

How Optimal Can We Get: Stochastic and Adversarial Reinforcement Learning

MIT PRIMES, Mentor: Mayuri Sridhar


Alicia Li and Mati Yablon

MIT

October 16, 2022

A Cute Robot in A Cute Maze

We (a cute robot) need to find the optimal path in this maze!



How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

A Cute Robot in A Cute Maze

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon


Background

Our Approach

Conclusion

References

We (a cute robot) need to find the optimal path in this maze!



We could try every path in the maze, but this is inefficient :(

A Cute Robot in A Cute Maze

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon


Background

Our Approach

Conclusion

References

We (a cute robot) need to find the optimal path in this maze!



We could try every path in the maze, but this is inefficient :(
Let's use Reinforcement Learning! Every time we take an **action**, we receive a **reward**, which shapes our future actions.

A Cute Robot in A Cute Maze

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon


Background

Our Approach

Conclusion

References

We (a cute robot) need to find the optimal path in this maze!



We could try every path in the maze, but this is inefficient :(
Let's use Reinforcement Learning! Every time we take an **action**, we receive a **reward**, which shapes our future actions.
Let's formalize this notion...

Markov Decision Processes

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Definition of MDP (Markov Decision Process)

$$\mathcal{M} := (\mathcal{S}, \mathcal{A}, R, \mathcal{P})$$

Markov Decision Processes

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Definition of MDP (Markov Decision Process)

$$\mathcal{M} := (\mathcal{S}, \mathcal{A}, R, \mathcal{P})$$

- \mathcal{S} is **state space**: Set of all states in which the agent may be
- \mathcal{A} is **action space**: Set of all actions which the agent may take in a state
- $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is **reward function**: Outputs the reward given to the agent when taking action a in state s
- $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is **transition dynamics function**: Outputs the probability of the agent transitioning to new state s' if it takes action a in state s

ϵ -greedy Policy

Definition of policy π

A policy π is a mapping of the state and action spaces to a probability that dictates the agent's behavior.

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

ϵ -greedy Policy

Definition of policy π

A policy π is a mapping of the state and action spaces to a probability that dictates the agent's behavior.

ϵ -greedy:

- Probability ϵ : sample random action
- Probability $1 - \epsilon$: take best perceived action $\arg \max_a Q(s, a)$.


ϵ -greedy Policy

Definition of policy π

A policy π is a mapping of the state and action spaces to a probability that dictates the agent's behavior.

ϵ -greedy:

- Probability ϵ : sample random action
- Probability $1 - \epsilon$: take best perceived action $\arg \max_a Q(s, a)$.



Q-values

Now how does RL work? Central goal is to learn an optimal policy (i.e. behavior)

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Q-values

Now how does RL work? Central goal is to learn an optimal policy (i.e. behavior)

Q-values store how “good” a state is

Approaches the expected value $Q(s_t, a_t) \approx \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t R_t]$.

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Q-values

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Now how does RL work? Central goal is to learn an optimal policy (i.e. behavior)

Q-values store how “good” a state is

Approaches the expected value $Q(s_t, a_t) \approx \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t R_t]$.

Learned via Bellman optimality equation:

$$Q(s_t, a_t) \leftarrow (1 - \alpha)Q(s_t, a_t) + \alpha(R_t + \gamma \max_a Q(s_{t+1}, a)).$$

Q-values

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Now how does RL work? Central goal is to learn an optimal policy (i.e. behavior)


Q-values store how “good” a state is

Approaches the expected value $Q(s_t, a_t) \approx \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t R_t]$.

Learned via Bellman optimality equation:


$$Q(s_t, a_t) \leftarrow (1 - \alpha)Q(s_t, a_t) + \alpha(R_t + \gamma \max_a Q(s_{t+1}, a)).$$

Heat map of learned Q-values:



Adversarial RL

What if something perturbs the MDP?



How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background


Our Approach

Conclusion

References

Adversarial RL

What if something perturbs the MDP?



Performance can be degraded by:

- Human biases
- Modeling errors
- Actual adversaries

Robust RL

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Definition

Robust RL aims to find the best-performing policy in the worst-case scenario. It can be framed as a 2-player zero-sum game.

Objective: Find the policy π that satisfies:

$$\max_{\pi} \min_{\mathcal{P}} \mathbb{E}_{\pi, \mathcal{P}} \left[\sum_t R_t \right],$$

where \mathcal{P} is the environment and R_t is the reward at time t .

Robust RL

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Definition

Robust RL aims to find the best-performing policy in the worst-case scenario. It can be framed as a 2-player zero-sum game.

Objective: Find the policy π that satisfies:

$$\max_{\pi} \min_{\mathcal{P}} \mathbb{E}_{\pi, \mathcal{P}} \left[\sum_t R_t \right],$$

where \mathcal{P} is the environment and R_t is the reward at time t .

Robust RL Methods Include:

- Injecting noise into the environment during training (Maximum Entropy)

Robust RL

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Definition

Robust RL aims to find the best-performing policy in the worst-case scenario. It can be framed as a 2-player zero-sum game.

Objective: Find the policy π that satisfies:

$$\max_{\pi} \min_{\mathcal{P}} \mathbb{E}_{\pi, \mathcal{P}} \left[\sum_t R_t \right],$$

where \mathcal{P} is the environment and R_t is the reward at time t .

Robust RL Methods Include:

- Injecting noise into the environment during training (Maximum Entropy)
- Train the agent in an environment with an adversary that corrupts the reward function

Best of Both Worlds

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

We want to perform well in **all** environments, not just worst-case scenarios...

Best of Both Worlds

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

We want to perform well in **all** environments, not just worst-case scenarios... Best of Both Worlds!

Definition

Best of Both Worlds: We want performance that degrades gracefully with an increasing corruption level, can be used in RL

Best of Both Worlds Methods:

- Layering algorithms designed for varying corruption levels

Problem Setting

Previous work [2] in Best-Of-Both-Worlds has focused on bandit MDPs We consider layered

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Problem Setting

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References


Previous work [2] in Best-Of-Both-Worlds has focused on bandit MDPs. We consider layered. For every sample, our adversary is able to:

- Corrupt the edges that victim traverses with probability p
- Corrupt that edge's reward by a maximum of δ each

Problem Setting

Previous work [2] in Best-Of-Both-Worlds has focused on bandit MDPs. We consider layered. For every sample, our adversary is able to:

- Corrupt the edges that victim traverses with probability p
- Corrupt that edge's reward by a maximum of δ each



Calculating Adversarial Budget to Switch Paths

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Adversary wants to make optimal path seem worse than some suboptimal path, how much budget does it have? (victim traverses each path equally)

Calculating Adversarial Budget to Switch Paths

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Adversary wants to make optimal path seem worse than some suboptimal path, how much budget does it have? (victim traverses each path equally)
Consider the following MDP:

Calculating Adversarial Budget to Switch Paths

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon


Background

Our Approach

Conclusion

References

Adversary wants to make optimal path seem worse than some suboptimal path, how much budget does it have? (victim traverses each path equally)
Consider the following MDP:



Calculating Adversarial Budget to Switch Paths

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon


Background

Our Approach

Conclusion

References

Adversary wants to make optimal path seem worse than some suboptimal path, how much budget does it have? (victim traverses each path equally)
Consider the following MDP:



Naive Approach: $p\delta$ each from corrupting AB up and CE down whenever paths 3 and 1 are traversed, yielding $2p\delta$

Calculating Adversarial Budget to Switch Paths

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background


Our Approach

Conclusion

References

Adversary wants to make optimal path seem worse than some suboptimal path, how much budget does it have? (victim traverses each path equally)

Consider the following MDP:




Naive Approach: $p\delta$ each from corrupting AB up and CE down whenever paths 3 and 1 are traversed, yielding $2p\delta$

Our Approach: $2p\delta +$ extra $\frac{1}{2}p\delta$ of “free corruption” from corrupting AC whenever path 2 is traversed

Adversarial Attack

Let's attack! Given that $p = 0.25$ and $\delta = 4; p\delta = 1$




Algorithm 1 Adversarial Attack: Switching Non-disjoint Paths

```
1:  $P \leftarrow P^*$   $\triangleright$  Current path to perturb; we will iterate and find a better one
2: for all  $P_i \in \mathcal{M} \wedge P_i \neq P^*$  do  $\triangleright$  Iterate only through suboptimal paths
3:    $b_i \leftarrow 2$   $\triangleright$  Add budget for the two disjoint edges on the paths we want to
   switch
4:   for all  $P_j \notin \{P_i, P^*\}$  do  $\triangleright$  Iterate through paths that are not optimal or  $P_i$ 
5:      $a_e^* = \infty$ 
6:     for all  $e \in P_j$  do
7:       if  $(e \in P^* \oplus e \in P_i) \wedge a_e < a_e^*$  then
8:          $a_e^* \leftarrow a_e$   $\triangleright$  Calculate budget
9:       end if
10:    end for
11:     $b_i \leftarrow b_i + \frac{1}{a_e^*}$ 
12:  end for
13:  if  $R(P_i) < R(P) \wedge R(P^*) - b_i p \delta < R(P_i)$  then
14:     $P \leftarrow P_i$   $\triangleright$  Compare current path to best path found thus far
15:  end if
16: end for
17: output  $P$ 
```

Adversarial Attack

Let's attack! Given that $p = 0.25$ and $\delta = 4; p\delta = 1$




Algorithm 1 Adversarial Attack: Switching Non-disjoint Paths

```
1:  $P \leftarrow P^*$   $\triangleright$  Current path to perturb; we will iterate and find a better one
2: for all  $P_i \in \mathcal{M} \wedge P_i \neq P^*$  do  $\triangleright$  Iterate only through suboptimal paths
3:    $b_i \leftarrow 2$   $\triangleright$  Add budget for the two disjoint edges on the paths we want to
   switch
4:   for all  $P_j \notin \{P_i, P^*\}$  do  $\triangleright$  Iterate through paths that are not optimal or  $P_i$ 
5:      $a_e^* = \infty$ 
6:     for all  $e \in P_j$  do
7:       if  $(e \in P^* \oplus e \in P_i) \wedge a_e < a_e^*$  then
8:          $a_e^* \leftarrow a_e$   $\triangleright$  Calculate budget
9:       end if
10:    end for
11:     $b_i \leftarrow b_i + \frac{1}{a_e^*}$ 
12:  end for
13:  if  $R(P_i) < R(P) \wedge R(P^*) - b_i p \delta < R(P_i)$  then
14:     $P \leftarrow P_i$   $\triangleright$  Compare current path to best path found thus far
15:  end if
16: end for
17: output  $P$ 
```

Budget of Switching:

Adversarial Attack

Let's attack! Given that $p = 0.25$ and $\delta = 4; p\delta = 1$



Algorithm 1 Adversarial Attack: Switching Non-disjoint Paths


```
1:  $P \leftarrow P^*$   $\triangleright$  Current path to perturb; we will iterate and find a better one
2: for all  $P_i \in \mathcal{M} \wedge P_i \neq P^*$  do  $\triangleright$  Iterate only through suboptimal paths
3:    $b_i \leftarrow 2$   $\triangleright$  Add budget for the two disjoint edges on the paths we want to
   switch
4:   for all  $P_j \notin \{P_i, P^*\}$  do  $\triangleright$  Iterate through paths that are not optimal or  $P_i$ 
5:      $a_e^* = \infty$ 
6:     for all  $e \in P_j$  do
7:       if  $(e \in P^* \oplus e \in P_i) \wedge a_e < a_e^*$  then
8:          $a_e^* \leftarrow a_e$   $\triangleright$  Calculate budget
9:       end if
10:    end for
11:     $b_i \leftarrow b_i + \frac{1}{a_e^*}$ 
12:  end for
13:  if  $R(P_i) < R(P) \wedge R(P^*) - b_i p \delta < R(P_i)$  then
14:     $P \leftarrow P_i$   $\triangleright$  Compare current path to best path found thus far
15:  end if
16: end for
17: output  $P$ 
```

Budget of Switching:

1 with 3: $2\frac{5}{6}$, not enough to switch paths :(

Adversarial Attack

Let's attack! Given that $p = 0.25$ and $\delta = 4; p\delta = 1$



Algorithm 1 Adversarial Attack: Switching Non-disjoint Paths

```
1:  $P \leftarrow P^*$   $\triangleright$  Current path to perturb; we will iterate and find a better one
2: for all  $P_i \in \mathcal{M} \wedge P_i \neq P^*$  do  $\triangleright$  Iterate only through suboptimal paths
3:    $b_i \leftarrow 2$   $\triangleright$  Add budget for the two disjoint edges on the paths we want to
   switch
4:   for all  $P_j \notin \{P_i, P^*\}$  do  $\triangleright$  Iterate through paths that are not optimal or  $P_i$ 
5:      $a_e^* = \infty$ 
6:     for all  $e \in P_j$  do
7:       if  $(e \in P^* \oplus e \in P_i) \wedge a_e < a_e^*$  then
8:          $a_e^* \leftarrow a_e$   $\triangleright$  Calculate budget
9:       end if
10:    end for
11:     $b_i \leftarrow b_i + \frac{1}{a_e^*}$ 
12:  end for
13:  if  $R(P_i) < R(P) \wedge R(P^*) - b_i p \delta < R(P_i)$  then
14:     $P \leftarrow P_i$   $\triangleright$  Compare current path to best path found thus far
15:  end if
16: end for
17: output  $P$ 
```


Budget of Switching:

1 with 3: $2\frac{5}{6}$, not enough to switch paths :(

2 with 3: 2, enough to switch paths :)

Adversarial Attack

Let's attack! Given that $p = 0.25$ and $\delta = 4; p\delta = 1$



Algorithm 1 Adversarial Attack: Switching Non-disjoint Paths

```
1:  $P \leftarrow P^*$   $\triangleright$  Current path to perturb; we will iterate and find a better one
2: for all  $P_i \in \mathcal{M} \wedge P_i \neq P^*$  do  $\triangleright$  Iterate only through suboptimal paths
3:    $b_i \leftarrow 2$   $\triangleright$  Add budget for the two disjoint edges on the paths we want to
   switch
4:   for all  $P_j \notin \{P_i, P^*\}$  do  $\triangleright$  Iterate through paths that are not optimal or  $P_i$ 
5:      $a_e^* = \infty$ 
6:     for all  $e \in P_j$  do
7:       if  $(e \in P^* \oplus e \in P_i) \wedge a_e < a_e^*$  then
8:          $a_e^* \leftarrow a_e$   $\triangleright$  Calculate budget
9:       end if
10:    end for
11:     $b_i \leftarrow b_i + \frac{1}{a_e^*}$ 
12:  end for
13:  if  $R(P_i) < R(P) \wedge R(P^*) - b_i p \delta < R(P_i)$  then
14:     $P \leftarrow P_i$   $\triangleright$  Compare current path to best path found thus far
15:  end if
16: end for
17: output  $P$ 
```

Budget of Switching:


1 with 3: $2\frac{5}{6}$, not enough to switch paths :(

2 with 3: 2, enough to switch paths :)

4 with 3: $2\frac{1}{2}$, enough to switch paths :)

Adversarial Attack

Let's attack! Given that $p = 0.25$ and $\delta = 4; p\delta = 1$



Algorithm 1 Adversarial Attack: Switching Non-disjoint Paths

```
1:  $P \leftarrow P^*$   $\triangleright$  Current path to perturb; we will iterate and find a better one
2: for all  $P_i \in \mathcal{M} \wedge P_i \neq P^*$  do  $\triangleright$  Iterate only through suboptimal paths
3:    $b_i \leftarrow 2$   $\triangleright$  Add budget for the two disjoint edges on the paths we want to
   switch
4:   for all  $P_j \notin \{P_i, P^*\}$  do  $\triangleright$  Iterate through paths that are not optimal or  $P_i$ 
5:      $a_e^* = \infty$ 
6:     for all  $e \in P_j$  do
7:       if  $(e \in P^* \oplus e \in P_i) \wedge a_e < a_e^*$  then
8:          $a_e^* \leftarrow a_e$   $\triangleright$  Calculate budget
9:       end if
10:    end for
11:     $b_i \leftarrow b_i + \frac{1}{a_e^*}$ 
12:  end for
13:  if  $R(P_i) < R(P) \wedge R(P^*) - b_i p \delta < R(P_i)$  then
14:     $P \leftarrow P_i$   $\triangleright$  Compare current path to best path found thus far
15:  end if
16: end for
17: output  $P$ 
```

Budget of Switching:

1 with 3: $2\frac{5}{6}$, not enough to switch paths :(

2 with 3: 2, enough to switch paths :)

4 with 3: $2\frac{1}{2}$, enough to switch paths :)

We choose to switch path 3 with path 4

Proof of Optimality (Sketch)

Our algorithm is **optimal**

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Proof of Optimality (Sketch)

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Our algorithm is **optimal**

- 1 Reduce showing that our algorithm picks the optimal path to showing our algorithm calculates budget optimally

Proof of Optimality (Sketch)

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Our algorithm is **optimal**

- 1 Reduce showing that our algorithm picks the optimal path to showing our algorithm calculates budget optimally
 - Suppose otherwise that our algorithm didn't pick path with lowest reward. This means we didn't calculate budget optimally for a path with lower reward. Thus, we will prove our algorithm picks set of corrupted edges optimally.

Proof of Optimality (Sketch)

Our algorithm is **optimal**

- 1 Reduce showing that our algorithm picks the optimal path to showing our algorithm calculates budget optimally
 - Suppose otherwise that our algorithm didn't pick path with lowest reward. This means we didn't calculate budget optimally for a path with lower reward. Thus, we will prove our algorithm picks set of corrupted edges optimally.
- 2 Picking just one edge in each traversal is optimal.

Proof of Optimality (Sketch)

Our algorithm is **optimal**

- 1 Reduce showing that our algorithm picks the optimal path to showing our algorithm calculates budget optimally
 - Suppose otherwise that our algorithm didn't pick path with lowest reward. This means we didn't calculate budget optimally for a path with lower reward. Thus, we will prove our algorithm picks set of corrupted edges optimally.
- 2 Picking just one edge in each traversal is optimal.
- 3 Our algorithm picks the edge that is optimal in every traversal.

Proof of Optimality (Sketch)

Our algorithm is **optimal**

- 1 Reduce showing that our algorithm picks the optimal path to showing our algorithm calculates budget optimally
 - Suppose otherwise that our algorithm didn't pick path with lowest reward. This means we didn't calculate budget optimally for a path with lower reward. Thus, we will prove our algorithm picks set of corrupted edges optimally.
- 2 Picking just one edge in each traversal is optimal.
- 3 Our algorithm picks the edge that is optimal in every traversal.
 - Suppose otherwise that there exists an edge set to corrupt that is more optimal. Consider edges that differ from algorithm's set to optimal set.

Proof of Optimality (Sketch)

Our algorithm is **optimal**

- 1 Reduce showing that our algorithm picks the optimal path to showing our algorithm calculates budget optimally
 - Suppose otherwise that our algorithm didn't pick path with lowest reward. This means we didn't calculate budget optimally for a path with lower reward. Thus, we will prove our algorithm picks set of corrupted edges optimally.
- 2 Picking just one edge in each traversal is optimal.
- 3 Our algorithm picks the edge that is optimal in every traversal.
 - Suppose otherwise that there exists an edge set to corrupt that is more optimal. Consider edges that differ from algorithm's set to optimal set.
 - These substitutions will not yield greater corruption since algorithm chooses edge on least number of paths, which guarantees the maximum amount.

Adversarial Algorithm Against ϵ -Greedy Victim

We have an adversarial strategy against a simple victim... now we consider a smart one!

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

Adversarial Algorithm Against ϵ -Greedy Victim

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

We have an adversarial strategy against a simple victim... now we consider a smart one!

What is the optimal strategy for an adversary against a victim with an ϵ -greedy policy?

Adversarial Algorithm Against ϵ -Greedy Victim

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

We have an adversarial strategy against a simple victim... now we consider a smart one!

What is the optimal strategy for an adversary against a victim with an ϵ -greedy policy?

- Can't assume equal path traversal, sample complexity is tricky

Adversarial Algorithm Against ϵ -Greedy Victim

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach


Conclusion

References

We have an adversarial strategy against a simple victim... now we consider a smart one!

What is the optimal strategy for an adversary against a victim with an ϵ -greedy policy?

- Can't assume equal path traversal, sample complexity is tricky
- Perturbing edges not in the optimal path or path to be switched has an effect, especially for small budget



Adversarial Algorithm Against ϵ -Greedy Victim

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach


Conclusion

References

We have an adversarial strategy against a simple victim... now we consider a smart one!

What is the optimal strategy for an adversary against a victim with an ϵ -greedy policy?

- Can't assume equal path traversal, sample complexity is tricky
- Perturbing edges not in the optimal path or path to be switched has an effect, especially for small budget



- Chebyshev's Inequality bound on expected reward of this strategy: it is less than $(r_1 + r_3) \cdot (N_1 + N_3) p \delta^2 \frac{1-p}{(r_1 - r_3)^2}$

Future Work

How Optimal
Can We Get:
Stochastic
and
Adversarial
Reinforcement
Learning

Alicia Li and
Mati Yablon

Background

Our Approach

Conclusion

References

- How does victim defend against adversary strategy outlined above using Best-of-Both-Worlds?
 - Devise layering algorithm for victim defense
- More generally: set up minimax between victim and adversary to fully describe their behaviors in the MDP
 - What is the value of corrupting a path that is neither the optimal path nor the path we are trying to switch with it? Is there value in confusing the victim in this way? When is this helpful?

Acknowledgements

We would like to thank...

- MIT PRIMES; Dr. Slava Gerovitch and Dr. Srinivasa Devadas for this wonderful opportunity
- Mayuri Sridhar for being an amazing mentor
- You!



References

- [1] Ben Eysenbach. *Maximum Entropy RL (Provably) Solves Some Robust RL Problems*. <https://bair.berkeley.edu/blog/2021/03/10/maxent-robust-rl/>. Accessed 29 June 2022.
- [2] Thodoris Lykouris, Vahab Mirrokni, and Renato Paes Leme. “Stochastic bandits robust to adversarial corruptions”. In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. 2018, pp. 114–122.
- [3] Lerrel Pinto et al. “Robust adversarial reinforcement learning”. In: *International Conference on Machine Learning*. PMLR. 2017, pp. 2817–2826.
- [4] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning, second edition: An Introduction*. 2018. ISBN: 9780262352703.