

© 2015 Key-Whan Chung

GAME THEORY WITH LEARNING FOR CYBER SECURITY
MONITORING

BY

KEY-WHAN CHUNG

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Electrical and Computer Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2015

Urbana, Illinois

Adviser:

Professor Ravishankar K. Iyer

ABSTRACT

Carefully crafted computer worms such as Stuxnet and recent data breaches on retail organizations (e.g., Target, Home Depot) are very sophisticated security attacks on critical cyber infrastructures. Such attacks are referred to as advanced persistent threat (APT), and are on constant rise with severe implications. In all these attacks, the presence of an attacker itself is difficult to detect as they log-in as legitimate users. Hence, these attacks comprising multiple actions are challenging to differentiate from benign and therefore common detection techniques have to deal with high false positive rates. While machine learning and game theoretic models have been applied for intrusion detection, machine learning techniques lack the ability to model the rationality of the players, while the game theoretic approaches rely on the strict assumption of full rationality and complete information. This thesis discusses an approach that proposes Q-Learning to model the decision process of a security administrator which addresses the joint limitations of using game theory and machine learning techniques for this problem. This work compares variations of Q-Learning with a traditional stochastic game model by performing a simulation under different pair of profiles for attackers and defenders using parameters derived from real incident data of a large computer organization. Analysis on the strengths and weaknesses of the algorithms, and how the parameters in the algorithms affect the performance are studied. Simulation results show that Naive Q-Learning, despite the restricted information on the opponent, better reduces the impact of an attacker compared to Minmax Q-Learning against all attackers, or Stochastic Games players against less rational opponents.

To my parents, for their love and support.

ACKNOWLEDGMENTS

I would like to express my deep appreciation to my advisers Professor Ravishankar K. Iyer and Professor Zbigniew T. Kalbarczk, for their support and guidance to achieve perfection. I also appreciate my fellow DEPEND groupmates for all their insightful comments and advice. In addition, I send thanks to my collaborators Dr. Charles A. Kamhoua and Dr. Kevin Kwiat at the Air Force Research Laboratory for the guidance and insights on conducting research on game theory. I would also like to thank my friends in UIUC for their company and encouragement. Lastly, I send my deep appreciation to my family for all their love and support.

TABLE OF CONTENTS

CHAPTER 1	INTRODUCTION	1
CHAPTER 2	MOTIVATION	5
2.1	Study on Security Incidents	5
2.2	Guessing the 2/3 of the Average Game	7
CHAPTER 3	ATTACK MODEL	8
3.1	Attacker	8
3.2	Defender	9
3.3	Attacker-Defender Interaction	10
3.4	Example: NCSA Incidents	10
CHAPTER 4	GAME MODEL	13
4.1	Stochastic Game	13
4.2	Minimax Q-Learning (MMQL)	15
4.3	Naive Q-Learning (NQL)	16
CHAPTER 5	EXPERIMENT	18
CHAPTER 6	EXPERIMENTAL RESULTS	20
6.1	Comparison between Algorithms	20
6.2	Time to Convergence for Different Learning Rates	23
6.3	Change in Optimal Policy	25
CHAPTER 7	RELATED WORK	26
7.1	Intrusion Detection System (IDS)	26
7.2	Learning in Cyber Security	27
7.3	Anomaly Based Intrusion Detection	28
7.4	Game Theory for Cyber Security	31
CHAPTER 8	LIMITATIONS AND APPLICATIONS	33
8.1	Limitations	33
8.2	Applications	33
CHAPTER 9	CONCLUSION	35
REFERENCES		37

CHAPTER 1

INTRODUCTION

Computer systems are tempting targets for attackers. Successful invasion of data and control can threaten the availability of the computer system and compromise the integrity and confidentiality of information processed or stored in the system. To defend systems from exploits, sensors and monitors are deployed at different layers of the system, and system and user activities are logged and audited to filter out suspicious or malicious activities and/or trigger an in-depth investigation or response to a potential attack.

While that type of monitoring and response has been an effective method for detecting hostile activities against systems, recent analysis shows a change in attack trends against cyber systems. For example, according to [1], attacks are not only increasing in number, but are also getting more sophisticated and intelligent, which is accelerating an increase in the number and variety of security measures applied to systems. However, naive deployment of more security monitors and policies does not always lead to better detection. While such increments bring better security coverage, they also increase the complexity of analysis and overhead in terms of performance degradation. In [2], an analysis of security incidents shows that a significant portion of alarms that trigger human investigation turn out to be false positives. In addition, it was shown that most of the incidents were detected only after actual damage was done. Those observations reveal a need for an automated intrusion response that can react to malicious actions threatening a computer system.

In this thesis, we propose a game-theoretic model to emulate the decision-making process in responding to cyber security incidents. Given an attack model

and a reward model based on expert knowledge, our approach determines the optimal action that minimizes damage. We focus on intrusion response; detection of zero-day attacks and unknown attacks are outside the scope of this paper. For modeling a security game, there exist different models. In this study, we use Markov games to model multi-step (multi-state) attacks. The solution of a Markov game as a stochastic game requires full information about the opponent in addition to the player's own information on reward functions [3]. While such modeling fits economic problems well, players in security games (e.g., attackers and system security administrators) often need to learn more about the opponents to supplement shortages in information. System vulnerabilities of attack targets are hidden from attackers, and system security administrators have no information about the attackers in terms of their ability and/or intent. In addition, we find it unrealistic to assume full rationality (an assumption that a player always makes the best decision that maximizes the reward) in attackers with different abilities.

In order to better model learning in security games, we considered variations of Q-Learning where Q-Learning [4] is a model-free reinforcement learning technique, used to learn the optimal policy that maximizes the expected reward (e.g., the monetary value of the information exploited or an estimated loss caused by compromised system availability). We claim that our approach is more realistic than pure game theoretic approaches [3], as it uses an algorithm that releases the restrictions on the rationality of the players or the completeness of information. However, its application is limited to optimizing a Markov decision process that lacks interaction among multiple players [5]. Littman [6] introduces the Minmax Q-Learning algorithm for Markov games. Unlike the traditional Markov game, Minmax Q-Learning does not require full information, and instead reinforces the decision model by supplementing the shortage of information with learning from history. Furthermore, [7] applies Naive Q-Learning for players with no information about the opponent. Q-Learning algorithms are discussed in more detail in Chapter 4.

With respect to learning, earlier work has shown the effectiveness of applying

machine learning for cyber security. However, we find that rationality has been neglected in the existing machine learning approaches. For the models trained on the dataset, the model becomes specific to the observations from history. Hence, such approaches take time to adapt to unforeseen patterns. Q-Learning on the other hand makes decisions based on a model that was derived from human intelligence. Moreover, assuming the possible incompleteness of the model, the approach reinforces the model by adapting to the patterns from the previous iterations.

In this paper, the performance of Q-Learning algorithms (MMQL, NQL) for detecting execution of a multistage attack is evaluated through comparisons of the cumulative earnings of the attacker after multiple iterations over the game. Simulating a security game of attackers and defenders with different abilities and knowledge, we show that Nave Q-Learning has a potential to minimize the loss against non-fully rational attackers.

The main contributions discussed in this thesis are:

- Motivates the approach through a study of real incidents. From an analysis of the Target data breach and earlier work on security incidents, we show the need for automation in incident response.
- Models the battle of an attacker and a defender as a security game. Using real incident data from the Organization X, we derive an attack model that reflects both the attacker's and defender's perspective, and we use the model to formulate a security game. In terms of a security game, this model represents the worst case where the attacker can perform all attacks shown in the dataset.
- Presents an experimental result showing the possibility of applying Nave Q-Learning for effectively learning the opponent's behavior and making a proper decision. Comparing the performance of different decision making algorithms, we present simulation results that show Naive Q-Learning performing better than algorithms with restricted assumptions, especially

against irrational attackers, and show that Naive Q-Learning performs as well as Minmax Q-Learning, despite the relatively limited information.

CHAPTER 2

MOTIVATION

In this chapter, we study a well-known security breach [8] and a set of incidents at the National Center of Supercomputing Applications (NCSA). From this study, we motivate a need for (i) an automated decision process modeling the rational decision making process and (ii) an approach that models the learning aspects of interaction in security games by recalling a well-known incidents showing the limitations of well-known assumptions on the rationality of players and the completeness of information given to each player in a game.

2.1 Study on Security Incidents

Unlike earlier attacks, whose goal was to invade a target and leave as quickly as possible after causing damage, recent attacks show that attackers are willing to remain in the system undetected. Hence, the attack sequences are designed to consist of a set of actions that are hard to differentiate from legitimate ones. These type of attacks are often called Advanced Persistent Threats (APT) [9]. In 2013, around the holiday season, a data breach on the cyber infrastructure of a major retail company was reported [8]. Through this attack, it was estimated that the information of 110 million customers, including personal and financial information, was stolen. Figure 2.1 visualizes the attack timeline based on the senate report [8]. From the initial breach into the system to the removal of malware after confirming the breach on December 15, the attackers resided in the system undetected for more than a month. We notice two interesting aspects of that incident. One is the decision model underneath the attack. The attack consisted of multiple

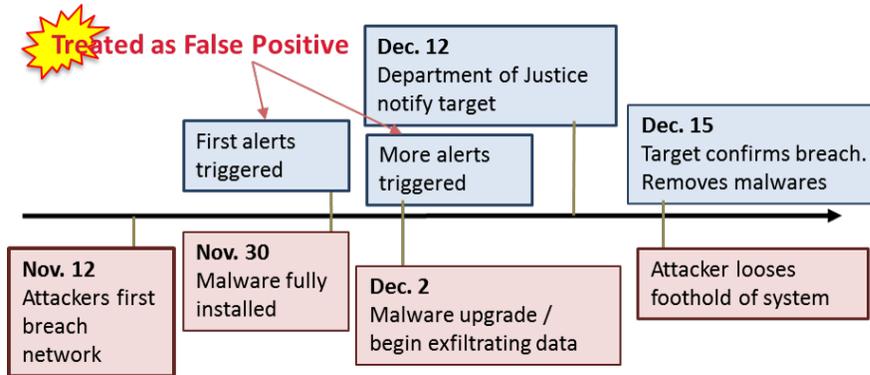


Figure 2.1: Attack and detection timeline of the Target data breach [8].

states (phases), and a decision on the next action to be taken had to be made in each state. From that decision model, we recognized the possibility of applying game-theoretic approaches to counteract malicious intent. Also, we concentrate on the fact that the system was able to detect the intrusion but no proper response was made for easily assuming the alarms as false positives. Often, alarms lack confidentiality as they rely on limited observation. Instead, an attack has to be understood as a sequence of events that calls for the detection/response model to encompass observations from varying dimensions.

In addition, we note the insights from a study of security incidents at a large computing organization [2]. Among the multiple discoveries, we focus on the fact that (i) approximately 50% of the incidents were detected only after the actual damage to the target system has happened, (ii) among roughly 140 alerts per day, a vast majority turn out to be a false positive only after being investigated by the security team. In addition, the description on the detection and response process shows that the process is still heavily dependent on human expertise that uses the monitoring results for making the decision.

From those observations, we see a need for an automated decision-making process to respond to potential attacks. Such a process should provide a decision based not only on the current observation, but also on results from the past and the expected result of taking each action available at the decision-making state. In this thesis, we discuss a method to automatically determine the response, given

the observations on the system states from a set of monitors.

2.2 Guessing the 2/3 of the Average Game

In [10], Nagel performs a guessing game. In the game, each person is asked to select an integer from 0 to 100; the person who picked a number closest to 2/3 of the average would win the game. Under the assumption on complete information and full rationality, the traditional game-theoretic approach suggests 0 would be the winning number. When asked to predict 2/3 of the average, a person would have no chance to win with a value over 66, and if all players are rational enough, the assumption is that 2/3 of 66 will not win the competition. After a number of iterations of recursive thinking, the resulting candidate is reduced to 0. However, the empirical results show that the winning value was not 0 or not close to 0. (The winning value in this study was 22.) From this study, we claim the traditional game theoretic model based on an assumption for full rationality (i.e., players always maximize their reward) and complete information (i.e., each player has complete information about the reward model and the strategy of the opponent) does not fit a game when such assumptions are released. We use this study as a motivation for introducing learning to traditional game theoretic models.

CHAPTER 3

ATTACK MODEL

In modeling an attack, we are considering parties with a conflict of interests: the attacker and the defender. The defender, often a system administrator, manages the system. The main interest of the defender is to secure the cyber infrastructure from malicious activities. The attacker, on the other hand, is a malicious opponent who attempts to compromise the target system. We model the interaction between the attacker and the defender based on the data of actual security incidents.

3.1 Attacker

The attacker is an opponent who accesses the system with the intention of threatening its security. Attacks can vary from a single action to a sequence of activities. In this thesis, we limit our interest to attacks that consist of multiple activities that lead to an ultimate goal.

Attack state AS_x represents the state of the attack, i.e., the depth/degree of intrusion. Each attack state is assigned a numeric value (reward) which quantifies the damage to the target system. The bigger the impact, the more severe the damage to the system and/or the greater the unauthorized control over the system. Transition from one state to another depends on the result of the action.

Activity A is a set of actions a_i available to the attacker. It can lead to malicious control over the system, or if the attacker decides to remain in the current state, the transition will result in a loop. The set of available activities in state AS_x is denoted by A_x . Therefore, A_x is a subset of A . The causal relation between activities and attack states can be represented as a state diagram.

Transition matrix $P_a(s, s')$ is the probability that an action from state s will lead to a transition to the next state s' . In an attack model, a transition matrix represents the probability of a successful attack. Depending on the monitoring system configured on the defender's side, an attack can be either detected or missed. The transition matrix models the uncertainty of the result of an action.

Immediate reward $R_a(s, s')$ is the reward of the attacker as a result of a transition from state s to s' for performing action a . The reward is a quantitative representation of the earnings that the attacker can get from a successful attack.

3.2 Defender

The defender is a party that is in charge of making proper responses to secure the system from malicious attacks. The defender has a set of monitors to protect the system. The main objective of this player is to make proper responses in a preemptive manner based on a limited view of the system status, relying on monitors.

Attack state DS_x represents the state of the attack from the defender's perspective. The observations that defenders use rely on the monitoring systems, and lack the granularity needed to reveal the details of users' actions.

Defender action D is a set of actions (d) available to the defender in a given state. For security incident detection and response, a monitor detects changes in system status. However, such detections do not directly map to the attacker's definite actions. The monitor may miss an action (false negative) or misidentify a benign action as malicious (false positive). Hence, the defender needs to take an appropriate action while relying on imperfect information. Assuming that there are proper responses for each action, we abstract the defender action to either "Response" or "No Response," where "No Response" is useful for monitored events that are hard to differentiate from benign ones, and/or events that do not cause immediate harm to the system.

Table 3.1: Timeline of Sample Attacks of Security Incidents at NCSA

ID	Time	Event
1	09:52	log in from known host using public key authentication
	09:59	changed “authorized_key”
	10:05	logged in from new host / updates “known_hosts”
	11:19	download malware from remote host
	11:20	attempt root escalation
2	17:34	log in through WinVNC exploit
	17:49	download malware from remote host
	17:52	connection to blacklisted command & control systems
3	07:38	access network with weak credentials
	07:39	download malware from remote host
	07:40	start bruteforce ssh scan
	09:15	attempt SSH scan on a DDoS black list

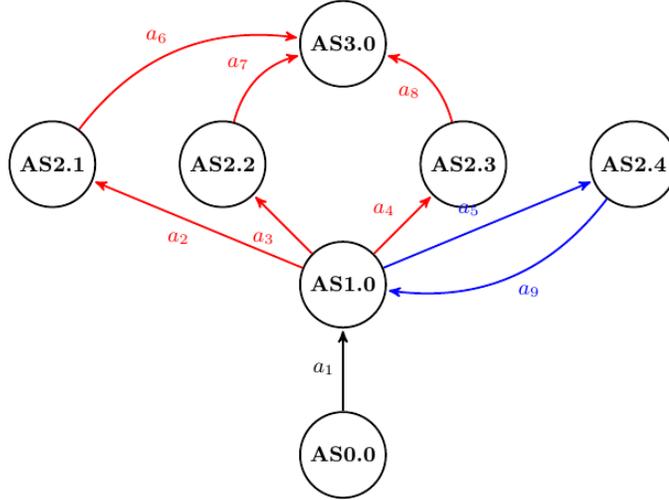
3.3 Attacker-Defender Interaction

While each attacker has a logic flow for making decisions, his or her decisions are not independent, but are related to the opponent’s decision process. Hence, we model the interaction between the two players. We discuss the interaction in more detail in Section 3.4 using the attack model derived from the NCSA incidents.

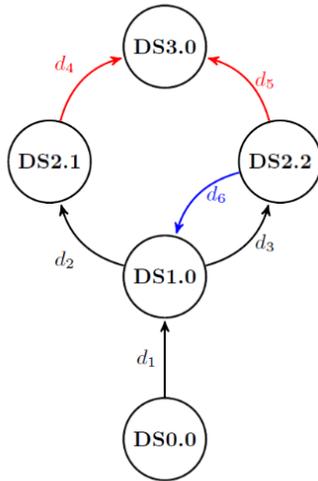
3.4 Example: NCSA Incidents

In [2], it was shown that 62% of the incidents in the study were detected only after attacks have already damaged the system, for the monitoring relies heavily on alerts from the IDS; no significant evidence is available until the actual attack has started. Analyzing the incidents at NCSA, we find a common pattern in the attack phase. For example, as shown in the sample incidents summarized in Table 3.1, suppose a malware file has to be downloaded for an attack to progress. Unless the attacker has downloaded well-known malware whose signature will match the malware detection database, the monitoring system will report the event as a general file download. Since downloading of a file is a general action often performed by benign users, it would be difficult to use a file download as

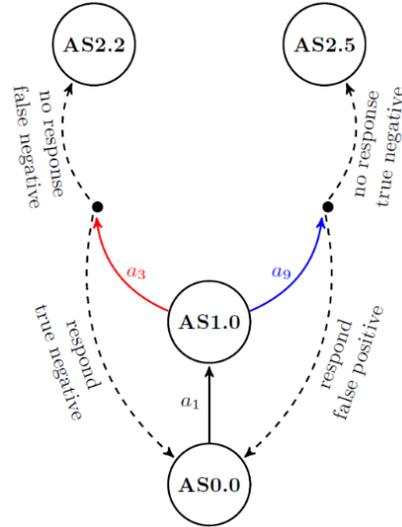
the basis for determining that an attack is occurring. However, once the malware has been executed, the system has already been exploited, with visible damage. Hence, it would be ideal if a useful decision on the response could be made at the download phase.



(a) Attack model from the attacker's perspective.



(b) Attack model from the defender's perspective.



(c) Snapshot of the security game.

Figure 3.1: State diagram representations of the attack phase (a, b) and a snapshot of the security game (c). Note that although (a) and (b) both represent the same attack model, they have different numbers of states and actions, because the defender has a limited view of the attacker's actions.

Figure 3.1 represents the attack model from the attacker’s and defender’s perspectives. Figure 3.1a models an attack in four phases. $AS_{0,0}$ is a base state. It is the default state, in which the attacker has not taken any malicious action. In this state, an attacker cannot be differentiated from a benign user. Using stolen credentials or already exploited backdoors, attackers gain access to the system, which leads to $AS_{1,0}$. In $AS_{1,0}$, the attacker has four options: to download malware (well-known, rare, or new) or download benign software. Downloading benign software would not give immediate benefit to the attacker, but we model it to represent a case in which the attacker is trying to obfuscate the user profile. A successful download leads to $AS_{2,X}$. Once the attacker has reached $AS_{2,1}$, $AS_{2,2}$, or $AS_{2,3}$, he or she can execute the malware and exploit the system. Our goal is to be able to make a proper response before the exploit happens. Figure 3.1b depicts the defender’s view of the same attack model depicted in Figure 3.1a. Because of the limitations of the monitoring system, the defender has a limited view of the files downloaded in $AS_{2,X}$. If the malware has a known signature predefined in the monitoring system, the file can be recognized as malware ($d2$). Otherwise, a malware download will be seen as a benign file download ($d3$).

In Figure 3.1c, we show a subset of the security game. Once an attacker has taken an action, the defender chooses his or her action based on the information from the monitoring system. An attacker’s action results in a transition to the intended state only if the defender does not make a proper response. Once the defender has responded to the observed action, the attacker is forced to transit to the default state. Assuming a zero-sum game, a successful attack will result in an immediate reward, and the defender will have a symmetric loss. As a result of the execution of the attack, the attack state will change accordingly. Otherwise, if the defender detects the attack and makes a proper response, the attack state will be reset to the default for the identified attacker. In that case, a reward will be assigned to the defender, with an equivalent loss to the attacker.

CHAPTER 4

GAME MODEL

In this chapter, the interaction of an attacker and a defender is discussed in terms of games. We refer to [3], [6] and [11] for the definitions and equations for formulating the game. The game consists of two rational players with conflict, and their goal is to maximize their reward by deriving the optimal policy for each state. The comparison of different methods is summarized in Table 4.1.

4.1 Stochastic Game

First we define the terminologies used for solving the game as a Stochastic game [3].

Set of actions A contains all possible actions, a that are available to the player. We use o for the opponent's action.

Reward $R^t(s, a, o)$ defines the immediate reward based on the attack state s and player's action, a and the opponent's action o in the t th iteration.

A stochastic game that consists of multiple stages is often called a Markov game. In a Markov game, the concepts of quality of state and value of state are introduced to represent the expected reward of the player's decision.

Value of state $V(s)$ is the expected reward when the player, starting from state s , follows the optimal policy. It is equivalent to the maximum reward that the player can expect, assuming that the opponent's action o will be the action that minimizes the expected reward. The player maximizes the value of state by deriving the optimal policy, i.e., the probability distribution among the actions

Table 4.1: Comparison on Different Approaches

	MG [3]	QL [4]	MMQL [6]	NQL [11]
Number of Agents	multiple	single	multiple	multiple
Required Opponent Info	full	n/a	limited	no
Learning	no	yes	yes	yes
Adapt to Opponent	no	n/a	partial	partial

available to the player in a given state.

$$V(s) = \max_{\pi} \min_{o' \in O_s} \sum_{a' \in A_s} \pi(s, a') Q(s, a', o') \quad (4.1)$$

Quality of state $Q(s, a, o)$ is the expected reward each player can gain by taking actions a and o from state s and then following the optimal policy from then on. The quality of state is a sum of the immediate reward from this iteration (R^{t-1}) and the reward expected as a result of transitioning to state s' ($V^t(s')$), which was derived from the previous t iterations. Note that the value of state is weighted by a discount factor (γ).

$$Q^{t+1}(s, a, o) = R^{t+1}(s, a, o) + \gamma V^t(s') \quad (4.2)$$

Discount factor γ is assigned by the user's intention on balancing between future and current rewards. A myopic player, who only considers current reward, is modeled by a value of 0 while 1 is assigned for a player who strives for a long-term high reward.

Optimal policy π is the set representing the probability distribution of actions ($\pi(s, \cdot)$) available at each state(s). It is chosen to maximize the value of state ($V(s)$) which represents the expected reward of the player if the player follows the optimal policy. The notation $\pi(s, a)$ indicates the likelihood of taking action a in state s where π is the overall distribution that maximizes the value of state

$(V(s))$.

$$\pi(s, \cdot) = \arg \max_{\pi'(s, \cdot)} \min_{o \in O_s} \sum_{a' \in A_s} \pi(s, a') Q(s, a', o') \quad (4.3)$$

4.2 Minimax Q-Learning (MMQL)

To solve stage-based games, a Markov game assumes full rationality and complete information about the opponent. However, an empirical study involving a guessing game has shown that the assumption on complete information and full rationality is not realistic in all cases [10].

In a security game, the assumption of complete information and rationality is even more unrealistic. In security games, players generally make decisions with limited information, and compensate for their lack of information with learning [12]. To account for that characteristic of security games, we apply Minimax Q-Learning as a decision-making algorithm. Instead for a need of complete information on the attack model, the Minimax Q-Learning algorithm allocates partial weight on its earlier results to combine knowledge of history, the actual earnings on the current iteration, and the future expected reward.

Quality of state for Minimax Q-Learning is defined as follows to embed the learning aspect into the algorithm.

$$Q^{t+1}(s, a, o) = \alpha Q^t(s, a, o) + (1 - \alpha) R^{t+1}(s, a, o) + \gamma V^t(s') \quad (4.4)$$

Learning rate α leverages the ability of the player by assigning a real value between 0 and 1. A learning rate of zero represents full learning ability for the player while a rate of one models the case where the player only considers only the most recent information. In full learning, the player would not consider the immediate reward $R(s, a, o)$ and the expected future award $V(ns)$ but keep the quality of state constant. To account for the absence of prior results to learn from

at the initial stage of the game, an α of 1.0 is assigned; α then decays as $Q(s, a, o)$ accumulates information on the performance of previous iterations [6].

Exploration rate exp is a distinct parameter for Q-Learning which determines the degree of variation from the optimal policy. Unlike the Markov game, in which the optimal solution is known from the initial iteration, Q-Learning has to learn the optimal policy by trial and error. The exploration rate determines the relative rate of the action not following the optimal policy to learn the results of different actions. An exp value closer to 0 results to a Markov game while a value closer to 1 means that the player will take random actions.

4.3 Naive Q-Learning (NQL)

In a security game, information about the opponent is not always available. The attacker often has information about the target system from public resources. However, the amount of information is limited. Similarly, the defender is playing a game against an unspecified opponent. In order to model this situation, Naive Q-Learning from [7] is applied. Naive Q-Learning optimizes the strategy without information about the opponent, such as the opponent's action o . It utilizes limited information of the immediate reward and its own information to derive the optimal policy.

Quality of state is updated accordingly to reflect the limited information. Note that the opponent's action is no longer considered for differentiating the Quality of state.

$$Q^{t+1}(s, a) = \alpha Q^t(s, a) + (1 - \alpha)R^{t+1}(s, a) + \gamma V^t(s') \quad (4.5)$$

Value of state is the maximum expected reward when following the optimal policy. Note that because of the lack of information about the opponent, o is no longer considered.

$$V(s) = \max_{\pi} \sum_{a' \in A_s} \pi(s, a') Q(s, a') \quad (4.6)$$

Optimal policy is the optimal policy that maximizes the value of state ($V(s)$). Note that the quality of state (Q) is only defined for s and a but not o .

$$\pi(s, \cdot) = \arg \max_{\pi'(s, \cdot)} \sum_{a' \in A_s} \pi(s, a') Q(s, a') \quad (4.7)$$

CHAPTER 5

EXPERIMENT

To evaluate and analyze the game, a simulation was performed. The simulation started with initialization of the value and quality of all states, optimal policy and the learning rate (α). For Q-Learning, initially there is no information that the algorithm can learn from. Therefore, α is set to 1.0 indicating that initially, the decision relies on the rationality (game model) only. As shown in Figure 5.1, the Markov game would evaluate the game model and determines the optimal policy before determining what action to take. Once the model has converged to a convergence coefficient (ϵ), the optimal policy is fixed. Q-Learning, on the other hand, needs to take actions before updating the quality and value of state. The next action is decided based on the *exp* value. A random action is chosen with a probability of *exp*; otherwise the action follows the optimal policy. Once the opponent's action has been determined in a similar manner, the quality of state is updated. Then through the use of linear programming, the optimal policy (π) and the value of state ($V(s)$) can be derived. Recall that the optimal policy is the probability distribution of available actions at a given state that maximizes the value of state.

Using the simulator, we compared the performances of the algorithms by assuming a random, Markov, Minimax Q-Learning and Naive Q-Learning attacker and pairing with a Markov, Minimax Q-Learning and Naive Q-Learning defenders. We did not consider the unrealistic situation in which a defender is a random player. For the attacker, on the other hand, a random player is a considerable assumption for modeling unprofessional attackers such as script kiddies. To compare the performance of different algorithms, we evaluated the accumulated immediate

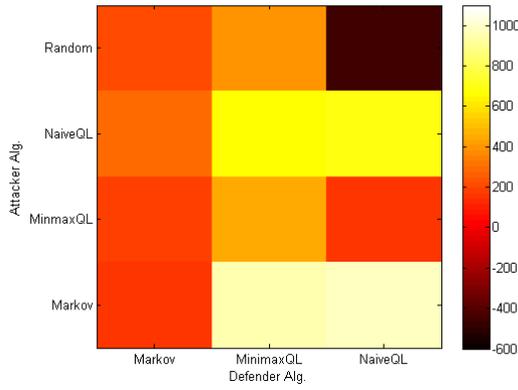
CHAPTER 6

EXPERIMENTAL RESULTS

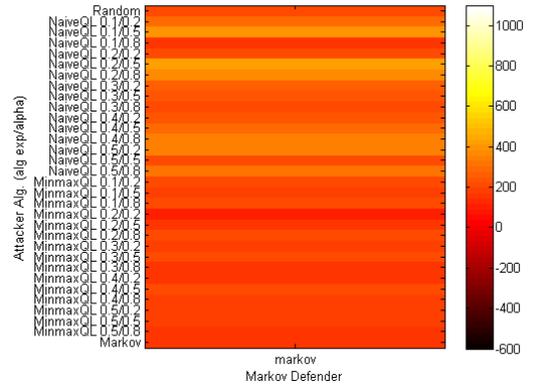
6.1 Comparison between Algorithms

First, we compare how each of the algorithms with varying parameters performs against different opponents. Figure 6.1a provides an overview of the comparison. We represent the accumulated reward of the attacker using a heat map in which a lighter (i.e., closer to white) color indicates higher reward. Looking at the first column, we confirm that the defender, on average, shows the best performance when the decision-making is based on Markov games. In addition, we find low variance within the column. From that observation, we see that when the defender has full information about the attacker and hence is playing a Markov game, the attacker’s choice of algorithm does not make a significant difference. That insight becomes obvious when we consider the assumption upon which the algorithm relies.

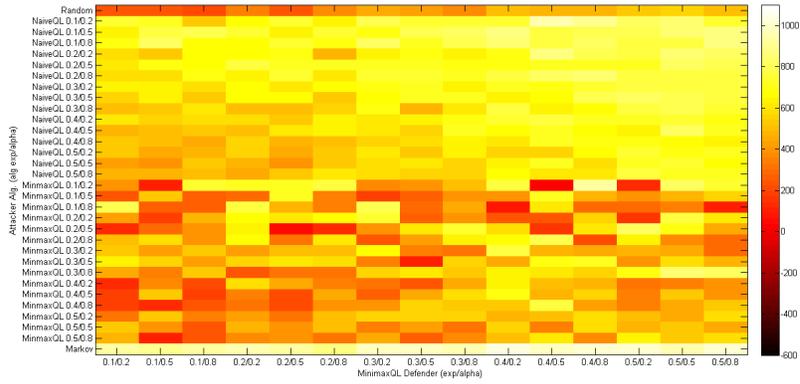
However, the same insight does not always apply to all algorithm pairs. An interesting observation is that Naive Q-Learning performs better than Minmax Q-Learning, despite the limited information. Comparing the second and third columns, we find that attacker performance when played against a Naive Q-Learning defender (third column) is represented by a darker color (lower accumulated reward). In addition, as shown in the upper-right corner of Figure 6.1a, we find that NQL performs better than the Markov game when played against a random attacker. Because the Markov game and Minmax Q-Learning algorithms have more information about the opponent, they are able to formulate a more accurate model.



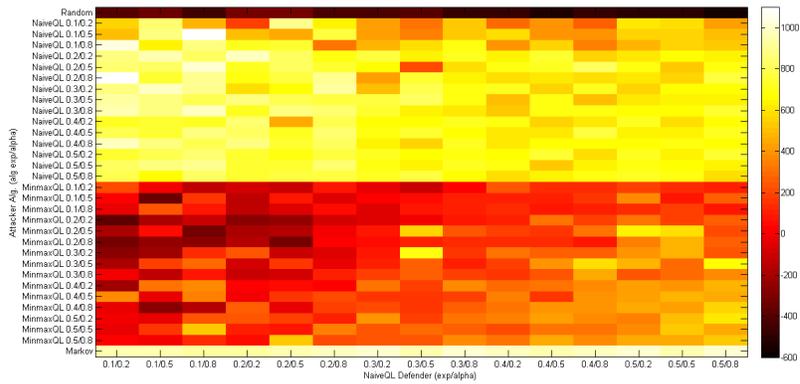
(a) Comparison for different algorithm pairs



(b) Performance of Markov defender



(c) Performance of MMQL defender



(d) Performance of NQL defender

Figure 6.1: Comparison on attacker's accumulated rewards.

On the other hand, as these models assume the rationality of the opponent, they cannot adapt well to the behavior of a random (irrational) attacker. Naive Q-Learning lacks the information about the opponent that would be needed to formulate a complete game model, and hence makes decisions based on its own decision model; that gives it the flexibility to adapt to the attacker’s empirical behavior. By observing the changes of the optimal policy, we confirm that Naive Q-Learning defenders converge to an optimal policy that has higher probability for counteraction against attacks with higher immediate reward.

Figures 6.1b through 6.1d show, in detail, how the attacker’s performance changes for different combinations of algorithms, exploration rates (*exp*), and learning rates (α). In Figure 6.1b, we confirm the insignificant impact of the parameters (*exp* and α) on the performance. While different exploration rates and learning rates are applied, the rows have low variance for the defender’s algorithm driving the game.

From Figures 6.1c and 6.1d, we can see how the parameters affect the performance of the decision-maker. When the attacker is playing the strongest algorithm (MG), the parameters have no impact. That claim can be confirmed by observing the consistent color across columns in the last row of both figures. While the different columns stand for different parameter pairs, no difference in attacker performance was found.

As defined in Section 4.2, Q-Learning introduces a new parameter called the exploration rate (*exp*). This rate defines the ratio of actions that are randomly chosen (rather than being chosen by following the optimal policy). From looking at Figures 6.1c and 6.1d, we find that there is no single trend for *exp*, but rather it depends on the algorithms for both players. When the defender is playing MMQL against an NQL attacker, as shown in the upper-half of Figure 6.1c, a higher exploration rate of the attacker leads to a higher accumulated reward, while a higher *exp* rate of the defender lowers the accumulated reward of the attacker. On the other hand, the lower-half of Figure 6.1d shows that when the defender is deploying NQL against an MMQL attacker, the defender reduces the accumulated

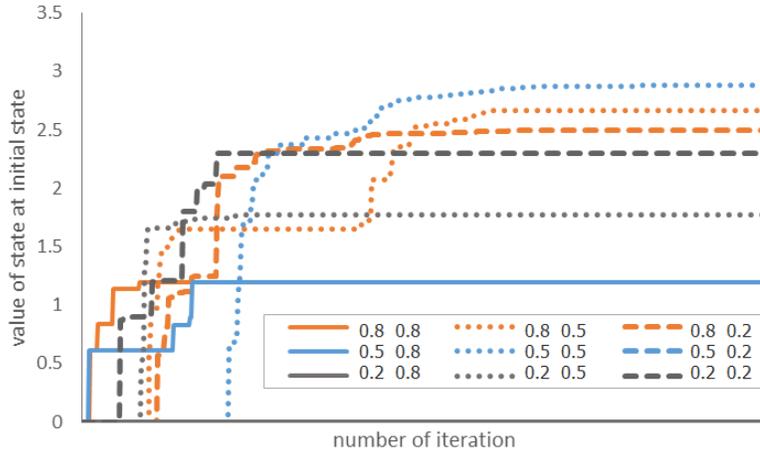


Figure 6.2: Comparison of value of state at initial state. The label indicates the learning rate pair (learning rate of the attacker, learning rate of the defender).

reward of the attacker with a smaller exploration rate, while the attacker gains more with a higher exploration rate. For both cases, we find that the defender only needs a minimal exploration rate to assure discovery of all possible actions. We find that deviating from the defender’s optimal policy does not benefit the defender.

Studying the impact of the learning rate, namely the relative weight between the game model and the learning model, we find no clear pattern in the accumulated reward. Instead, we find a potential relationship to the time to convergence, which we discuss in Section 6.2.

6.2 Time to Convergence for Different Learning Rates

In Figure 6.2, we can see how the learning rate affects the time to convergence. When we assume that the attacker is playing the game under a consistent strategy, the time to convergence indicates the time it takes for the defender to derive the strategy that minimizes the loss. From Figure 6.1b to Figure 6.1d, we saw that the learning rate of the players did not have a significant impact on the accumulated reward of the attacker. To confirm how the learning rate affects the

decision-making process, we compare the values of the state at the initial state of the attack model. That value of state ($V(s_0)$) represents the expected reward that the attacker can earn when the attacker follows the optimal policy from the starting state and thereafter. Once $V(s_0)$ becomes constant, we claim that the optimal policy of the defender becomes constant. Note that this analysis does not assign meanings to the value of state for its nature of representing the expected reward, not the actual reward that has or could be earned.

From the formulation of Minmax Q-Learning algorithms, a zero-sum game is expected to converge [6]. In this experiment, we checked whether the learning rate affects the time to convergence. Recall that a low learning rate indicates intensive learning, as more weight is assigned to the previous quality of state than to the sum of the immediate reward and the future expected reward. From Figure 6.2, we observe two things. One is that the expected reward of the attacker is larger if the defender’s learning rate is larger than or equal to that of the attacker. That insight can be confirmed in the figure through comparison of “0.8 0.5,” “0.2 0.5,” and “0.5 0.5,” and comparison of “0.8 0.8” and “0.2 0.8.” Recall that “0.8 0.5” is interpreted as “Attacker with learning rate 0.8 against a defender with learning rate 0.5”. Based on that observation, we can verify that if both parties apply Minmax Q-Learning for decision-making, then it is more likely for the defender to be able to protect the system against a Minmax Q-Learning attacker when the defender weights learning more than rationality. Another observation from the figure is that the game converges faster for the attackers with lower learning rates. The learning rate of the opponent (the defender, in this case) also affects the time to convergence of the player. While there is a slight deviation between 0.5 and 0.8, a lower learning rate for the defender (opponent) also accelerates the time to convergence of the attacker. A similar analysis for Naive Q-Learning players is not applicable, because the Naive Q-Learning algorithm is not guaranteed to converge.

Table 6.1: Optimal Policy of the Defender

Defender		Defender Optimal Policy					
		na	o1	o2	o3	o4	o5
Makov	AS0		0.4	0.6			
	AS1				0.4	0.6	
	AS2				0.4		0.6
MMQL	AS0		0.64	0.36			
	AS1				0.64	0.36	
	AS2				0.6		0.4
NQL	AS0			1.0			
	AS1					1.0	
	AS2						1.0

6.3 Change in Optimal Policy

Table 6.1 summarizes the optimal policy of different defenders when played against a Markov attacker. Noting that a successful attack in a2, a4 and a5 results to an immediate reward of 2 while 1 on the others, we conclude that Minimax Q-Learning results to a more passive optimal policy while the Minimax Q-Learning defender is protecting the system against aggressive attackers.

From analysis on the final optimal policy, we find that none of the attacker algorithms assign a probability to taking no action. In security games, the actions from attackers do not purely consist of malicious actions. There are intervals between attacks, and some intelligent attackers hide their malicious intentions by executing benign activities. While our approach was intended to model those intentions, the simulation results show that the algorithm considered the cost of a false negative (i.e., missing a malicious action) to be higher than the cost of responding to a false positive. Such a decision can change based on the reward model defined in the game model.

CHAPTER 7

RELATED WORK

7.1 Intrusion Detection System (IDS)

An Intrusion Detection System (IDS) is a hardware or software system that monitors cyber infrastructure to detect suspicious activity to alert system/network administrators [13]. Based on what the IDS concentrates on, it can be classified as either a Network based Intrusion Detection System (NIDS) or a Host based Intrusion Detection System (HIDS). NIDS scans the traffic that enters or leaves a network and tries to detect any malicious action. The detector can identify the host operating systems, applications from the packets, and detect attacks from the application to network layer. HIDS, on the other hand, monitors individual hosts. Depending on the type of monitors that are deployed on each host, the IDS can monitor network traffic for that host, system logs, running processes, file access/modifications etc. and the detector determines malicious actions from such data collected from different hosts. The type of attacks that HIDS can detect includes malicious code execution that can permit unauthorized access or escalate privileges, and can also detect buffer overflow by monitoring certain sequences of instructions or memory access points. Moreover HIDS type systems can also perform network traffic analysis but are limited to the traffic to and from the hosts that are being monitored. An HIDS can also perform files system monitoring. However, as Sharma et al. [2] claim, the use of such monitoring is limited for the overhead and performance degradation. Signature based IDS is the traditional IDS that detects attacks from known signatures. It seeks to find known attacks by looking for specific patterns or signatures within available information. For

instance, an IDS system checks the hash value of files being transmitted into the system and sees if the value matches any of those in the malware hash registry. Also, it can check the file type and can generate alerts if any executable file is downloaded. However, for the need of a reference to compare to, such detection is limited to specified, well-known attacks but is not capable of new, unfamiliar or obfuscated attacks [14]. With the varying attacks, each of which is getting sophisticated, the traditional IDS faces a limitation on defining all the possible signatures which asks for an alternative approach.

7.2 Learning in Cyber Security

With the computer systems getting complicated and customized, attackers can no longer rely on common knowledge for performing an attack. Instead, to determine the system status and vulnerabilities, attackers probe the system and learn about the system from side channels. Bethecourt [12] presents a probe response attack for locating network sensors. By probing an IP address that would lead to reporting the monitoring results to the Internet Storm Center (ISC), and later checking the reports the paper claims that the attacker can identify the identity of the sensors within less than a week. In [15], Shmatikov and Wang show how the attackers gain knowledge about the detection by generating detectable attacks and monitoring how a collaborative intrusion detection system generates an alert. Jana and Shmatikov [16], on the other hand, demonstrates a side channel attack where the authors show the possibility of exploiting the memory footprints to exploit the secret of applications through inference.

In addition, a set of recent attacks motivates a need for learning on the detector side. One example is the Target data breach discussed in Section 2.1. From this attack, we have learned that attackers are designing sophisticated attacks for a specific target. For such attacks that are hard to differentiate from benign activities, traditional IDS approaches face limitations. Another example is the Stuxnet attack. Stuxnet is a worm discovered in June 2010 that was targeting

Table 7.1: Stuxnet’s novel characteristics [17]

Aspect	Stuxnet	Common Malware
Targeting	Extremely selective	Indiscriminate
Type of target	Industrial control systems	Computers
Size	500 KBytes	Less than 1 Mbyte
Probable initial infection vector	Removable flash drive	Internet and other networks
Exploits	Four zero-days	Possibly one zero-day

the control system of Iran’s nuclear facilities. It is known to spread via MS Windows, and target Siemens industrial control systems, and is the first discovered malware that intruded industrial systems. As shown in Table 7.1, Stuxnet differs from ordinary malware. While malwares from the past try to infect as many computers as possible, Stuxnet targets a specific system. Also while common malware has exploited possibly a single zero-day vulnerability, Stuxnet is known to have exploited four zero-day attacks. For such characteristics of Stuxnet, Chen et al. [17] claim that the attack has become more sophisticated. It targeting a specific system requires remarkable knowledge of the system, and the creators would have needed to know the target configuration. Also from an observation on the number of exploits included in this attack, we can easily determine the amount of effort involved in designing this attack. For such workload put into an attack, it is getting difficult to detect such an attack from traditional approaches and this trend requires a change in detection methods.

7.3 Anomaly Based Intrusion Detection

To overcome the limitation of the signature based IDS approach on detecting attacks that are growing in size and variation, anomaly based IDS has become a hot topic in the research area of computer security. Before we discuss previous work utilizing computational intelligence, we briefly introduce common algorithms applied to security problems. We summarize the models based on the discussions in [18] and [19].

Figure 7.1 [19] compares the four common algorithms, Naive Bayes (NB), Hid-

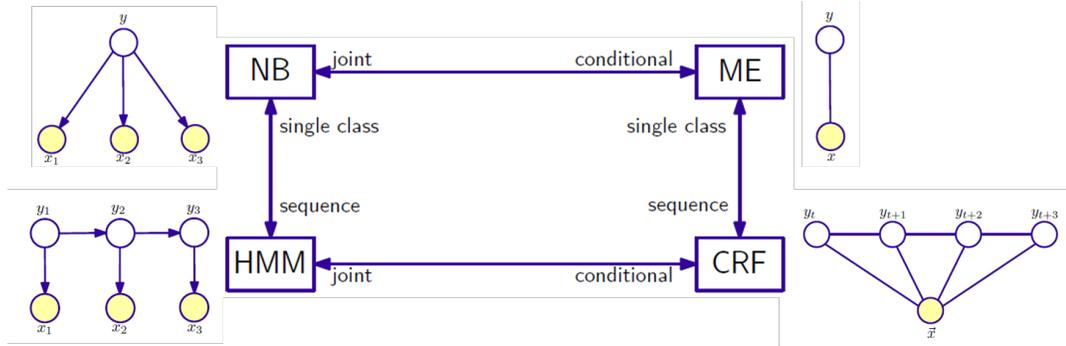


Figure 7.1: Comparison of common machine learning algorithms

den Markov models (HMM), Maximum Entropy model (ME), and Conditional Random Fields (CRF). The naive Bayes model classifies a variable in dependence of several feature values by modeling the joint probability. Comparing the conditional probability among available outputs, the classifier labels the output with the highest joint probability.

The maximum entropy model, on the other hand, models the conditional probability. Based on the principle of maximum entropy¹, the model aims to find the output label that maximizes the conditional entropy.

While the naive Bayes model and the maximum entropy model might be effective for modeling single class variables, such an approach has limitations on modeling the multistate attacks discussed in Chapter 2. The hidden Markov model expands the naive Bayesian model to predict a sequence of output variables. Unlike the naive Bayes model that depended only on the observation variables, HMM brings the output variable from the previous state in computing the joint probability. However, this model is only applicable for a certain number of cases for assuming the conditional independence between observation variables.

Conditional random fields can be understood as a sequence version of the maximum entropy model. Similar to the extension from a naive Bayes model to a hidden Markov model, CRF models the conditional probability of the sequence

¹If incomplete information about a probability distribution is available, the only unbiased assumption that can be made is a distribution which is as uniform as possible given the available information [19].

of states, given an observation vector and outputs a sequence of labels that maximizes the conditional entropy.

Amor et al. [20] proposed a naive Bayesian learning model for intrusion detection and compared the performance with a model that utilizes a decision tree. Like many other similar work, it uses data from KDD '99 [21] which contains network intrusion logs ranging from Denial of Service (DoS), user to root, remote to user and probing attacks. In this model a subset of the 41 features in the dataset is given as an input and the model is trained. Then for the test or unlabeled data, the attack class is chosen to maximize the joint probability. Experimental results show that the accuracy of the naive Bayes model turns out to be similar to a decision tree model despite the weak assumption of the NB model of the independence assumption.

While Amor has shown that the naive Bayes model performs well despite the weak assumption it is based on, there are other works that implement anomaly based intrusion detection. Ourston et al. [22] model a detector using a hidden Markov model. HMM is an expanded probabilistic model of the NB, that eliminates the independence assumption and models a sequence of output states. The work is based on the fact that a network attack can be modeled into different phase of attacks. Here it is divided into probing, consolidating, exploiting and then compromised state.

Gupta et al. [23], on the other hand, present a different model for anomaly based intrusion detection. It not only utilizes a different algorithm but also applies a layered approach for better accuracy. In this model, the authors train a conditional random field, which models the conditional probability for a sequence of output classes given a sequence of input sequences. For a layered approach, the paper borrows the concept from the airport security model where a number of checks is performed in a sequential manner. The detector is layered on the attack class and each layer is trained with relevant features for each output class. Then for classification, each event is evaluated for the first attack class and any that does not pass the filter are considered malicious. Only the events that pass

through all four filters are labeled normal. This model compares to HMMs for less computational overhead. Gupta et al. show that while use of conditional random fields itself might not have improved the performance in a significant manner, layering the detector takes the main part of boosting the detection performance.

7.4 Game Theory for Cyber Security

Modeling attacker intent or attack flow in a graphical model has been a well-studied problem in security. For the applicability and variation of game theoretic models, numerous approaches exist for modeling security as a game. Bier [24] provides a good study of a defender securing a set of potential targets with limited resource. Bier solves the problem in a resource constrained environment to answer policy questions to better secure the physical target against threat of terror. Though this paper solves a physical security problem, it provides a good framework for cyber security games. Liu et al. present a game theoretic model to infer attacker intent, objectives, and strategies (AIOS)[25]. Considering the incomplete knowledge of each players on the opponents, the authors choose a Bayesian game model. The model also introduces the state of the attack. The attack state is normally predicted from observable events.

A number of previous works apply traditional game theoretic approach for cyber security problems. In [26] and [27], the authors apply a Markov decision process (MDP) to secure information sharing in online social networks. In addition, [28] and [29] apply a Markov game framework for optimal data management in online social network. Nguyen et al. [30], apply fictitious play [31] for solving a security game with incomplete information. Similar to [7], information on the opponent's payoff matrix is not available, hence the agent derives the belief of its prediction on the opponent's strategy by monitoring the result of its own move. However, such an approach has limitations on applying to a Markov game consisting of numerous states. Alpcan and Basar present a comprehensive study on modeling security games under different level of information about the opponent

Table 7.2: Summary of Game Theoretic Approaches for Security Games

	Information	Repeated Game	Attack Type
Lye 2002 [32]	complete	stochastic	battle
Liu 2006 [33]	incomplete	no (static Bayesian)	single
Alpcan 2006 [11]	complete/incomplete	stochastic	single
Zonouz2009 [34]	complete	yes	single
Becker 2011 [35]	complete	no	single
Markov	complete	stochastic	battle
MMQL	incomplete	stochastic + learning	battle
NQL	incomplete	stochastic + learning	battle

[11]. In their model, the interactions of two players are modeled as a stochastic (Markov) game. Each player has an option of attack/no attack or respond/not respond. They present a simulation model under three different conditions: perfect information about the system (Q learning), partial (action set, transaction history) information about the opponent (Minimax Q-Learning) and no information about the opponent (naive Q learning). In Table 7.2, we compare the works more relevant to the approach discussed in this thesis.

CHAPTER 8

LIMITATIONS AND APPLICATIONS

8.1 Limitations

Unlike many approaches using machine learning [36], [37], our approach is not intended to detect new attacks. Instead, our game theoretic approach, like other decision-making applications, suggests a likelihood of taking a certain action to maximize the benefit of the security administrator.

Also, while our approach is based on the attack and reward model, because of the lack of agreement on security metrics, there is no true measure for quantifying the rewards of a successful attack or attack detection. Therefore, the current configuration relies on expert knowledge to enumerate the potential damage or overhead for taking a certain action.

One last limitation is that the attack model's performance depends on expert knowledge. Because the decision-making process is based on the attack model, the granularity and completeness of the attack model affect the performance. That limitation is intrinsic to pure game-theoretic models, and we claim that our approach can compensate for that shortcoming of the attack model by adding the ability to learn from past iterations. However, lack of coverage of undefined actions or states still remains as a limitation.

8.2 Applications

Utilizing Q-Learning for optimizing decisions, the algorithm returns the optimal policy (i.e., the probability distribution). The optimal policy represents the prob-

ability distribution of the possible actions that would maximize the benefit of the defender. An action with high probability indicates either the high potential earning and/or the high possibility of occurrence according to the previous sequence of actions. The optimal policy can be used to prioritize the alarms from the monitoring results and/or to implement preemptive defense methods. The reward model and its optimal policy can be used to assign security scores for each alarm generated in the monitoring system. Such score can prioritize the alarms upon severity and probability of true positives. Furthermore, if the prediction from learning and rationality becomes precise enough, an automatic response method can be applied from the decision using this approach. While the discussion in this thesis was focused on automated response upon intrusion detection, we plan to test and evaluate our work by embedding it in a security analytics testbed [38]. With real incident data fed into the testbed and factor graph [39] detecting malicious intentions; our approach will then determine the right response to take. By combining our game theory based response model with the detection framework, we expect to verify the timeliness of intrusion detection and response and determine the accuracy and impact of misdetection.

CHAPTER 9

CONCLUSION

In this thesis, we presented our approach for modeling the decision-making process of cyber security monitoring using a game-theoretic approach. To reflect the realistic conditions of decision-making in a security game, we considered variations of Q-Learning algorithms. Minmax and Naive Q-Learning, compared to traditional Markov games, are more realistic when applied to security games, because they relax the requirement for full information about the opponent. We compensated for the lack of information by enabling learning of the optimal policy, which has the advantage that it resembles situations in which attackers probe system vulnerabilities (through techniques like scanning) and defenders train and renew security policies and devices based on earlier data. We noted that the rich literature in online learning theory lacks efforts to reason about the pattern to capture the rationality of attackers in security games.

From the experiments based on simulation, it was shown that Naive Q-Learning performs well against irrational (non-Markov) attackers, i.e., random decision-makers or attackers based on probing and learning. When played against a Markov attacker, the Naive Q-Learning approach was able to perform at least as well as a Minmax Q-Learning defender. In the real space of security games, players, especially the defenders, have limited ability to obtain information about their opponents. Any parties with access to the system are potential attackers, and their ability and knowledge not only vary but are hidden. Hence, a Markovian attacker, which represents the worst case for the defender, is unrealistic. The simulation results show that despite the limited information on which decisions are based, our approach is promising compared to the traditional Markov game

approach and Minmax Q-Learning.

While we present the possibility of Naive Q-Learning as a decision-making logic in security games, some limitations remain. For the lack of agreement in metrics that represent impacts, there is no clear definition of the reward model. In the simulation, the reward was abstracted as the relative severity under an assumption of a zero-sum game, indicating that a player's reward is his or her opponent's loss. In addition, for the dependency of the parameters to the opponent, it is necessary to do further study on how to tune the parameters of a Naive Q-Learning algorithm against an attacker from real data, and to test the approach embedded in a framework with real-time logs and specific detection logic.

REFERENCES

- [1] S. J. Templeton and K. Levitt, “A requires/provides model for computer attacks,” in *Proceedings of the 2000 Workshop on New Security Paradigms*. ACM, 2001, pp. 31–38.
- [2] A. Sharma, Z. Kalbarczyk, J. Barlow, and R. Iyer, “Analysis of security data from a large computing organization,” in *Dependable Systems & Networks (DSN), 2011 IEEE/IFIP 41st International Conference on*. IEEE, 2011, pp. 506–517.
- [3] L. S. Shapley, “Stochastic games,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 39, no. 10, p. 1095, 1953.
- [4] C. J. Watkins and P. Dayan, “Q-learning,” *Machine Learning*, vol. 8, no. 3-4, pp. 279–292, 1992.
- [5] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2009, vol. 414.
- [6] M. L. Littman, “Markov games as a framework for multi-agent reinforcement learning.” in *ICML*, vol. 94, 1994, pp. 157–163.
- [7] T. Alpcan and T. Basar, “A game theoretic approach to decision and analysis in network intrusion detection,” in *Decision and Control, 2003. Proceedings. 42nd IEEE Conference on*, vol. 3. IEEE, 2003, pp. 2595–2600.
- [8] U.S. Senate, Committee on Commerce, Science and Transportation, “A ‘kill chain’ analysis of the 2013 target data breach,” Tech. Rep., March 2014.
- [9] S. Bodmer, M. Kilger, G. Carpenter, and J. Jones, *Reverse Deception: Organized Cyber Threat Counter-Exploitation*. McGraw Hill Professional, 2012.
- [10] R. Nagel, “Unraveling in guessing games: An experimental study,” *The American Economic Review*, pp. 1313–1326, 1995.
- [11] T. Alpcan and T. Basar, “An intrusion detection game with limited observations,” in *12th Int. Symp. on Dynamic Games and Applications, Sophia Antipolis, France*, 2006.

- [12] J. Bethencourt, J. Franklin, and M. Vernon, “Mapping internet sensors with probe response attacks,” in *USENIX Security*, 2005.
- [13] J. P. Anderson, “Computer security threat monitoring and surveillance,” Technical report, James P. Anderson Company, Fort Washington, Pennsylvania, Tech. Rep., 1980.
- [14] P. Garcia-Teodoro, J. Diaz-Verdejo, G. Maciá-Fernández, and E. Vázquez, “Anomaly-based network intrusion detection: Techniques, systems and challenges,” *Computers & Security*, vol. 28, no. 1, pp. 18–28, 2009.
- [15] V. Shmatikov and M.-H. Wang, “Security against probe-response attacks in collaborative intrusion detection,” in *Proceedings of the 2007 Workshop on Large Scale Attack Defense*. ACM, 2007, pp. 129–136.
- [16] S. Jana and V. Shmatikov, “Memento: Learning secrets from process footprints,” in *Security and Privacy (SP), 2012 IEEE Symposium on*. IEEE, 2012, pp. 143–157.
- [17] T. M. Chen and S. Abu-Nimeh, “Lessons from Stuxnet,” *Computer*, vol. 44, no. 4, pp. 91–93, 2011.
- [18] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.
- [19] R. Klinger and K. Tomanek, *Classical Probabilistic Models and Conditional Random Fields*. TU, Algorithm Engineering, 2007.
- [20] N. B. Amor, S. Benferhat, and Z. Elouedi, “Naive bayes vs decision trees in intrusion detection systems,” in *Proceedings of the 2004 ACM symposium on Applied Computing*. ACM, 2004, pp. 420–424.
- [21] M. Tavallae, E. Bagheri, W. Lu, and A.-A. Ghorbani, “A detailed analysis of the KDD cup 99 data set,” in *Proceedings of the Second IEEE Symposium on Computational Intelligence for Security and Defence Applications 2009*, 2009.
- [22] D. Ourston, S. Matzner, W. Stump, and B. Hopkins, “Applications of hidden Markov models to detecting multi-stage network attacks,” in *System Sciences, 2003. Proceedings of the 36th Annual Hawaii International Conference on*. IEEE, 2003, pp. 10–19.
- [23] K. K. Gupta, B. Nath, and R. Kotagiri, “Layered approach using conditional random fields for intrusion detection,” *Dependable and Secure Computing, IEEE Transactions on*, vol. 7, no. 1, pp. 35–49, 2010.
- [24] V. Bier, S. Oliveros, and L. Samuelson, “Choosing what to protect: Strategic defensive allocation against an unknown attacker,” *Journal of Public Economic Theory*, vol. 9, no. 4, pp. 563–587, 2007.

- [25] P. Liu, W. Zang, and M. Yu, “Incentive-based modeling and inference of attacker intent, objectives, and strategies,” *ACM Transactions on Information and System Security (TISSEC)*, vol. 8, no. 1, pp. 78–118, 2005.
- [26] J. White, J. S. Park, C. A. Kamhoua, and K. A. Kwiat, “Social network attack simulation with honeytokens,” *Social Network Analysis and Mining*, vol. 4, no. 1, pp. 1–14, 2014.
- [27] J. White, J. Park, C. Kamhoua, and K. Kwiat, “Game theoretic attack analysis in online social network (OSN) services,” in *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. ACM, 2013, pp. 1012–1019.
- [28] C. A. Kamhoua, K. A. Kwiat, and J. S. Park, “A game theoretic approach for modeling optimal data sharing on online social networks,” in *CCE*, 2012, pp. 1–6.
- [29] J. S. Park, K. A. Kwiat, C. A. Kamhoua, J. White, and S. Kim, “Trusted online social network (OSN) services with optimal data management,” *Computers & Security*, vol. 42, pp. 116–136, 2014.
- [30] K. C. Nguyen, T. Alpcan, and T. Basar, “Security games with incomplete information,” in *Communications, 2009. ICC’09. IEEE International Conference on*. IEEE, 2009, pp. 1–6.
- [31] D. Fudenberg and D. Levine, “Learning in games,” *European Economic Review*, vol. 42, no. 3, pp. 631–639, 1998.
- [32] K.-W. Lye and J. M. Wing, “Game strategies in network security,” in *Foundations of Computer Security*, 2002, p. 13.
- [33] Y. Liu, C. Comaniciu, and H. Man, “A Bayesian game approach for intrusion detection in wireless ad hoc networks,” in *Proceeding from the 2006 Workshop on Game Theory for Communications and Networks*. ACM, 2006, p. 4.
- [34] S. A. Zonouz, H. Khurana, W. H. Sanders, and T. M. Yardley, “RRE: A game-theoretic intrusion response and recovery engine,” in *Dependable Systems & Networks, 2009. DSN’09. IEEE/IFIP International Conference on*. IEEE, 2009, pp. 439–448.
- [35] S. Becker, J. Seibert, D. Zage, C. Nita-Rotaru, and R. State, “Applying game theory to analyze attacks and defenses in virtual coordinate systems,” in *Dependable Systems & Networks (DSN), 2011 IEEE/IFIP 41st International Conference on*. IEEE, 2011, pp. 133–144.
- [36] R. A. Maxion and T. N. Townsend, “Masquerade detection augmented with error analysis,” *Reliability, IEEE Transactions on*, vol. 53, no. 1, pp. 124–147, 2004.

- [37] F. Bergadano, D. Gunetti, and C. Picardi, “User authentication through keystroke dynamics,” *ACM Transactions on Information and System Security (TISSEC)*, vol. 5, no. 4, pp. 367–397, 2002.
- [38] P. Cao, E. C. Badger, Z. T. Kalbarczyk, R. K. Iyer, A. Withers, and A. J. Slagell, “Towards an unified security testbed and security analytics framework,” in *Proceedings of the 2015 Symposium and Bootcamp on the Science of Security*. ACM, 2015, p. 24.
- [39] P. Cao, E. Badger, Z. Kalbarczyk, R. Iyer, and A. Slagell, “Preemptive intrusion detection: Theoretical framework and real-world measurements,” in *Proceedings of the 2015 Symposium and Bootcamp on the Science of Security*. ACM, 2015, p. 5.