

## Artificial Intelligence I

*Dr Sean Holden*

Notes on *games* (*adversarial search*)

Copyright © Sean Holden 2002-2010.

## Solving problems by search: playing games

How might an agent act when *the outcomes of its actions are not known* because an *adversary is trying to hinder it*?

- This is essentially a more realistic kind of search problem because we do not know the exact outcome of an action.
- This is a common situation when *playing games*: in chess, draughts, and so on an opponent *responds* to our moves.
- We don't know what their response will be, and so the outcome of our moves is not clear.

Game playing has been of interest in AI because it provides an *idealisation* of a world in which two agents act to *reduce* each other's well-being.

## Playing games: search against an adversary

Despite the fact that games are an idealisation, game playing can be an excellent source of hard problems. For instance with chess:

- The average branching factor is roughly 35.
- Games can reach 50 moves per player.
- So a rough calculation gives the search tree  $35^{100}$  nodes.
- Even if only different, legal positions are considered it's about  $10^{40}$ .

*So: in addition* to the uncertainty due to the opponent:

- We can't make a complete search to find the best move...
- ... so we have to act even though we're not sure about the best thing to do.

## Playing games: search against an adversary

And chess isn't even very hard:

- *Go* is *much* harder than chess.
- The branching factor is about 360.

Until very recently it has resisted all attempts to produce a good AI player.

See:

[senseis.xmp.net/?MoGo](http://senseis.xmp.net/?MoGo)

and others.

## Playing games: search against an adversary

It seems that games are a step closer to the complexities inherent in the world around us than are the standard search problems considered so far.

The study of games has led to some of the most celebrated applications and techniques in AI.

We now look at:

- How game-playing can be modelled as *search*.
- The *minimax algorithm* for game-playing.
- Some problems inherent in the use of minimax.
- The concept of  $\alpha - \beta$  *pruning*.

*Reading:* Russell and Norvig chapter 6.

## Perfect decisions in a two-person game

Say we have two players. Traditionally, they are called *Max* and *Min* for reasons that will become clear.

- We'll use *noughts and crosses* as an initial example.
- Max moves first.
- The players alternate until the game ends.
- At the end of the game, prizes are awarded. (Or punishments administered—**EVIL ROBOT** is starting up his favourite chain-saw...)

This is exactly the same game format as chess, Go, draughts and so on.

## Perfect decisions in a two-person game

Games like this can be modelled as search problems as follows:

- There is an *initial state*.




- There is a set of *operators*. Here, Max can place a cross in any empty square, or Min a nought.
- There is a *terminal test*. Here, the game ends when three noughts or three crosses are in a row, or there are no unused spaces.
- There is a *utility* or *payoff* function. This tells us, numerically, what the outcome of the game is.

This is enough to model the entire game.

## Perfect decisions in a two-person game


We can *construct a tree* to represent a game. From the initial state Max can make nine possible moves:



Then it's Min's turn...

### Perfect decisions in a two-person game


For each of Max's opening moves Min has eight replies:



And so on...

This can be continued to represent *all* possibilities for the game.


### Perfect decisions in a two-person game



At the leaves a player has won or there are no spaces. Leaves are *labelled* using the utility function.

### Perfect decisions in a two-person game

How can Max use this tree to decide on a move? Consider a much simpler tree:




If Max is rational he will play to reach a position with the *biggest utility possible*

But if Min is rational she will play to *minimise* the utility available to Max.

### The minimax algorithm

There are two moves: Max then Min. Game theorists would call this one move, or two *ply* deep.

The *minimax algorithm* allows us to infer the best move that the current player can make, given the utility function, by working backward from the leaves.



As Min plays the last move, she *minimises* the utility available to Max.


## The minimax algorithm

Min takes the final move:

- If Min is in game position 1, her best choice is move 3. So from Max's point of view this node has a utility of 2.
- If Min is in game position 2, her best choice is move 3. So from Max's point of view this node has a utility of 6.
- If Min is in game position 3, her best choice is move 1. So from Max's point of view this node has a utility of 1.
- If Min is in game position 4, her best choice is move 4. So from Max's point of view this node has a utility of 4.

## The minimax algorithm

Moving one further step up the tree:



We can see that Max's best opening move is move 2, as this leads to the node with highest utility.

## The minimax algorithm

*In general:*

- Generate the complete tree and label the leaves according to the utility function.
- Working from the leaves of the tree upward, label the nodes depending on whether Max or Min is to move.
- If *Min* is to move label the current node with the *minimum* utility of any descendant.
- If *Max* is to move label the current node with the *maximum* utility of any descendant.

If the game is  $p$  ply and at each point there are  $q$  available moves then this process has (surprise, surprise)  $O(q^p)$  time complexity and space complexity linear in  $p$  and  $q$ .

## Making imperfect decisions

We need to avoid searching all the way to the end of the tree. *So:*

- We generate only part of the tree: instead of testing whether a node is a leaf we introduce a *cut-off* test telling us when to stop.
- Instead of a utility function we introduce an *evaluation function* for the evaluation of positions for an incomplete game.

The evaluation function attempts to measure the expected utility of the current game position.

### Making imperfect decisions

How can this be justified?

- This is a strategy that humans clearly sometimes make use of.
- For example, when using the concept of *material value* in chess.
- The effectiveness of the evaluation function is *critical*...
- ... but it must be computable in a reasonable time.
- (In principle it could just be done using minimax.)

The importance of the evaluation function can not be understated—it is probably the most important part of the design.

### The evaluation function

Designing a good evaluation function can be extremely tricky:

- Let's say we want to design one for chess by giving each piece its material value: pawn = 1, knight/bishop = 3, rook = 5 and so on.
- Define the evaluation of a position to be the difference between the material value of black's and white's pieces

$$\text{eval}(\text{position}) = \sum_{\text{black's pieces } p_i} \text{value of } p_i - \sum_{\text{white's pieces } q_i} \text{value of } q_i$$

This seems like a reasonable first attempt. Why might it go wrong?

### The evaluation function

Consider what happens at the start of a game:

- Until the first capture the evaluation function gives 0, so in fact we have a *category* containing many different game positions with equal estimated utility.
- For example, all positions where white is one pawn ahead.
- The evaluation function for such a category should perhaps represent the probability that a position chosen at random from it leads to a win.

So in fact this seems highly naive...

### The evaluation function

Ideally, we should consider *individual positions*.

If on the basis of past experience a position has 50% chance of winning, 10% chance of losing and 40% chance of reaching a draw, we might give it an evaluation of

$$\text{eval}(\text{position}) = (0.5 \times 1) + (0.1 \times -1) + (0.4 \times 0) = 0.4.$$

Extending this to the evaluation of categories, we should then weight the positions in the category according to their likelihood of occurring.

Of course, we *don't know* what any of these likelihoods are...

### The evaluation function

Using material value can be thought of as giving us a *weighted linear evaluation function*

$$\text{eval}(\text{position}) = \sum_{i=1}^n w_i f_i$$

where the  $w_i$  are *weights* and the  $f_i$  represent *features* of the position.  
In this example

$f_i$  = value of the  $i$ th piece

$w_i$  = number of  $i$ th pieces on the board

where black and white pieces are regarded as different and the  $f_i$  are positive for one and negative for the other.

### The evaluation function

Evaluation functions of this type are very common in game playing.

There is no systematic method for their design.

Weights can be chosen by allowing the game to play itself and using *learning* techniques to adjust the weights to improve performance.

By using more carefully crafted features we can give *different evaluations* to *individual positions*.

### $\alpha - \beta$ pruning


Even with a good evaluation function and cut-off test, the time complexity of the minimax algorithm makes it impossible to write a good chess program without some further improvement.

- Assuming we have 150 seconds to make each move, for chess we would be limited to a search of about 3 to 4 ply whereas...
- ...even an average human player can manage 6 to 8.

Luckily, it is possible to prune the search tree *without affecting the outcome* and *without having to examine all of it*.

### $\alpha - \beta$ pruning


Returning for a moment to the earlier, simplified example:



The search is depth-first and left to right.

### $\alpha - \beta$ pruning

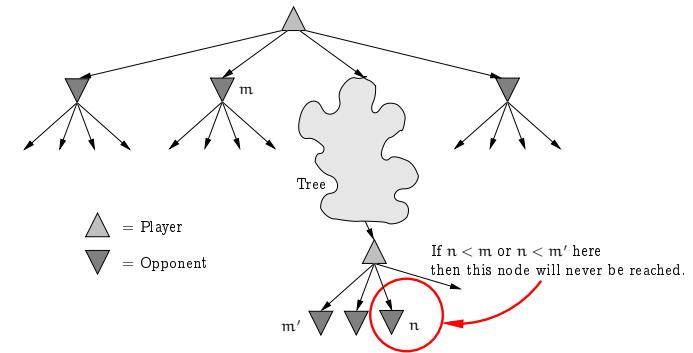
The search continues as previously for the first 8 leaves.



Then we note: if *Max* plays move 3 then *Min* can reach a leaf with utility at most 1.


So: we don't need to search any further under *Max*'s opening move 3. This is because the search has *already established* that *Max* can do better by making opening move 2.

### $\alpha - \beta$ pruning in general



So: once you've established that  $n$  is sufficiently small, you don't need to explore any more of the corresponding node's children.

### $\alpha - \beta$ pruning in general



So: once you've established that  $n$  is sufficiently large, you don't need to explore any more of the corresponding node's children.

### $\alpha - \beta$ pruning in general

The search is depth-first, so we're only ever looking at *one path through the tree*.

We need to keep track of the values  $\alpha$  and  $\beta$  where

$\alpha$  = the *highest* utility seen so far on the path for *Max*

$\beta$  = the *lowest* utility seen so far on the path for *Min*

Assume *Max* begins. Initial values for  $\alpha$  and  $\beta$  are

$$\alpha = -\infty$$

and

$$\beta = +\infty.$$

## $\alpha - \beta$ pruning in general

So: we start with the function call

$\text{max}(-\infty, +\infty, \text{root})$

where max is the function

```

max(alpha, beta, node)
{
  if (node is at cut-off)
    return evaluation(node);
  else
    {
      for (each successor n' of node)
        {
          alpha = maximum(alpha, min(alpha, beta, n'));
          if (alpha >= beta)
            return beta;          // pruning happens here.
        }
      return alpha;
    }
}
    
```

## $\alpha - \beta$ pruning in general


The function min is

```

min(alpha, beta, node)
{
  if (node is at cut-off)
    return evaluation(node);
  else
    {
      for (each successor n' of node)
        {
          beta = minimum(beta, max(alpha, beta, n'));
          if (beta <= alpha)
            return alpha;          // pruning happens here.
        }
      return beta;
    }
}
    
```

## $\alpha - \beta$ pruning in general

Applying this to the earlier example and keeping track of the values for  $\alpha$  and  $\beta$  you should obtain:



## How effective is $\alpha - \beta$ pruning?

(Warning: the theoretical results that follow are somewhat idealised.)

A quick inspection should convince you that the *order* in which moves are arranged in the tree is critical.

So, it seems sensible to try good moves first:

- If you were to have a perfect move-ordering technique then  $\alpha - \beta$  pruning would be  $O(q^{p/2})$  as opposed to  $O(q^p)$ .
- so the branching factor would effectively be  $\sqrt{q}$  instead of  $q$ .
- We would therefore expect to be able to search ahead *twice as many moves as before*.

However, this is not realistic: if you had such an ordering technique you'd be able to play perfect games!



### How effective is $\alpha - \beta$ pruning?

If moves are arranged at random then  $\alpha - \beta$  pruning is:

- $O((q/\log q)^p)$  asymptotically when  $q > 1000$  or...
- ...about  $O(q^{3p/4})$  for reasonable values of  $q$ .

In practice simple ordering techniques can get close to the best case. For example, if we try captures, then threats, then moves forward *etc.*

Alternatively, we can implement an iterative deepening approach and use the order obtained at one iteration to drive the next.

### A further optimisation: the transposition table

Finally, note that many games correspond to *graphs* rather than *trees* because the same state can be arrived at in different ways.

- This is essentially the same effect we saw in heuristic search: recall *graph search* versus *tree search*.
- It can be addressed in a similar way: store a state with its evaluation in a hash table—generally called a *transposition table*—the first time it is seen.

The transposition table is essentially equivalent to the *closed list* introduced as part of graph search.

This can vastly increase the effectiveness of the search process, because we don't have to evaluate a single state multiple times.