

Games – adversarial search

In which we examine the problems that arise when we try to plan ahead in a world where other agents are planning against us

Piotr Fulmański

Institut Nauk Ekonomicznych i Informatyki,
Państwowa Wyższa Szkoła Zawodowa w Płocku, Polska

December 9, 2015

- 1 Introduction to lecture
- 2 Game in AI
- 3 Optimal decisions in games – minimax
- 4 Optimal decisions in games – alpha-beta cutoff
- 5 Imperfect real-time decisions
- 6 Stochastic games
- 7 Partially observable games

Competitive environments – goals in conflict

- In this part we cover competitive environments, in which the agents' **goals are in conflict**, giving rise to adversarial search problems often known as games.
- We begin with a definition of the **optimal move** and an algorithm for finding it.
- We then look at techniques for choosing a good move when time is limited.
- We also discuss games that include an element of chance.
- We also discuss games that include elements of imperfect information.

Competitive environments – goals in conflict

- In this part we cover competitive environments, in which the agents' **goals are in conflict**, giving rise to adversarial search problems often known as games.
- We begin with a definition of the **optimal move** and an algorithm for finding it.
- We then look at techniques for choosing a good move when time is limited.
- We also discuss games that include an element of chance.
- We also discuss games that include elements of imperfect information.

Competitive environments – goals in conflict

- In this part we cover competitive environments, in which the agents' **goals are in conflict**, giving rise to adversarial search problems often known as games.
- We begin with a definition of the **optimal move** and an algorithm for finding it.
- We then look at techniques for choosing a good move when time is limited.
- We also discuss games that include an element of chance.
- We also discuss games that include elements of imperfect information.

Competitive environments – goals in conflict

- In this part we cover competitive environments, in which the agents' **goals are in conflict**, giving rise to adversarial search problems often known as games.
- We begin with a definition of the **optimal move** and an algorithm for finding it.
- We then look at techniques for choosing a good move when time is limited.
- We also discuss games that include an element of chance.
- We also discuss games that include an element of imperfect information.

Competitive environments – goals in conflict

- In this part we cover competitive environments, in which the agents' **goals are in conflict**, giving rise to adversarial search problems often known as games.
- We begin with a definition of the **optimal move** and an algorithm for finding it.
- We then look at techniques for choosing a good move when time is limited.
- We also discuss games that include an element of chance.
- We also discuss games that include elements of imperfect information.

Competitive environments – goals in conflict

- In this part we cover competitive environments, in which the agents' **goals are in conflict**, giving rise to adversarial search problems often known as games.
- We begin with a definition of the **optimal move** and an algorithm for finding it.
- We then look at techniques for choosing a good move when time is limited.
- We also discuss games that include an element of chance.
- We also discuss games that include elements of imperfect information.

AI games

In AI, the most common games are of a rather specialized kind, what game theorist call deterministic, turn-taking, two-player, zero-sum.

In our terminology, this means

- deterministic,
- fully observable,
- agents act alternately,
- utility values at the end of the game are always opposite.

AI games

In AI, the most common games are of a rather specialized kind, what game theorists call deterministic, turn-taking, two-player, zero-sum.

In our terminology, this means

- deterministic,
- fully observable,
- agents act alternately,
- utility values at the end of the game are always opposite.

AI games

In AI, the most common games are of a rather specialized kind, what game theorists call deterministic, turn-taking, two-player, zero-sum.

In our terminology, this means

- deterministic,
- fully observable,
- agents act alternately,
- utility values at the end of the game are always opposite.

AI games

In AI, the most common games are of a rather specialized kind, what game theorists call deterministic, turn-taking, two-player, zero-sum.

In our terminology, this means

- deterministic,
- fully observable,
- agents act alternately,
- utility values at the end of the game are always opposite.

AI games

In AI, the most common games are of a rather specialized kind, what game theorists call deterministic, turn-taking, two-player, zero-sum.

In our terminology, this means

- deterministic,
- fully observable,
- agents act alternately,
- utility values at the end of the game are always opposite.

AI games

In AI, the most common games are of a rather specialized kind, what game theorists call deterministic, turn-taking, two-player, zero-sum.

In our terminology, this means

- deterministic,
- fully observable,
- agents act alternately,
- utility values at the end of the game are always opposite.

AI games

In AI, the most common games are of a rather specialized kind, what game theorists call deterministic, turn-taking, two-player, zero-sum.

In our terminology, this means

- deterministic,
- fully observable,
- agents act alternately,
- utility values at the end of the game are always opposite.

AI games

In AI, the most common games are of a rather specialized kind, what game theorists call deterministic, turn-taking, two-player, zero-sum.

In our terminology, this means

- deterministic,
- fully observable,
- agents act alternately,
- utility values at the end of the game are always opposite.

AI games

In AI, the most common games are of a rather specialized kind, what game theorists call deterministic, turn-taking, two-player, zero-sum.

In our terminology, this means

- deterministic,
- fully observable,
- agents act alternately,
- utility values at the end of the game are always opposite.

AI games

In AI, the most common games are of a rather specialized kind, what game theorists call deterministic, turn-taking, two-player, zero-sum.

In our terminology, this means

- deterministic,
- fully observable,
- agents act alternately,
- utility values at the end of the game are always opposite.

Why games?

Games are interesting because they are too hard to solve.

Example: chess

- An average branching factor of about 35.
- Games often go to 50 moves.
- So the search tree has about 35^{100} or 10^{154} nodes^a.
- Fortunately the search graph has „only“ about 10^{40} distinct nodes.

^aThe number of atoms in the entire observable universe is estimated to be within the range of 10^{78} to 10^{82} (see

<http://www.universetoday.com/36302/atoms-in-the-universe/>,

http://en.wikipedia.org/wiki/Observable_universe).

Why games?

Games are interesting because they are too hard to solve.

Example: chess

- An average branching factor of about 35.
- Games often go to 50 moves.
- So the search tree has about 35^{100} or 10^{154} nodes^a.
- Fortunately the search graph has „only” about 10^{40} distinct nodes.

^aThe number of atoms in the entire observable universe is estimated to be within the range of 10^{78} to 10^{82} (see

<http://www.universetoday.com/36302/atoms-in-the-universe/>,

http://en.wikipedia.org/wiki/Observable_universe).

Why games?

Games are interesting because they are too hard to solve.

Example: chess

- An average branching factor of about 35.
- Games often go to 50 moves.
- So the search tree has about 35^{100} or 10^{154} nodes^a.
- Fortunately the search graph has „only” about 10^{40} distinct nodes.

^aThe number of atoms in the entire observable universe is estimated to be within the range of 10^{78} to 10^{82} (see

<http://www.universetoday.com/36302/atoms-in-the-universe/>,

http://en.wikipedia.org/wiki/Observable_universe).

Why games?

Games are interesting because they are too hard to solve.

Example: chess

- An average branching factor of about 35.
- Games often go to 50 moves.
- So the search tree has about 35^{100} or 10^{154} nodes^a.
- Fortunately the search graph has „only” about 10^{40} distinct nodes.

^aThe number of atoms in the entire observable universe is estimated to be within the range of 10^{78} to 10^{82} (see

<http://www.universetoday.com/36302/atoms-in-the-universe/>,

http://en.wikipedia.org/wiki/Observable_universe).

Why games?

Games are interesting because they are too hard to solve.

Example: chess

- An average branching factor of about 35.
- Games often go to 50 moves.
- So the search tree has about 35^{100} or 10^{154} nodes^a.
- Fortunately the search graph has „only” about 10^{40} distinct nodes.

^aThe number of atoms in the entire observable universe is estimated to be within the range of 10^{78} to 10^{82} (see

<http://www.universetoday.com/36302/atoms-in-the-universe/>,

http://en.wikipedia.org/wiki/Observable_universe).

Why games?

Games are interesting because they are too hard to solve.

Example: chess

- An average branching factor of about 35.
- Games often go to 50 moves.
- So the search tree has about 35^{100} or 10^{154} nodes^a.
- Fortunately the search graph has „only” about 10^{40} distinct nodes.

^aThe number of atoms in the entire observable universe is estimated to be within the range of 10^{78} to 10^{82} (see

<http://www.universetoday.com/36302/atoms-in-the-universe/>,

http://en.wikipedia.org/wiki/Observable_universe).

Why games?

Games, like the real world, therefore require **the ability to make *some* decision even when calculating the *optimal* decision is infeasible.** Games also penalize inefficiency severely.

Definition

A game can be formally defined as a kind of search problem with the following elements

- The **initial state**, which specifies how the game is set up at the start.
- Function $\text{player}(s)$, which defines which player has the move in a state s .
- Function $\text{actions}(s)$, which returns the set of legal moves in a state.
- The **transition model**, $\text{result}(s,a)$, which defines the result of a move.
- A **terminal test**, $\text{terminalTest}(s)$, which is true when the game is over and false otherwise (equivalent of a goal test function).
- Function $\text{payoff}(s,p)$, also called an **objective function** or **utility function**, which returns final numeric value for a game that ends in terminal state s for a player p (equivalent of a cost function).

Definition

A game can be formally defined as a kind of search problem with the following elements

- The **initial state**, which specifies how the game is set up at the start.
- Function $\text{player}(s)$, which defines which player has the move in a state s .
- Function $\text{actions}(s)$, which returns the set of legal moves in a state.
- The **transition model**, $\text{result}(s,a)$, which defines the result of a move.
- A **terminal test**, $\text{terminalTest}(s)$, which is true when the game is over and false otherwise (equivalent of a goal test function).
- Function $\text{payoff}(s,p)$, also called an **objective function** or **utility function**, which returns final numeric value for a game that ends in terminal state s for a player p (equivalent of a cost function).

Definition

A game can be formally defined as a kind of search problem with the following elements

- The **initial state**, which specifies how the game is set up at the start.
- Function $\text{player}(s)$, which defines which player has the move in a state s .
- Function $\text{actions}(s)$, which returns the set of legal moves in a state.
- The **transition model**, $\text{result}(s,a)$, which defines the result of a move.
- A **terminal test**, $\text{terminalTest}(s)$, which is true when the game is over and false otherwise (equivalent of a goal test function).
- Function $\text{payoff}(s,p)$, also called an **objective function** or **utility function**, which returns final numeric value for a game that ends in terminal state s for a player p (equivalent of a cost function).

Definition

A game can be formally defined as a kind of search problem with the following elements

- The **initial state**, which specifies how the game is set up at the start.
- Function $\text{player}(s)$, which defines which player has the move in a state s .
- Function $\text{actions}(s)$, which returns the set of legal moves in a state.
- The **transition model**, $\text{result}(s,a)$, which defines the result of a move.
- A **terminal test**, $\text{terminalTest}(s)$, which is true when the game is over and false otherwise (equivalent of a goal test function).
- Function $\text{payoff}(s,p)$, also called an **objective function** or **utility function**, which returns final numeric value for a game that ends in terminal state s for a player p (equivalent of a cost function).

Definition

A game can be formally defined as a kind of search problem with the following elements

- The **initial state**, which specifies how the game is set up at the start.
- Function $\text{player}(s)$, which defines which player has the move in a state s .
- Function $\text{actions}(s)$, which returns the set of legal moves in a state.
- The **transition model**, $\text{result}(s,a)$, which defines the result of a move.
- A **terminal test**, $\text{terminalTest}(s)$, which is true when the game is over and false otherwise (equivalent of a goal test function).
- Function $\text{payoff}(s,p)$, also called an **objective function** or **utility function**, which returns final numeric value for a game that ends in terminal state s for a player p (equivalent of a cost function).

Definition

A game can be formally defined as a kind of search problem with the following elements

- The **initial state**, which specifies how the game is set up at the start.
- Function $\text{player}(s)$, which defines which player has the move in a state s .
- Function $\text{actions}(s)$, which returns the set of legal moves in a state.
- The **transition model**, $\text{result}(s,a)$, which defines the result of a move.
- A **terminal test**, $\text{terminalTest}(s)$, which is true when the game is over and false otherwise (equivalent of a goal test function).
- Function $\text{payoff}(s,p)$, also called an **objective function** or **utility function**, which returns final numeric value for a game that ends in terminal state s for a player p (equivalent of a cost function).

Definition

A game can be formally defined as a kind of search problem with the following elements

- The **initial state**, which specifies how the game is set up at the start.
- Function $\text{player}(s)$, which defines which player has the move in a state s .
- Function $\text{actions}(s)$, which returns the set of legal moves in a state.
- The **transition model**, $\text{result}(s,a)$, which defines the result of a move.
- A **terminal test**, $\text{terminalTest}(s)$, which is true when the game is over and false otherwise (equivalent of a goal test function).
- Function $\text{payoff}(s,p)$, also called an **objective function** or **utility function**, which returns final numeric value for a game that ends in terminal state s for a player p (equivalent of a cost function).

Definition

The initial state, actions and results define the **game tree** for the game – a game where

- the nodes are game states and
- the edges are moves.

Example

Game tree for the game of tic-tac-toe.

Definition

The initial state, actions and results define the **game tree** for the game – a game where

- the nodes are game states and
- the edges are moves.

Example

Game tree for the game of tic-tac-toe.

What is an optimal solution in games?

- In a normal search problem, the optimal solution would be a **sequence of actions** leading to a goal state.
- In adversarial search, opponent has something to say about it.
- Therefore we must find a **contingent strategy**, which specifies our move in the initial state and then moves in the states resulting from every possible response by our opponent.

What is an optimal solution in games?

- In a normal search problem, the optimal solution would be a **sequence of actions** leading to a goal state.
- In adversarial search, opponent has something to say about it.
- Therefore we must find a **contingent strategy**, which specifies our move in the initial state and then moves in the states resulting from every possible response by our opponent.

What is an optimal solution in games?

- In a normal search problem, the optimal solution would be a **sequence of actions** leading to a goal state.
- In adversarial search, opponent has something to say about it.
- Therefore we must find a **contingent strategy**, which specifies our move in the initial state and then moves in the states resulting from every possible response by our opponent.

What is an optimal solution in games?

- In a normal search problem, the optimal solution would be a **sequence of actions** leading to a goal state.
- In adversarial search, opponent has something to say about it.
- Therefore we must find a **contingent strategy**, which specifies our move in the initial state and then moves in the states resulting from every possible response by our opponent.

How to find optimal strategy

- One ply tree.
- Expansion of the previous tree: two ply (one move deep).
- Example with two moves deep tree.

How to find optimal strategy

- One ply tree.
- Expansion of the previous tree: two ply (one move deep).
- Example with two moves deep tree.

How to find optimal strategy

- One ply tree.
- Expansion of the previous tree: two ply (one move deep).
- Example with two moves deep tree.

How to find optimal strategy

- One ply tree.
- Expansion of the previous tree: two ply (one move deep).
- Example with two moves deep tree.

Minimax algorithm

```
function miniMax(state) return an action
{
  v := max(state)
  return the action in actions(state) with value v
}
```

```
function max(state) return a payoff value
{
  if(terminalTest(state)) then
    return payoff(state)
  v := -infty
  for each a in actions(state) do
    v := max(v, min(result(state,a)))
  return v
}
```

```
function min(state) return a payoff value
{
  if(terminalTest(state)) then
    return payoff(state)
  v := +infty
  for each a in actions(state) do
    v := min(v, max(result(state,a)))
  return v
}
```

Optimal decisions in games

Optimal decision in multiplayer games

Let us examine how to extend the minimax idea to multiplayer games.

First, we need to replace the single value for each node with a vector of values. For terminal states, this vector gives the utility of the state from each player's viewpoint.

Optimal decisions in games

Optimal decision in multiplayer games

Let us examine how to extend the minimax idea to multiplayer games.

First, we need to replace the single value for each node with a vector of values. For terminal states, this vector gives the utility of the state from each player's viewpoint.

Optimal decisions in games

Optimal decision in multiplayer games: two players

Tree

```
      *      me (max level: I want to maximize my payoff)
     / \
2 *   * 3   opponent (min level: opponent wants to minimize MY payoff,
 / \ / \   so he/she doesn't care about his/her payoff. This is
4 2 1 3   correct with the assumptions we made at the beginning:
          every time I get a point, my opponent lose it and vice
          versa)
```

has to be replaced by

```
      *      me (max level)
     / \
[2,?] *   * [3,?]  opponent (min level)
    / \ / \ |
    /  | | \ +-+ Payoff for opponent, but I don't care about this.
    /  | | \
[4,?][2,?][1,?][3,?]
```

Optimal decisions in games

Optimal decision in multiplayer games

Optimal decision in multiplayer games

Example with one move deep tree for three players.

We can create game tree for three players follow the idea from previous slide: in every node we have a payoff vector, but we care only about our payoff values and try to maximize it while our opponents want to minimize it. In consequence every level corresponding to opponents choice is a MIN level, while every level corresponding to our choice is a MAX level.

Optimal decisions in games

Optimal decision in multiplayer games

The minimax algorithm performs a complete depth-first exploration of the game tree. This is a real problem, because the number of game states it has to examine is exponential in the depth of the tree. Unfortunately, we can't eliminate the exponent, but it turns out we can effectively cut it in half. **The trick is that it is possible to compute the correct minimax decision without looking at every node in the game tree.** The technique we examine is called **alpha-beta cutoff** (or alpha-beta pruning). When applied to a standard minimax tree, it returns the same move as minimax would, but **cutoff away branches that cannot possibly influence the final decision.**

alpha-beta cutoff – intuition

Intuition (see separate file)

Optimal decisions in games

Definition of alpha and beta parameter

- α the value of the best choice we have found so far at any choice point along the path for MAX (so, α is the highest value so far).
- β the value of the best choice we have found so far at any choice point along the path for MIN (so, β is the lowest value so far).

Optimal decisions in games

alpha-beta cutoff with alpha and beta parameter

The most important to remember when working with alpha-beta is:

- init α as $-\infty$
- init β as $+\infty$
- every time in MAX node we update α with the formula
 $\alpha := \text{maximum}(\alpha, v)$ ¹
- every time in MIN node we update β with the formula
 $\beta := \text{minimum}(\beta, v)$
- we cut whenever $\alpha \geq \beta$

Example with alpha and beta (see separate file).

¹ v is the payoff for current node and some action.

alpha-beta cutoff and move ordering

The effectiveness of alpha-beta pruning is highly dependent on the order in which the states are examined (see separate file).

Optimal decisions in games

alpha-beta cutoff and move ordering

The effectiveness of alpha-beta pruning is highly dependent on the order in which the states are examined (see separate file).

Minimax with alpha-beta cutoff algorithm

```
function minMaxAlphaBetaCutoff(state) return an action{
  v := max(state, -infty, +infty)
  return the action in actions(state) with value v
}
```

```
function max(state, alpha, beta) return a utility value{
  if (terminalTest(state)) then
    return utility(state)
  v := -infty
  for each a in actions(state) do
    v := maximum(v, min(result(state, a), alpha, beta))
    alpha := maximum(alpha, v)

    if (alpha >= beta) then //cutoff
      return v

  return v
}
```

```
function min(state, alpha, beta) return a utility value{
  if (terminalTest(state)) then
    return utility(state)
  v := +infty
  for each a in actions(state) do
    v := minimum(v, max(result(state, a), alpha, beta))
    beta := minimum(beta, v)

    if (alpha >= beta) then //cutoff
      return v
```

Imperfect real-time decisions

Imperfect real-time decisions

- **Evaluation function.** This function returns an **estimate** of the expected utility of the game from a given position, just as the heuristic functions in previous part of this lecture.
- **Cutting off search** (depth limit). Replace
`if(terminalTest(state)) then return utility(state)`
with
`if(cutoffTest(state,depth)) then return eval(state)`
- **Forward pruning** – some moves at a given node are pruned immediately without further consideration.
- **Search versus lookup.**

Imperfect real-time decisions

Imperfect real-time decisions

- **Evaluation function.** This function returns an **estimate** of the expected utility of the game from a given position, just as the heuristic functions in previous part of this lecture.
- **Cutting off search** (depth limit). Replace
`if(terminalTest(state)) then return utility(state)`
with
`if(cutoffTest(state,depth)) then return eval(state)`
- **Forward pruning** – some moves at a given node are pruned immediately without further consideration.
- **Search versus lookup.**

Imperfect real-time decisions

Imperfect real-time decisions

- **Evaluation function.** This function returns an **estimate** of the expected utility of the game from a given position, just as the heuristic functions in previous part of this lecture.
- **Cutting off search** (depth limit). Replace
`if(terminalTest(state)) then return utility(state)`
with
`if(cutoffTest(state,depth)) then return eval(state)`
- **Forward pruning** – some moves at a given node are pruned immediately without further consideration.
- **Search versus lookup.**

Imperfect real-time decisions

Imperfect real-time decisions

- **Evaluation function.** This function returns an **estimate** of the expected utility of the game from a given position, just as the heuristic functions in previous part of this lecture.
- **Cutting off search** (depth limit). Replace
`if(terminalTest(state)) then return utility(state)`
with
`if(cutoffTest(state,depth)) then return eval(state)`
- **Forward pruning** – some moves at a given node are pruned immediately without further consideration.
- Search versus lookup.

Imperfect real-time decisions

Imperfect real-time decisions

- **Evaluation function.** This function returns an **estimate** of the expected utility of the game from a given position, just as the heuristic functions in previous part of this lecture.
- **Cutting off search** (depth limit). Replace
`if(terminalTest(state)) then return utility(state)`
with
`if(cutoffTest(state,depth)) then return eval(state)`
- **Forward pruning** – some moves at a given node are pruned immediately without further consideration.
- **Search versus lookup.**

Backgammon and chance nodes

Although one player (say white) knows his own legal moves are, he does not know what opponent is going to roll and thus does not know what opponent's legal moves will be. That means white cannot construct a standard game tree of the sort we saw in chess or tic-tac-toe. A game tree in backgammon must include **chance nodes** in addition to MAX and MIN nodes^a. Terminal nodes and MAX and MIN nodes work exactly the same way as before (because the dice roll is known). **For chance nodes, we compute the expected value, which is the sum of the value over all outcomes, weighted by the probability of each chance action.** This leads to generalize minimax for deterministic games to an **expectiminimax** for game with chance node.

^aSo we have nodes as follow: MAX — CHANCE — MIN — CHANCE — MAX —

...

Partially observable games

Partially observable games – battleships

In deterministic partially observable games, uncertainty about the state of the board arises entirely from the lack of access to the choices made by the opponent. This class includes games such as Battleships.

Partially observable games

Partially observable games – card games

Card games provide many examples of stochastic partial observability, where the missing information is generated randomly. For example, in many games, cards are dealt randomly at the beginning of the game, with each player receiving a hand that is not visible to the other players.

At first sight, it might seem that these card games are just like dice games: the cards are dealt randomly and determine the moves available to each player, but all the „dice” are rolled at the beginning.

Even though this analogy turns out to be incorrect, it suggests an effective algorithm: consider all possible deals of the invisible cards; solve each one as if it were a fully observable game; and then choose the move that has the best outcome averaged over all the deals.

Partially observable games

Partially observable games – card games

Card games provide many examples of stochastic partial observability, where the missing information is generated randomly. For example, in many games, cards are dealt randomly at the beginning of the game, with each player receiving a hand that is not visible to the other players.

At first sight, it might seem that these card games are just like dice games: the cards are dealt randomly and determine the moves available to each player, but all the „dice” are rolled at the beginning.

Even though this analogy turns out to be incorrect, it suggests an effective algorithm: consider all possible deals of the invisible cards; solve each one as if it were a fully observable game; and then choose the move that has the best outcome averaged over all the deals.

Partially observable games

Partially observable games – card games

Card games provide many examples of stochastic partial observability, where the missing information is generated randomly. For example, in many games, cards are dealt randomly at the beginning of the game, with each player receiving a hand that is not visible to the other players. At first sight, it might seem that these card games are just like dice games: the cards are dealt randomly and determine the moves available to each player, but all the „dice” are rolled at the beginning.

Even though this analogy turns out to be incorrect, it suggests an effective algorithm: consider all possible deals of the invisible cards; solve each one as if it were a fully observable game; and then choose the move that has the best outcome averaged over all the deals.