


Games II: The Minimax Algorithm

15-150 M21

Lecture 0726
26 July 2021




- Implemented playable games in SML
- Our game implementation consisted of:
 - ▶ A `GAME` (specifying rules, how to make moves, etc.)
 - ▶ `PLAYERs` (plays a particular `GAME`, provides function `next_move` assigning a “choice” of move to each state)
 - ▶ Games are refereed by a `CONTROLLER`, who facilitates play between two `PLAYERs` playing the same game.
- Implemented `Nim`, where states were of the form `(s, Minnie)` or `(s, Maxie)` for `s: int` nonnegative. A move is a positive `int` `i` which is less than or equal to `Int.min(3, s)`.

We'll deal with 4 different kinds of players:

- Human players (our game library includes utilities to accept user input to determine `next_move`)
- Directly-implemented players (`NimPlayer` from tomorrow's lab)
- MiniMax players (this lecture)
- Alphabeta players (Wednesday's lecture, `games` homework)

0 How to Build Smart PLAYERS

What do we mean by smart?

 **dis·cern·ing**
/dəˈsɜrnɪŋ/
adjective
having or showing good judgment.
"the restaurant attracts discerning customers"

Similar: discriminating selective judicious tasteful refined cultivated

We want to design our PLAYERS such that their `next_move` function makes decisions which generally lead to it winning the game more often.


Demonstration:

RunNim.play RunNim.HvM

vs.

RunNim.play RunNim.HvP

What do we mean by smart?

 **dis·cern·ing**
/dəˈsɜrnɪŋ/
adjective

having or showing good judgment.
"the restaurant attracts discerning customers"

Similar: discriminating selective judicious tasteful refined cultivated ▼

We want to design our `PLAYERs` such that their `next_move` function makes decisions which generally lead to it winning the game more often.


So what we want to do is build a player who “knows what’s good for her”: who is able to assess the moves available to her, decide which one has the most favorable outcome, and make the corresponding move her `next_move`.

Formally, we make sense of games mathematically by examining the corresponding *game tree*. A game tree is a finitely-branching tree where

- The nodes represent *game states*
- The edges represent *moves*
- The root node is the current state of the game, and the rest of the tree represents different outcomes achievable by a certain series of moves from the two players
- The children of a given node are the states reachable from that game state by the current player making a valid move.

We'll call our players '**Maxie**' and '**Minnie**'.

**I want to
Maximize!**



MAXIE



MINNIE

**I want to
Minimize!**

Demonstration: Nim Game Tree

Observation:

**A good player is thinking a few
moves into the future**

Problem:

It's not always tractable to search through the *entire* game tree

- Problem: it's impractical (and often impossible) to visit every node of the tree
- Solution: explore some of the tree, and guess
 - ▶ Have a fixed 'search depth' d
 - ▶ Explore the top d levels of the tree (i.e. the game states than can be reached from the current one in d moves or fewer)
 - ▶ When you hit your search depth, use your knowledge of the game to assign an appropriate value to that state, and treat that value as the value of the node.
- More precisely: we'll have a function `estimate` which takes a game state (for instance, a value of type `Nim.State.t`) and returns a "guess" of the goodness or badness of that state.

An *estimator* for a game G is a function assigning “guesses” to each state to (perhaps roughly) indicate who’s winning.

- The “guesses” will usually be numerical (e.g. `ints`): lower numbers better for **Minnie**, larger numbers better for **Maxie**. The scale is arbitrary: all that matters is the relative ordering of states.
- The goal here is to induce an ordering on states, i.e. articulate a sense in which states are “better” or “worse” than each other (from one player’s perspective).
- We want “better” to mean “more likely to win” (as best as possible)
- A given GAME will have many possible estimators, with varying degrees of sophistication, and which may weight different factors differently.

0726.0 (lib/game/estimate/ESTIMATOR.sig)

```
2 signature ESTIMATOR =
3 sig
4   structure Game : GAME
5
6   type guess
7   datatype est = Definitely of Game.Outcome.t
8                 | Guess of guess
9   val compare : est * est -> order
10  val toString : guess -> string
11
12  val estimate : Game.State.t -> guess
```

- Note that the only operation on values of type `est` is comparison (the function `compare`). We don't – in general – require `guesses` to be numbers at all, we just require that they be ordered.
- We **transparently** ascribe to this signature. While we don't require in general that `guess` is implemented as `int` or `real`, if we do happen to implement it that way we want to have access to the associated methods (e.g. from the basis structures `Int` and `Real`).

Nim has a perfect estimator


Player p can guarantee a win from $(s, \text{flip } p)$

iff

$$s \bmod 4 \cong 1$$

(remember

```
fun flip Maxie = Minnie | flip Minnie = Maxie)
```



Nim has a perfect estimator

So, assuming the other player plays optimally, whoever's turn it is when s is of the form $(4*k)+1$ for some $k : \text{int}$ will *lose*.

```
(* recall a value of Nim.State.t is (s,p)
   for some nonnegative int s and either
   p=Minnie or p=Maxie *)
(* estimate : Nim.State.t -> int *)
fun estimate (s,p) =
  case (s mod 4, p) of
    (1,Minnie) => 1
  | (1,Maxie) => ~1
  | (_,Minnie) => ~1
  | (_,Maxie) => 1
```



how to build an ESTIMATOR



This is somewhat *too* clean of an example: most games don't have perfect estimators. Rather, the best we can do is make pretty good guesses! To design an estimator, we'll usually use some combination of simple heuristics and more sophisticated theory.

For instance, here's a common heuristic for chess: for a chess piece p , let $v(p)$ be given by the following chart

Symbol					
Piece	pawn	knight	bishop	rook	queen
Value	1	3	3	5	9

Then put

$$\text{estimate}(S) = \left(\sum_{\substack{\text{Pieces } p \text{ Maxie} \\ \text{has in play (in } S)}} v(p) \right) - \left(\sum_{\substack{\text{Pieces } p \text{ Minnie} \\ \text{has in play (in } S)}} v(p) \right)$$

1 The MiniMax Algorithm

- We should assign each node an estimator guess, its “value” .
- The value of a node should reflect who’s winning from that node, which depends on the moves available from that state.
- From there, we can fill in the rest of the game tree by assuming the players play optimally

Demonstration: Minimax

Fix a search depth d .

- 1 Traverse the game tree down to the d -th level. (For every node encountered where the game is over, assign such nodes the value `Definitely` of whoever the winner is.)
- 2 Call the estimator to assign values to the d -th level.
- 3 Work upwards, assigning values to nodes according to the Minnie and Maxie principles described above
 - ▶ For **Minnie** nodes: the value should be the *minimum* of the values of the child nodes
 - ▶ For **Maxie** nodes: the value should be *maximum* of the values of the child nodes.

Once we've filled all the way to the top of the tree (our current state), then we can decide which move to make based on the estimated values.

2 SML Implementation

0726.3 (lib/game/estimate/MiniMax.fun)

```
2 functor MiniMax (Settings : SETTINGS) :> PLAYER
```

0726.1 (lib/game/estimate/SETTINGS.sig)

```
2 signature SETTINGS =  
3 sig  
4  
5     structure Est : ESTIMATOR  
6  
7     val search_depth : int  
8  
9 end
```

0726.2 (lib/game/core/PLAYER.sig)

```
2 signature PLAYER =  
3 sig  
4  
5     structure Game : GAME  
6  
7     val next_move : Game.State.t -> Game.Move.t  
8  
9 end
```

0726.3 (lib/game/estimate/MiniMax.fun)

```
2 functor MiniMax (Settings : SETTINGS) :> PLAYER
3 where Game = Settings.Est.Game =
4 struct
5     structure Est = Settings.Est
6     structure Game = Est.Game
```

0726.4 (lib/game/estimate/MiniMax.fun)

```
9  type edge = Game.Move.t * Est.est
10 fun valueOf ((_, value) : edge) = value
11 fun moveOf ((move, _) : edge) = move
12 fun max ((m1, v1):edge, (m2, v2):edge):edge =
13     case Est.compare (v1, v2) of
14         LESS => (m2, v2)
15     | _ => (m1, v1)
16 fun min ((m1, v1):edge, (m2, v2):edge):edge =
17     case Est.compare (v1, v2) of
18         GREATER => (m2, v2)
19     | _ => (m1, v1)
```

Encoding the difference between Minnie and Maxie

```
reduce1 : ('a * 'a -> 'a) -> 'a Seq.seq -> 'a
REQUIRES: g is total and associative, S is nonempty
ENSURES:
reduce1 g ⟨x1, ..., xn⟩ ≅ g(x1, g(x2, g(..., , xn)))
```


0726.5 (lib/game/estimate/MiniMax.fun)

```
22 (* choose : Player.t -> edge Seq.seq -> edge
23 *)
24 fun choose Player.Maxie = Seq.reduce1 max
    | choose Player.Minnie = Seq.reduce1 min
```



```
fun even 0 = true
  | even n = odd (n-1)
and odd 0 = false
  | odd n = even (n-1)
```

Maxie



0726.6 (lib/game/estimate/MiniMax.fun)

```
27 (* search : int -> G.State.t -> edge *)
28 (* REQUIRES: d > 0 *)
29 fun search (d:int) (s:Game.State.t) : edge =
30     choose (Game.player s)
31     (
32     Seq.map
33         (fn m =>
34             (m,
35              evaluate (d-1) (Game.play(s,m)))
36             )
37         (Game.moves s)
38     )
```

0726.7 (lib/game/estimate/MiniMax.fun)

```
41 (* evaluate:int -> Game.status -> Est.est *)
42 (* REQUIRES: d >= 0 *)
43 and evaluate (d : int) (st : Game.status) :
Est.est =
44   case st of
45     Game.Playing s => (
46       case d of
47         0 => Est.Guess (Est.estimate s)
48         | _ => valueOf (search d s)
49       )
50     | Game.Done oc => Est.Definitely oc
```

0726.8 (lib/game/estimate/MiniMax.fun)

```
1  val next_move =  
2    moveOf o search Settings.search_depth
```

Tomorrow's Lab:

Estimators & Minimax

Optimizing the complexity of minimax, to avoid unnecessary work

Thank you!