middle, which are harder to understand (near 0 expectation) require much long expert-human involvement.

4. CONCLUSION

Cyber risk comes from threats of attack to the organizations that rely on information systems for operations and for the storage of organizational/trade/industry secrets. The gate-set concept provides a framework to issue warnings of cyber-attacks and quantify a threat impact based on the behavior profiles of potential attackers. This paper presents a method to evaluate actor behavior/actions to recommend an optimal security policy. The gate-set concept can be used in communication to non-experts, and allows an independent evaluation of security policies. The cybersecurity environment can thus be improved using artificial intelligence technologies. Large datasets generated by security stacks are evaluated and updated. Intelligent access control leverages the strengths of the super-agent (including both AI/robot and expert/human decision makers) in order to create optimal gate-set policies. This approach gives system defenders an effective tool to conduct cyber risk management under uncertainty, involving both an AI system and a "human in the loop".

References

[1] Paté-Cornell, M.E., M.A. Kuypers, M.D. Smith and P.J. Keller. "Cyber Risk Management for Critical Infrastructure: A Risk Analysis model and Three Case Studies", Risk Analysis, DOI: 10.1111/risa.12844, Internal article ID: 14264471. Risk Analysis, Vol 38, issue 2, pp. 226-241, February 2018.
[2] Li W. Li and R. B. Vaughn. Cluster security research involving the modeling of network exploitations using exploitation graphs. In Cluster Computing and the Grid, 2006. CCGRID 06. Sixth IEEE International Symposium on, volume 2 pages 26–26. IEEE, 2000
[3] Hill D. and J. T. Lynn. Adaptive system and method for responding to computer network security attacks, July 11 2000 US Patent 6,088,804.
[4] Xie, P. J. H. Li, X. Ou, P. Liu, and R. Levy. Using Bayesian networks for cyber security analysis. In 2010 IEEE/IFIP International Conference on Dependable Systems Networks (DSN), pages 211–220. IEEE, 2010.
[5] Roesch M. et al. Snort: Lightweight intrusion detection for networks. In LISA, volume 99, pages 229–238, 1999.
[6] Abraham S. and S. Nair. A novel architecture for predictive cybersecurity using non-homogenous

Markov models. Trustcom/BigDataSE/ISPA, 2015 IEEE, volume 1, pages 774–781. IEEE, 2015.
[7] Abraham S. and S. Nair. Predictive cyber-security analytics framework: A non-homogenous Markov model for security quantification. arXiv preprint arXiv , 2015
[8] Abraham S. and S. Nair. Cyber security analytics: a stochastic model for security quantification using absorbing Markov chains. *Journal of Communications*, 9(12):899–907, 2014.
[9] Hink R., J. M. Beaver, M. A. Buckner, T. Morris, U. Adhikari, and S. Pan. Machine learning for power system disturbance and cyber-attack discrimination. Resilient Control Systems (ISRCS), 2014 7th International Symposium on, pages 1-8. IEEE, 2014.
[10] Ostheimer J., S. Chowdury and S. Iqbal. An alliance of human and machines for machine learning: hybrid intelligent systems and their design principles. Technology in Society, ISSN: 0160-791X, 66, 101647-, 2021.
[11] Dua S and X Du. Data mining and machine learning in cybersecurity. CRC press, 2016.
[12] Berghel. Oh, what a tangled web: Russian hacking, fake news, and the 2016 US presidential election. Computer 50(9):87–91, 2010
[13] Haggard and J. R. Lindsay. North Korea and the Sony hack: exporting instability through cyberspace. 2015.
[14] Zrahia A. Threat intelligence sharing between cybersecurity vendors: Network dyadic, and agent views. Journal of Cybersecurity, 4(1):tyy008, 2018.
[15] Jacobs J. and B. Rudis. Data-Driven Security: Analysis, Visualization and Dashboards. John Wiley & Sons, 2014.
[16] Najafabadi M., F. Villanustre, T. M. Khoshgoftaar, N. Seliya, R. Wald, andMuharemagic. Deep learning applications and challenges in big data analytics. Journal of Big Data, 2(1):1, 2015.
[17] Bilge L., D. Balzarotti, W. Robertson, E. Kirda, and C. Kruegel. Disclosure: detecting botnet command and control servers through large-scale netflow analysis.
[18] Patcha A. and J.-M. Park. An overview of anomaly detection techniques: Existing solutions and latest technological trends. Computer networks, 51(12):34, 2018.
[19] Oberheide J., E. Cooke, and F. Jahanian. Rethinking antivirus: Executable analysis in the network cloud. In HotSec, 2007.
[20] Ranka., J. National cyber range. Technical Report, DEFENSE ADVANCED RESEARCH PROJECTS AGENCY ARLINGTON VA STRATEGIC TECHNOLOGY OFFICE (STO), 2011.
[21] Ramachandran M., N. Feamster, and S. Vempala. Filtering spam with behavioral blacklisting. In Proceedings of the 14th ACM conference on Computer and communications security, pages 342–351. ACM, 2007
[22] Raiffa H. Decision Analysis: Introductory lectures on choices under uncertainty. Addison Wesley, Cambridge University Press, 1970.
[23] Berman, Daniel S., et al. "A survey of deep learning methods for cyber security." Information 10.4 (2019): 122.
[24] Prasad, Ramjee, and Vandana Rohokale. "Artificial intelligence and machine learning in cyber security." Cyber Security: The Lifeline of Information and Communication Technology. Springer, Cham, 2020. 231-247.

Acknowledgment

During this research, Lt. Col. Faber was funded by the US Department of Defense, and Professor Paté-Cornell by the Burt and Deedee McMurtry chair in the department of Management Science and Engineering at Stanford.