



What are Games? Which Kinds Exist?

ational

Jniversitv

Pascal Bercher

A **game** consists of a set of one or more **players**, a set of **moves** for the players, and a specification of **payoffs** (outcomes) for each combination of **strategies** (also called policy).

What kinds of restrictions can games have?

- Perfect information vs. imperfect information
- (One-player games vs.) Two-player games vs. multi-player games
- Zero-sum games vs. non-zero-sum games
- Games with chance (randomness) vs. games without chance





time an **optimal strategy** can be found for games of **perfect information** by enumerating paths of a **game tree**. However, in practice this can only be done for small games.





5.32

What's a Strategy?

What are we Looking For?

0000

What are we looking for?

- Game AI (strategy) vs. game theoretic outcome!
- Just because we have an AI that beats all humans, it doesn't mean the game is solved!

What's the game theoretic outcome?

- The outcome of the game assuming all players play rational.
- Rationality = optimization of expected reward.
- $\bullet~$ Outcome is known? \rightarrow The respective game is "solved".



• If the game tree is "sufficiently small" we can search in it to find

Consider two players, MAX and MIN. MAX tries to maximize his/her

We assume that the players are rational, taking turns, and that the

MiniMax: How to Solve Small Games?

Using search to solve a game:

and extract a strategy.

But we still need to do that efficiently!

own score, and player MIN tries to minimize it.

Solving Small Games



MiniMax: The MiniMax Algorithm

The MiniMax algorithm allows each player to compute their optimal move on a game tree of alternating MAX and MIN nodes.

The value of a node is the payoff for a game that is played optimally from that node until the end of the game.

max-value(s)

```
if state s is a leaf then

\ \ return payoff(s)

v := -\infty

forall successor states s' of s do

\ \ v := \max \{v, \min\text{-value}(s')\}

return v
```

Pascal Bercher

min-value(s)

```
if state s is a leaf then

\ \ return payoff(s)

v := \infty

forall successor states s' of s do

\ \ v := \min \{v, \max\text{-value}(s')\}

return v
```

Australian National University

game is zero-sum.

Solving Small Games

000000

MAX player plays X, MIN plays O. Outcomes (black boxes) are from the perspective of the MAX player (i.e., 1 is a win, -1 a loss, 0 a draw).



MiniMax: Example: Tic Tac Toe

000000

MAX player plays X, MIN plays O. Outcomes (black boxes) are from the perspective of the MAX player (i.e., 1 is a win, -1 a loss, 0 a draw).





MiniMax: Example: Tic Tac Toe

MAX player plays X, MIN plays O. Outcomes (black boxes) are from the perspective of the MAX player (i.e., 1 is a win, -1 a loss, 0 a draw).



Solving Small Games

000000

MAX player plays X, MIN plays O. Outcomes (black boxes) are from the perspective of the MAX player (i.e., 1 is a win, -1 a loss, 0 a draw).



MiniMax: Example: Tic Tac Toe

000000

MAX player plays X, MIN plays O. Outcomes (black boxes) are from the perspective of the MAX player (i.e., 1 is a win, -1 a loss, 0 a draw).





MiniMax: Example: Tic Tac Toe

00000

MAX player plays X, MIN plays O. Outcomes (black boxes) are from the perspective of the MAX player (i.e., 1 is a win, -1 a loss, 0 a draw).



Solving Small Games

MAX player plays X, MIN plays O. Outcomes (black boxes) are from the perspective of the MAX player (i.e., 1 is a win, -1 a loss, 0 a draw).



MiniMax: Example: Tic Tac Toe

Solving Small Games

000000

MAX player plays X, MIN plays O. Outcomes (black boxes) are from the perspective of the MAX player (i.e., 1 is a win, -1 a loss, 0 a draw).



Vation Games? Solving Small Games Games with Chance Solving Large Games Game Al Success S MiniMax: Example: Tic Tac Toe MAx player plays X, MIN plays O. Outcomes (black boxes) are from the perspective of the MAx player (i.e., 1 is a win, -1 a loss, 0 a draw).



MiniMax: Example: Tic Tac Toe

00000

Solving Small Games

MAX player plays X, MIN plays O. Outcomes (black boxes) are from the perspective of the MAX player (i.e., 1 is a win, -1 a loss, 0 a draw).



Solving Small Games

0000

MAX player plays X, MIN plays O. Outcomes (black boxes) are from the perspective of the MAX player (i.e., 1 is a win, -1 a loss, 0 a draw).



 α/β Pruning: Can we do better?

000000

- MiniMax suffers from the problem that the number of game states it has to examine is *always* exponential in the number of moves.
- α/β pruning is a method for reducing the number of nodes that need to be evaluated by only considering nodes that may be reached in game play.
- Alpha-beta pruning places bounds on the values appearing anywhere along a path:
 - α is the best (highest) value found so far for MAX
 - β is the best (lowest) value found so far for MIN
 - α and β propagate down the game tree.
 - v propagates up the game tree.

Solving Small Games MiniMax: Space and Time Complexities What is the runtime of MiniMax? • Time: All nodes have to be visited! How many are there? • Assume each game ends after *d* moves (tree depth). Each player has at most *b* moves (branching factor) \rightarrow Runtime is in $O(b^d)$ (exponential!)

What is the space requirement of MiniMax?

- We perform a depth-first search!
- So only the longest path needs to be stored.
- \rightarrow Space is in $O(b \cdot d)$ (linear)

Australian National University		
	Pascal Bercher	12.32

		Solving Small Games	Games with Chance	Solving Large Games	Game AI Success Story
α/β P	runing:	The MiniMax Algo	rithm Extended	By α/β Pruning	

Keep in mind:

- α is the best value found so far for MAX, initialize with $-\infty$.
- β is the best value found so far for MIN, initialize with ∞ .

max-value(s, α, β) min-value(s, α, β) if state s is a leaf then return payoff(s) $v := -\infty$ forall successor states s' of s do $v := \max \{v, \min\text{-value}(s', \alpha, \beta)\}$ if $v > \beta$ then return v $\alpha := \max{\{\alpha, \nu\}}$ return v

if state s is a leaf then return payoff(s) $v := \infty$ forall successor states s' of s do $v := \min \{v, \max\text{-value}(s', \alpha, \beta)\}$ if $v < \alpha$ then return v $\beta := \min \{\beta, \nu\}$ return v















- *b*, the branching factor (available moves per state)
- *d*, the depth (number of moves until game ends)
- For some games that is simply too large!
- So, let's take a look at some examples...

ustraliar

Pascal Bercher

lational

University



The "Size" of Games: Tic Tac Toe

Examples for (estimated) number of reachable (game) states: (Source: https://en.wikipedia.org/wiki/Game_complexity)

- Rough maximum: $3^9 = 19,683$ (including invalid states)
- Actual maximum: 5, 478

Australiar

Vational

University

Pascal Bercher

- Maximum after duplicating symmetries: 765
- There are still 26, 830 possible games!
 (For those states with eliminated duplicates.)
 What's a "game"?
 A path in the MiniMax tree!



Examples for (estimated) number of reachable (game) states:

(Source: https://en.wikipedia.org/wiki/Shannon_number)

- Some maximum: 5 · 10⁵²
- Lower limit on game tree size: 10¹²³
- More conservative estimate on lower limit of game tree size, eliminating obvious bad moves: 10⁴⁰

How to deal with large games?

Pascal Bercher

Jniversit

So, what to do for (too) large games?

- Don't compute the entire game tree!
- Stop at certain nodes and estimate their payoff! But how?
 - hand-crafted heuristics



- Title: Artificial Intelligence: A Modern Approach (3rd Ed.)
- Authors: Stuart Russel and Peter Norvig

Solving Large Games

URL: https://aima.cs.berkeley.edu/

Estimate per piece:

- pawn: 1 pt
- knight/bishop: 3 pts
- rook: 5 pts

queen: 9 pts



How to deal with large games?

Pascal Bercher

Australiar National

University

27.32

29.32

So, what to do for (too) large games?

- Don't compute the entire game tree!
- Stop at certain nodes and estimate their payoff! But how?
 - hand-crafted heuristics
 - learned heuristics

Machine learning techniques are often used to find a good static evaluation function based on a linear combination of features:

Solving Large Games

$$\hat{v}(s) = w_1 f_1(s) + \cdots + w_n f_n(s)$$

Note the similarity to chess!

- $w_1 = 1$, $f_1(s) =$ number of pawns in s
- $w_2 = 3$, $f_2(s) =$ number of knights/bishops in s
- ▶ ..



Solving Large Games

- hand-crafted heuristics
- learned heuristics
- simulate a game, use the outcome as estimate

Monte-Carlo Tree Search is a well-known algorithm exploiting this idea. It works in four phases:

- Selection (select a non-terminal leaf based on current strategy)
- Expansion (expand the selected node)
- Simulation (play a random game to the end)
- Backpropagation (use the outcome to update strategy)

Interested? See, e.g.,

https://www.youtube.com/watch?v=UXW2yZndl7U (15:30, lecture by Dr. John Levine from Univ. of Strathclyde)

	Australian National University					
		Pascal Bercher	29.32			

		nes? Solving Small Gam						Game AI Success Story ●○	
	Game Al Success Story								
	Game Al Success Story								
	stralian								
- Aus Nat Uni	ional versity Pasc	al Bercher							31.32

- In standard MiniMax or *alpha/beta* pruning, we make a **terminal** test to obtain the payoff, or continue expanding. With heuristics, we instead make a cut-off test to check whether we should stop expansion and *estimate* the payoff of the current node.
- What about using a fixed depth as cut-off test? → Suffers from the horizon problem:



Title: Artificial Intelligence: A Modern Approach (3rd Ed.) Authors: Stuart Russel and Peter Norvig URL: https://aima.cs.berkeley.edu/ White can promote a pawn into a queen

Solving Large Games

White can promote a pawn into a queen on its next move! So the cut-off test should be negative in this state.

Australian National University Pascal Bercher 30.32

0000 000000000 000 0

Mile Stones in AI Game Playing

- 1959 Arthur Samuel develops Checkers playing program
- 1997 IBM's Deep Blue chess machine beats Garry Kasparov
- 2007 Checkers solved by University of Alberta
- 2011 IBM's Watson wins Jeopardy! requiring natural language understanding
- 2015 Deep reinforcement learning algorithms learn to play Atari arcade games from scratch
- 2016 Google DeepMind's AlphaGo beats Lee Sedol, Korea
- 2017 AlphaZero learns Go, Chess, and Shogi from scratch (and beats AlphaGo)

Pascal Bercher



Game AI Success Story

- Australian National University

Sources / Copyright

Australiar Vational **Jniversit**

Pascal Bercher

Sources / Copyright 0000



Taken from https://www.brettspiele-report.de/cublino/ (We argue for fair use.)



Considered fair use for and by Wikipedia (https://en.wikipedia.org/ wiki/Blokus). We also consider it fair dealing (Australian equivalent to US's fair use) for illustrating the game in this lecture.



Picture is *public domain*

https://www.publicdomainpictures.net/en/view-image.php? image=163476&picture=finished-go-game



By Stuart Russel and Peter Norvig from their book Artificial Intelligence: A Modern Approach (3rd Ed.) Availble freely for teaching on https://aima.cs.berkeley.edu/



Picture is public domain https://commons.wikimedia.org/wiki/File:Puissance4_01.svg





Photo by Emile Perron on Unsplash https://unsplash.com/photos/_jXn-gNzuGo

Photo by A. Yobi Blumberg on Unsplash



Photo by Michał Parzuchowski on Unsplash https://unsplash.com/photos/oT-XbATcoTQ



Photo by Macau Photo Agency on Unsplash https://unsplash.com/photos/as5EWdBWKqk Unsplash pictures are free to use, see https://unsplash.com/license

Australiar National University Pascal Bercher

Sources / Copyright 0000

1.4

3.4





Picture is public domain https://nl.wikipedia.org/wiki/Checkers



Taken from the YouTube video https://www.youtube.com/watch?v= 67w0QCMr9EY by channel wikipedia tts, licensed under CC BY (https: //creativecommons.org/licenses/by/3.0/legalcode).



Considered fair use for and by Wikipedia (https://en.wikipedia.org/ wiki/Space_Invaders). We also consider it fair dealing (Australian equivalent to US's fair use) for illustrating the game in this lecture.

