Harvard CS 182: Artificial Intelligence October 12, 2022

Lecture 11: Stackelberg Security Games

Lecturer: Ariel Procaccia *Author: Lauren Cooke*

1 Introduction

We can apply game theory to all sorts of real-world problems as we attempt to write algorithms to solve them. In this case, we apply the Stackelberg game structure to problems where attackers want to attack a target, and decide upon a strategy given observations of how defenders protect their targets. Our goal is to then devise an optimal strategy for defending our targets from attackers who can observe the defense strategy. We call this particular application of game theory a **Stackelberg Security Game**.

2 Stackelberg Games

A **Stackelberg game** is a game between two players, a leader and a follower, that occurs sequentially. The leader will go first and commit to a strategy, while the follower observes the leader's committed strategy before choosing a strategy to commit to.

Example 1 *Consider the game grid in Table 1 where the row player can choose to play the top or bottom row, and the column player can choose to play the left or right column. The first value of the ordered pair represents the outcome for the row player and the second value of the ordered pair represents the outcome for the column player:*

$(1,1)$	$(3,0)$
$(0,0)$	$(2,1)$

Table 1: Example Game

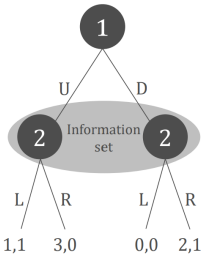

Figure 1: Stackelberg extensive form

In the scenario where both players are trying to play optimally for themselves, our row player will always choose the top row. Knowing that the row player will be optimal and choose the top row, our column player will choose the best outcome for himself out of the two choices in the top row, meaning that this game will be locked at $(1,1)$ *as a* ***Nash Equilibrium****.*

What's unique about a Stackelberg game is that the leader is not locked into a an individually dominant strategy. All that a Stackelberg game will guarantee is that ***the follower will choose a best response strategy to the strategy that the leader is using.*** *Say that our row player is the leader. In this example, the leader can commit to playing the bottom row, forcing our follower to choose the right column for an optimal outcome. This ability*

to commit allowed our leader to achieve a better outcome at (2,1). The leader announcing their strategy commitment before the follower chooses a strategy creates an extensive form game (Figure 2). Note that ***the follower will only be able to observe the strategy of the leader****, meaning that should the leader choose a mixed strategy where they choose each row with a certain probability, the follower won't know for certain which row that the leader has chosen.*

Using a mixed strategy where the leader commits to each row with an equal probability of 0.5, we can calculate our leader's expected utility. Given that we break ties in favor of the leader and our follower has an equal probability of reaching their maximum utility from choosing either column, our follower will commit to the right column. Given this commitment, we can find the total expected utility for the leader by summing up our utility options multiplied by the probability that each option occurs:

$$(.5)(3) + (.5)(2) = 1.5 + 1 = 2.5$$

3 Computing Stackelberg Equilibrium

3.1 Definitions

Let the leader of our stackelberg game play with a mixed strategy called x_1. We can define the set of **best response strategies** that the follower can use as $B_2(x_1)$.

$$B_2(x_1) = \text{argmax}_{S_2 \in S} u_2(x_1, s_2)$$

Where s_2 represents the selected follower strategy, S represents the set of strategies, S_2 represents the set of pure strategies that player 2 can play, and $u_2(x_1, s_2)$ represents the utility for the follower given that the leader uses strategy x_1 and the follower uses the strategy s_2. In other words, the set of best response strategies given the leader's input strategy will be the set of pure strategies that player 2 can play that maximize the utility for the follower. It is important to note that the **best response strategy is always a pure strategy**.

In a **strong Stackelberg equilibrium**, the leader plays a mixed strategy to maximize their utility given that the follower will choose the best response strategy. Formally:

$$\text{argmax}_{x_1 \in \Delta(S)} \text{max}_{s_2 \in B_2(x_1)} u_1(x_1, s_2)$$

Where $\Delta(S)$ is the set of mixed strategies and $u_1(x_1, s_2)$ represents the utility for the leader given that the leader uses strategy x_1 and the follower uses the best response strategy s_2. Note that given this definition, the follower is assumed to choose the best response strategy that is best for the leader.

3.2 Computation

In two-player normal form games, we can compute a Strong Stackelberg Equilibrium in polynomial time using a linear program:

$$\max \sum_{s_1 \in S} x(s_1) u_1(s_1, s_2^*)$$

$$\text{s.t } \forall s_2 \in S, \sum_{s_1 \in S} x(s_1) u_2(s_1, s_2^*) \geq \sum_{s_1 \in S} x(s_1) u_2(s_1, s_2)$$

$$\sum_{s_1 \in S} x(s_1) = 1$$

In other words, we are trying to maximize the expected utility for the leader given that the leader plays some mixed strategy $x(s_1)$ that represents the probability that each possible pure leader strategy $s_1 \in S$ is played with and our follower chooses a best response strategy s_2^*. Note that our variables are the $x(s_1)$s, meaning that our goal is to find the probability assignment for the mixed strategy for the leader that will maximize the leader's expected utility. We want to solve this linear program for every possible follower strategy $s_2 \in S$, so that we can choose the leader strategy that gives us the highest utility among these different follower options. We maximize our leader utility subject to the following constraints:

Our first constraint ensures that the selected follower strategy is optimal given the leader strategy because for all follower strategies $s_2 \in S$, we enforce that the expected utility for the follower when playing strategy s_2^*, $\sum_{s_1 \in S} x(s_1) u_2(s_1, s_2^*)$, is greater than or equal to the expected utility for the follower when using any other strategy s_2.

Our second constraint ensures that our mixed strategy is valid, meaning that the sum of all probabilities that each leader strategy is played as defined in $x(s_1)$ is 1.

4 Security Games

4.1 Model

In a security game, there are a set of **targets** $T = \{1, \ldots, n\}$, a set of m security **resources** that we will call Ω, and a set of **schedules** $\Sigma \subseteq 2^T$, or subsets of targets. We are also given $A(\omega)$, or the set of schedules that ω can be assigned to for every resource $\omega \in \Omega$. Here, in assigning a resource to a schedule, defense protects every target in that schedule. Note that given how resources are allocated to schedule, we get a set of **coverage probabilities** for all targets that we can call $c = (c_1, \ldots, c_n)$. Given this setup, the attacker then gets to chose one target to attack.

Per target $t \in T$, we have four measurable values of utility:

1. $u_d^+(t)$ the defender's utility if defense protects the target t and t is attacked
2. $u_d^-(t)$ the defender's utility if defense does not protect the target t and t is attacked
3. $u_a^+(t)$ the attacker's utility if defense protects the target t and t is attacked
4. $u_a^-(t)$ the attacker's utility if defense does not protect the target t and t is attacked

Intuitively, we know that $u_d^+(t) \geq u_d^-(t)$ and that $u_a^+(t) \leq u_a^-(t)$ because an attack will have a greater impact if there is no defense.

With these terms, we can then define the expected utility to the defender $u_d(t, c)$ and the expected utility to the attacker $u_a(t, c)$ if target t is attacked under coverage probability setup c:

$$u_d(t, c) = (c_t)(u_d^+(t)) + (1 - c_t)(u_d^-(t))$$
$$u_a(t, c) = (c_t)(u_a^+(t)) + (1 - c_t)(u_a^-(t))$$

Example 2 *Consider the following Stackelberg Security Game:*

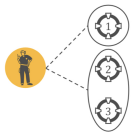

Figure 2: schedules

	Attacker	Attacker	Defender	Defender
Target	Covered	Uncovered	Covered	Uncovered
1	0	4	0	-2
2	0	3	0	-1
3	0	2	0	-5

Table 2: Example Stackelberg Security Game

In this case, we can calculate the utility for the defender for the target that is attacked. First, we find the probability of assigning a single resource to target 1:

To do this, we want to find the probability of assigning this resource to each target by ensuring that the utility of the attacker will be equal for each target that they could choose. Note that in this particular case, the attacker will never choose to attack target 3, as it is covered with the same probability as target 2 given the schedules and has a strictly lower utility than target 3. Therefore, we only need to ensure that the utility for the attacker when they attack target 1 is equal to the utility for the attacker when they attack target 2. We use this setup to solve for $\mathbb{P}(d_1)$*:*

$$u_a(1, c_1) = u_a(2, c_2)$$
$$\mathbb{P}(d_1)(u_a^+(1)) + \mathbb{P}(d_2)(u_a^-(1)) + \mathbb{P}(d_3)(u_a(1)^-) = \mathbb{P}(d_1)(u_a(2)^-) + \mathbb{P}(d_2)(u_a(2)^+) + \mathbb{P}(d_3)(u_a(2)^-)$$
$$\mathbb{P}(d_1)(0) + \mathbb{P}(d_2)(4) + (0)(4) = \mathbb{P}(d_1)(3) + \mathbb{P}(d_2)(0) + (0)(3)$$
$$\mathbb{P}(d_2)(4) = \mathbb{P}(d_1)(3)$$

We also know that the probabilities of each outcome must sum to 1 to complete the definition of a mixed strategy, giving us this system of equations:

$$\mathbb{P}(d_2)(4) = \mathbb{P}(d_1)(3)$$
$$\mathbb{P}(d_1) + \mathbb{P}(d_2) = 1$$
$$\mathbb{P}(d_1) = \frac{4}{7}$$
$$\mathbb{P}(d_2) = \frac{3}{7}$$

So, $c_1 = 1 - c_2 = \frac{4}{7}$ *and* $c_2 = \frac{3}{7}$*. Now, we can calculate the utility for the defender given that the attacker chooses either target 1 or target 2 to attack, and define the utility of the defender to be the greater of the two payoffs given that ties are broken in the leader's favor:*

$$u_d(t,c) = (c_t)(u_d^+(t)) + (1-c_t)(u_d^-(t))$$

$$\begin{aligned}
u_d(1,c) &= (c_1)(u_d^+(1)) + (1-c_1)(u_d^-(1)) \\
&= \left(\frac{4}{7}\right)(0) + \left(\frac{3}{7}\right)(-2) \\
&= -\frac{6}{7} \\
u_d(2,c) &= (c_2)(u_d^+(2)) + (1-c_2)(u_d^-(1)) \\
&= \left(\frac{3}{7}\right)(0) + \left(\frac{4}{7}\right)(-1) \\
&= -\frac{4}{7}
\end{aligned}$$

So, $-\frac{4}{7} > -\frac{6}{7}$*, and therefore the defender's utility will be* $-\frac{4}{7}$*.*