



# New Approach on 2048

Game Optimization  
Mon. May 1st, 2023

## Team Members

- Minho Cho (MEng, Data Analytics Concentration)
- Jason Andrade (MEng, Software Engineering Concentration)
- Alyssa Lowe (MEng, Software Engineering Concentration)
- Ethan Weaver (MEng, Data Analytics Concentration)

# Presentation Outline

- Introduction
- Problem Statement
- Approaches
- Implementation Details
- Demo
- Results
- Lesson Learned
- Future Work

# Introduction

- The game 2048 is a popular web based strategy game
- The objective of the game is to create a tile with the value of 2048, or greater
- A player can move all tiles on the board in one of 4 direction: up, down, left, or right
- If two tiles collide, they combine to form a tile of twice their value
- After each move, a new tile of value 2 or 4 is generated in a random empty space




Figure 1. Game Play

# Problem Statement

- Build an Expert Agent based on a Learning Agent using **different algorithms** to maximize game score making **optimal movements**.
- To compare different types of AI algorithms to see which performs best at the game 2048.




Figure 2. Losing Game Play

# Random Agent (Baseline)


Makes completely random moves

<b>Average Score</b>	<b>Standard Deviation</b>
962	473

<b>Tile Value</b>	<b>16</b>	<b>32</b>	<b>64</b>	<b>128</b>	<b>256</b>
<b>Success Rate</b>	99.99%	99.52%	88.7%	46.05%	4.22%

# Approaches

- Expectimax Algorithm



- Q Learning Algorithm

- Temporal Difference (TD) Algorithm

```

1: function PLAY GAME
2:   score ← 0
3:   s ← INITIALIZE GAME STATE
4:   while ¬IS TERMINAL STATE(s) do
5:     a ← arg maxa' ∈ A(s) EVALUATE(s, a')
6:     r, s', s'' ← MAKE MOVE(s, a)
7:     if LEARNING ENABLED then
8:       LEARN EVALUATION(s, a, r, s', s'')
9:     score ← score + r
10:    s ← s''
11:  return score
12:
13: function MAKE MOVE(s, a)
14:  s', r ← COMPUTE AFTERSTATE(s, a)
15:  s' ← ADD RANDOM TILE(s')
16:  return (r, s', s'')
  
```

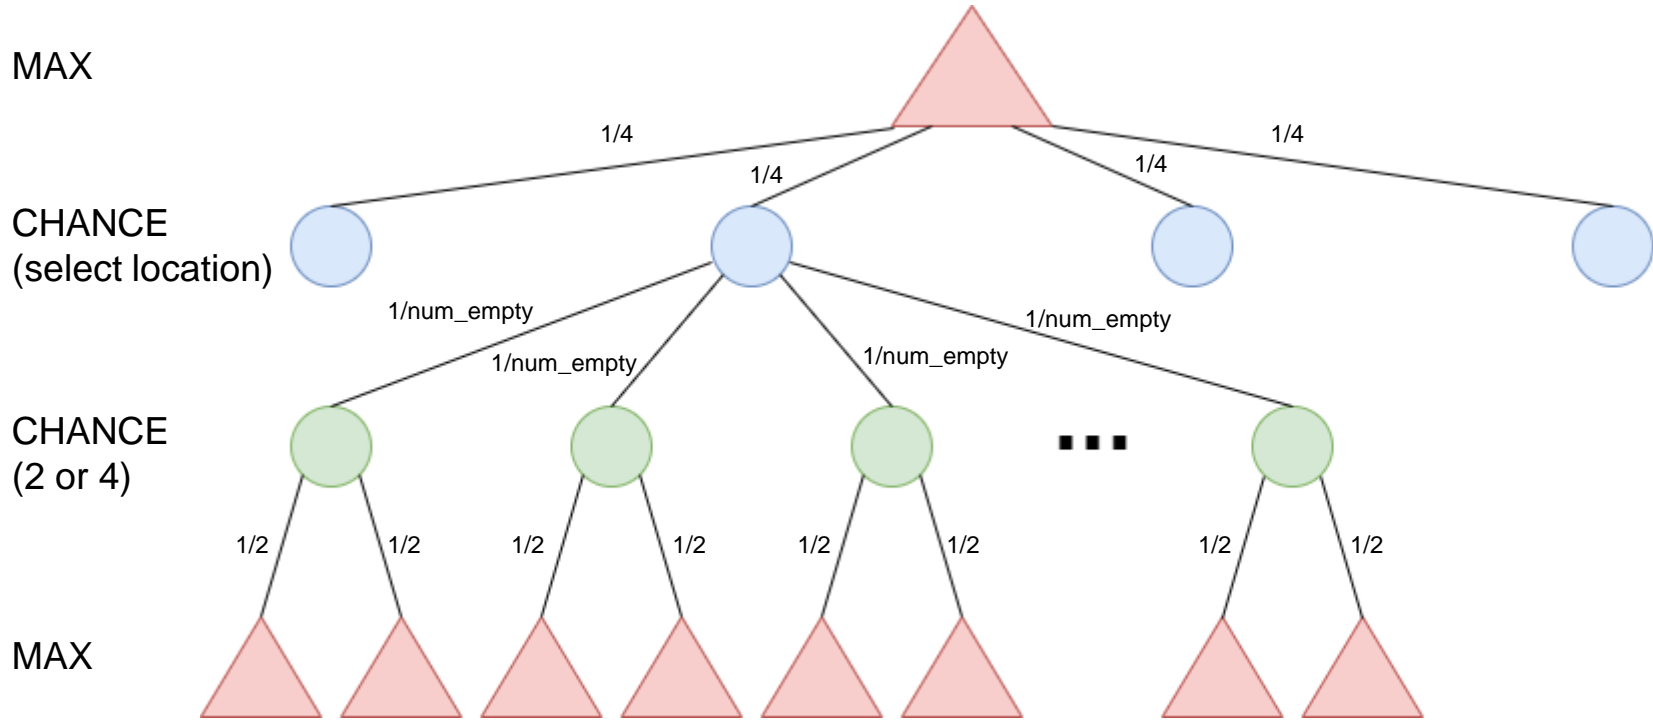
```

1: function EVALUATE(s, a)
2:   return Va(s)
3:
4: function LEARN EVALUATION(s, a, r, s', s'')
5:   vnext ← maxa' ∈ A(s'') Va'(s'')
6:   Va(s) ← Va(s) + α(r + vnext - Va(s))
  
```

```

1: function EVALUATE(s, a)
2:   s', r ← COMPUTE AFTERSTATE(s, a)
3:   return r + V(s')
4:
5: function LEARN EVALUATION(s, a, r, s', s'')
6:   anext ← arg maxa' ∈ A(s'') EVALUATE(s'', a')
7:   s'next, rnext ← COMPUTE AFTERSTATE(s'', anext)
8:   V(s') ← V(s') + α(rnext + V(s'next) - V(s'))
  
```

# Expectimax Implementation





# Expectimax Evaluation Function

Number of Empty  
Cells on the Board

	2		
	2		

Number of Merges  
Possible

		2	4
		2	4
16	8	2	4
64	32	16	4

Monotonicity of  
Rows and Columns

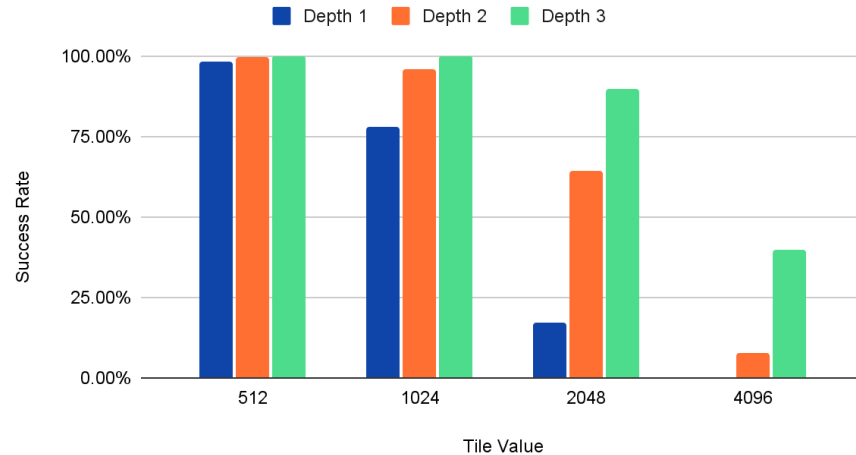
8	32	64	512
4	8	16	256
2	4	8	32
		4	8

$$E(x) = 350 * num\_empty + 800 * merges\_possible + 20 * monotonicity$$

# Expectimax Results

Depth	Average Score	Standard Deviation
1	15,883	6,841
2	27,962	12,632
3	40,369	16,109

Success Rate of Tile Values



# Q-Learning

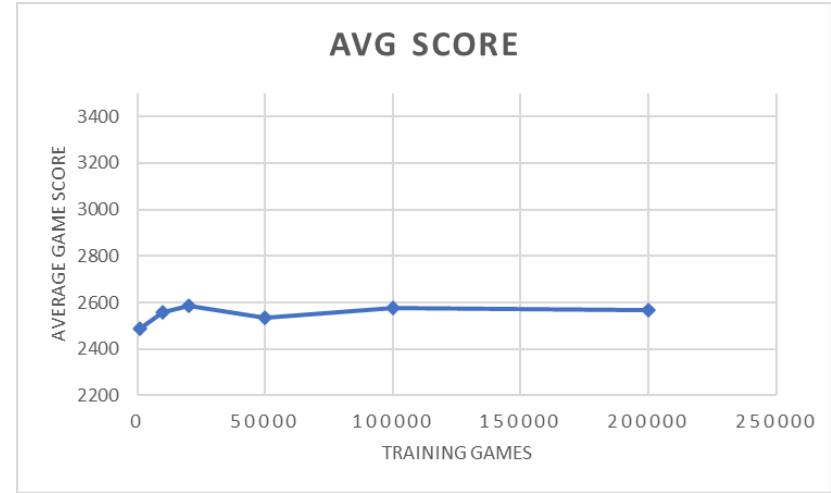
- Two states considered:
  - “current state” -  $s$ ,
  - “nextstate” -  $s'$
- Q-learning evaluates the state by tracking the rewards for a state given an action.
- The “Learn function” adjusts the current states reward within the given action

```
1: function PLAY GAME
2:    $score \leftarrow 0$ 
3:    $s \leftarrow$  INITIALIZE GAME STATE
4:   while  $\neg$ IS TERMINAL STATE( $s$ ) do
5:      $a \leftarrow \arg \max_{a' \in A(s)} \text{EVALUATE}(s, a')$ 
6:      $r, s', s'' \leftarrow$  MAKE MOVE( $s, a$ )
7:     if LEARNING ENABLED then
8:       LEARN EVALUATION( $s, a, r, s', s''$ )
9:      $score \leftarrow score + r$ 
10:     $s \leftarrow s''$ 
11:   return  $score$ 
12:
13: function MAKE MOVE( $s, a$ )
14:    $s', r \leftarrow$  COMPUTE AFTERSTATE( $s, a$ )
15:    $s'' \leftarrow$  ADD RANDOM TILE( $s'$ )
16:   return ( $r, s', s''$ )
```

```
1: function EVALUATE( $s, a$ )
2:   return  $V_a(s)$ 
3:
4: function LEARN EVALUATION( $s, a, r, s', s''$ )
5:    $v_{next} \leftarrow \max_{a' \in A(s'')} V_{a'}(s'')$ 
6:    $V_a(s) \leftarrow V_a(s) + \alpha(r + v_{next} - V_a(s))$ 
```

# Q-Learning Algorithm Results

alpha = 0.0025, real games = 1000



Number of Training Games	Average Score	Standard Deviation	% to reach 512
1,000	2486.62	1213.5	4.6%
10,000	2557.628	1220.7	4.3%
100,000	2,577.9	1,214.1	4.7%
200,000	2568.39	1210.6	5.2%

# Temporal Difference Afterstate Learning

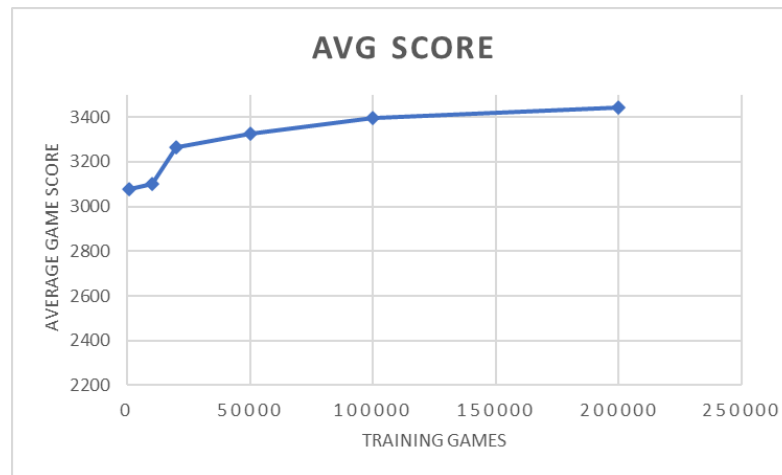
- Three states considered:
  - “current state” -  $s$ ,
  - “afterstate” -  $s'$
  - “nextstate” -  $s''$
- TD afterstate learning evaluates the result of moves based on the reward assigned to an afterstate.
- The “Learn function” adjusts an afterstates value by determining the possible reward of the next afterstate if an optimal move is made.

```
1: function PLAY GAME
2:    $score \leftarrow 0$ 
3:    $s \leftarrow$  INITIALIZE GAME STATE
4:   while  $\neg$ IS TERMINAL STATE( $s$ ) do
5:      $a \leftarrow \arg \max_{a' \in A(s)} \text{EVALUATE}(s, a')$ 
6:      $r, s', s'' \leftarrow$  MAKE MOVE( $s, a$ )
7:     if LEARNING ENABLED then
8:       LEARN EVALUATION( $s, a, r, s', s''$ )
9:      $score \leftarrow score + r$ 
10:     $s \leftarrow s''$ 
11:   return  $score$ 
12:
13: function MAKE MOVE( $s, a$ )
14:    $s', r \leftarrow$  COMPUTE AFTERSTATE( $s, a$ )
15:    $s'' \leftarrow$  ADD RANDOM TILE( $s'$ )
16:   return ( $r, s', s''$ )
```

```
1: function EVALUATE( $s, a$ )
2:    $s', r \leftarrow$  COMPUTE AFTERSTATE( $s, a$ )
3:   return  $r + V(s')$ 
4:
5: function LEARN EVALUATION( $s, a, r, s', s''$ )
6:    $a_{next} \leftarrow \arg \max_{a' \in A(s'')} \text{EVALUATE}(s'', a')$ 
7:    $s'_{next}, r_{next} \leftarrow$  COMPUTE AFTERSTATE( $s'', a_{next}$ )
8:    $V(s') \leftarrow V(s') + \alpha(r_{next} + V(s'_{next}) - V(s'))$ 
```

# TD-Learning Algorithm Results

alpha = 0.0025, real games = 1000



Number of Training Games	Average Score	Standard Deviation	% to reach 512	% to reach 1024
1,000	3077.408	1551.6	13.7%	0.1%
10,000	3098.74	1568.2	13.7%	0.1%
100,000	3394.1	1652.3	18.5%	0.03%
200,000	3441.46	1683.82	19.1%	0.5%

# Results Comparison

Algorithm	Average Score (Depth 3/100,000 Games)	Standard Deviation (Depth 3/100,000 Games)
Expectimax	40,369	16,109
TD-Learning	3,394	1,652
Q-Learning	2,577	1,214
Random	962	473

- Expectimax far outperforms the other methods since it has more information in the form of its evaluation function
- TD-learning has an overall better average score per learning games than Q-learning.
- TD-learning has a better learning rate than Q-learning
- Overall, results coincide with the amount of information that each algorithm has access to.


# Lessons Learned

- Time and computing resource limitations:
  - When using personal computing devices, some group members are not able to run all algorithms implemented
    - For example: running expectimax at depth 3 for one game would take up to 3 hours for some, or approx 30 minutes for others
    - Having to run games overnight to generate result
- Different algorithms perform better or worse depending on the specific game mechanics and data available for training
  - Small changes to the Expectimax Evaluation Function can have huge effects on the performance of the agent



# Future Work

- Improve the efficiency of all algorithms or at least improve upon best performing algorithm, expectimax, to maximize the expected outcome with other evaluation metrics such as max tile reached vs moves taken vs average score
- Similar to Pacman course projects, can continue to experiment with other reinforcement learning techniques that are best leveraged for gamification
- Use cloud computing resources proposed by Professor during class instead of using personal devices



# Q&A