Two-player games

www.npr.org

Setup

@Helge: Pressed record?

@Class: Logged into iClicker cloud?

Overview

First half:

- ***Shortest paths cont.***
- ***Two-player games***

… all about traversing trees efficiently

+ Some debugging tips

Second half:

- ***Physical simulation basics***
  - *setting and definitions*
- ***Efficient & precise simulation***
  - *today: what can go wrong?*

… the core of every game?

End of the day: ***be able to implement efficient shortest path, two-player AI, and to simulate flying pebbles (for A3!)***

Breadth-first vs. A*

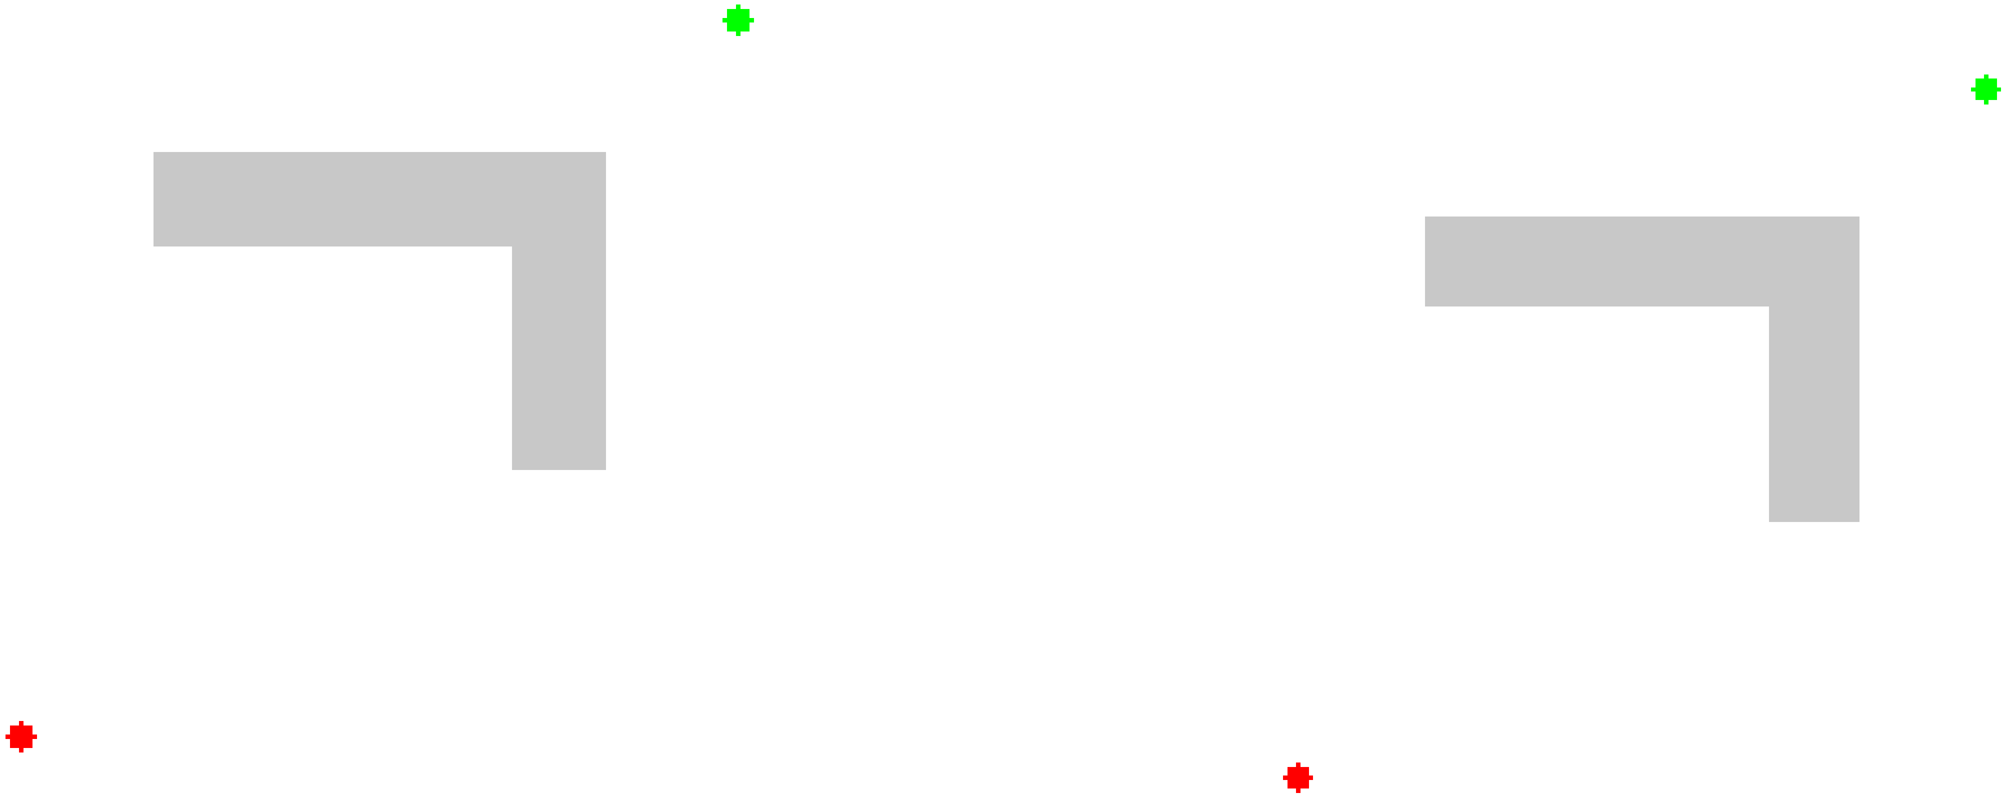

A* Search

- A* search takes into account both
  - *c(p)* = cost of path *p* to current node
  - *h(p)* = heuristic value at node *p (estimated "remaining" path cost)*
- Let *f(p)* = *c(p)* + *h(p).*
  - *f(p)* is an estimate of the cost of a path from the start to a goal via *p*.

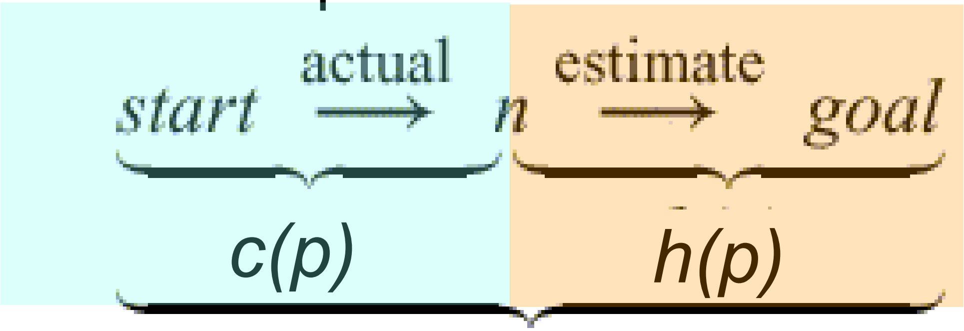

A* always chooses the path on the frontier with the lowest estimated distance from the start to a goal node constrained to go via that path.

A* Example

Init:

- ***Put starting node on open list:*** Lo = {6}
- ***Set its cost to 0:*** c[6] = 0
- ***Set closed list to empty list:*** Lc = {}

Step 1:

- ***Find node with smallest f on the list, call it q:*** q = 6
- ***Find q's "successors":*** sucs = {3,4,7}
- ***For each successor u:*** for u in sucs …
  - c(u) = c(q) + d(q,u)   c[3] = c[6] + 1 = 1
    c[4] = c[6] + 1.4 = 1.4
    c[7] = c[6] + 1 = 1
  - h(u) = d(g, u)   h[3] = 3.6   f[3] = c[3] + h[3] = 4.6
    f(u) = c(u) + h(u)   h[4] = 2.8   f[4] = c[4] + h[4] = 4.2
    h[7] = 3.6   f[7] = c[7] + h[7] = 4.6
- ***add successors to open list and move q to closed:***
  Lo = {3,4,7}; Lc = {6}

Frontier (open list)

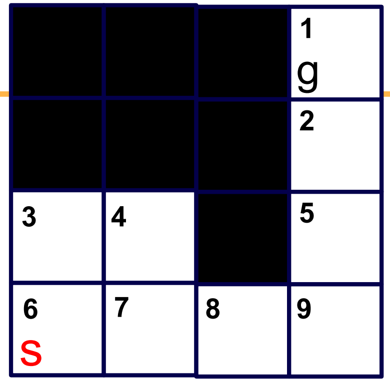

Step cost c

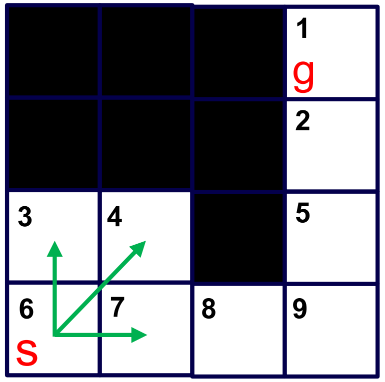

Heuristic dist. h

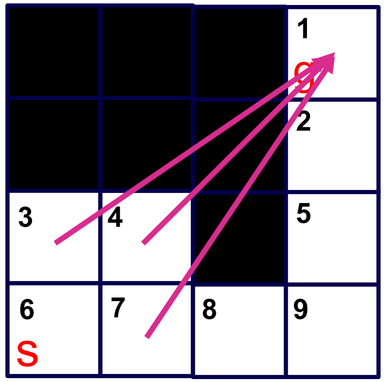

A* Example

Step 2: Lo = {3,4,7}; Lc = {6}

- ***Find node with smallest f on Lo, call it q:***
  - f[3] = 4.6
    f[4] = 4.2 -> q = 4
    f[7] = 4.6
- ***Find q's "successors":*** sucs = {3,6,7,8}
- for u in sucs…
  - c_tmp[3] = c[4] + 1 = 2.4 > c[3] = 1, skip
    c_tmp[6] = c[4] + 1.4 = 2.8 > c[6] = 0, skip
    c_tmp[7] = c[4] + 1 = 2.4 > c[7] = 1, skip
    c_tmp[8] = c[4] + 1.4 = 2.4 not in Lo or Lc, select c[8] = c_tmp[8]
  - Update heuristic and estimated cost f:
    h[8] = 3.2
    f[8] = c[8] + h[8] = 5.6
- ***add successors to open list and move q to closed list:***
  Lo = {3,7,8}; Lc = {6,4}

Frontier (open list)

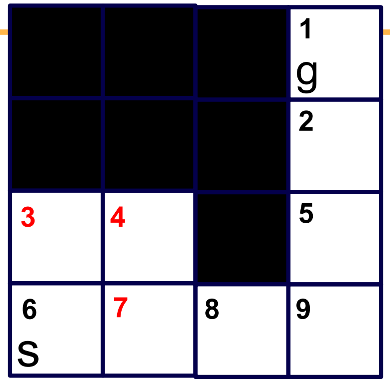

Step cost c

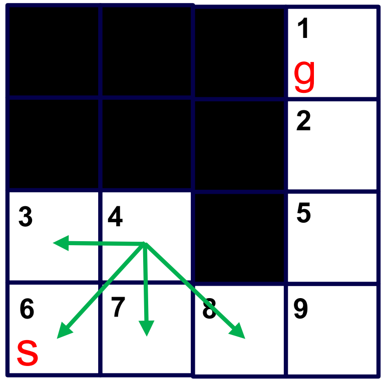

Heuristic dist. h

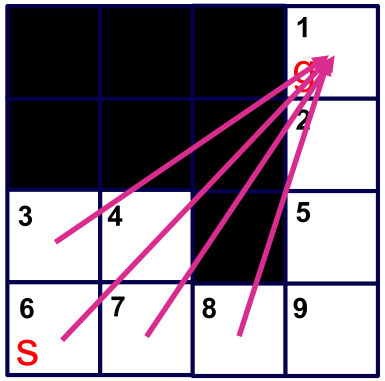

A* Example

Step 3: Lo = {3,7,8}; Lc = {6,4}

- ***Find node with smallest f on Lo, call it q:***
  - f[3] = 4.6 -> q = 3
    f[7] = 4.6
    f[8] = 5.6
- ***Find q's "successors":*** sucs = {4,6,7}
- for u in sucs…
  - c_tmp[4] = c[3] + 1 = 2 > c[4] = 1.4, skip
    c_tmp[6] = c[3] + 1.4 = 2.4 > c[6] = 0, skip
    c_tmp[7] = c[3] + 1 = 2 > c[7] = 1, skip
- ***add successors to open list? no successors!***
- ***move q to closed list:***
  Lo = {7,8};
  Lc = {6,4,3}

Frontier (open list)

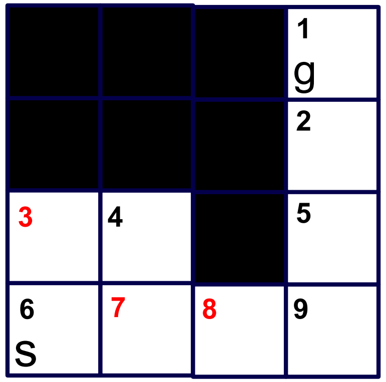

Step cost c

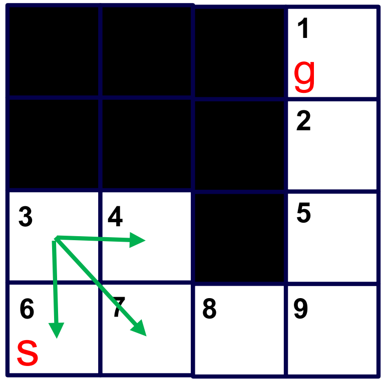

Heuristic dist. h

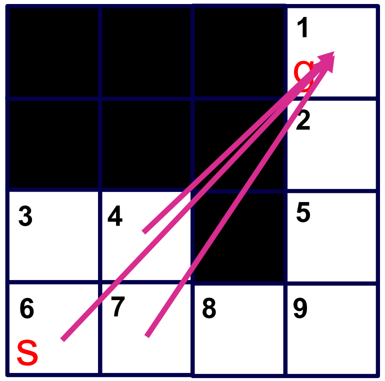

A* Example

Step 4: Lo = {7,8}; Lc = {6,4,3}

- ***Find node with smallest f on Lo, call it q:***
  - f[7] = 4.6   ->   q = 7
    f[8] = 5.6
- ***Find q's "successors":*** sucs = {3,4,6,8}
- **for u in sucs…**
  - c_tmp[3] = c[7] + 1.4 = 2.4   >   c[3] = 1, skip
    c_tmp[4] = c[7] + 1   = 2   >   c[4] = 1, skip
    c_tmp[6] = c[7] + 1   = 2   >   c[6] = 0, skip
    c_tmp[8] = c[7] + 1   = 2   >   c[8] = 2.4, select new c[8] = 2
- ***add successors to open list? Already there!***
- ***move q to closed list:***
  Lo = {8};
  Lc = {6,4,3,7}

Frontier (open list)

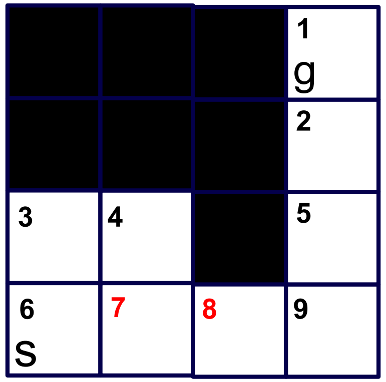

Step cost c

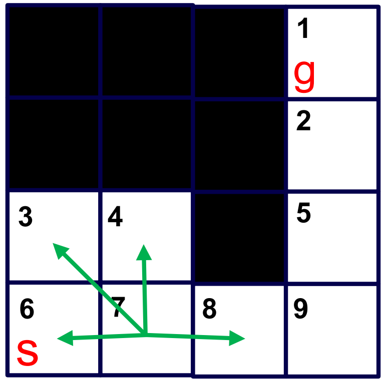

Heuristic dist. h

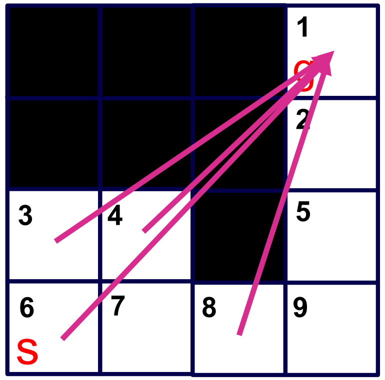

Keep track of your parents

- ***We neglected parent-child relation in previous slides…***

Lc = {6,4,3}   Lo = {8};

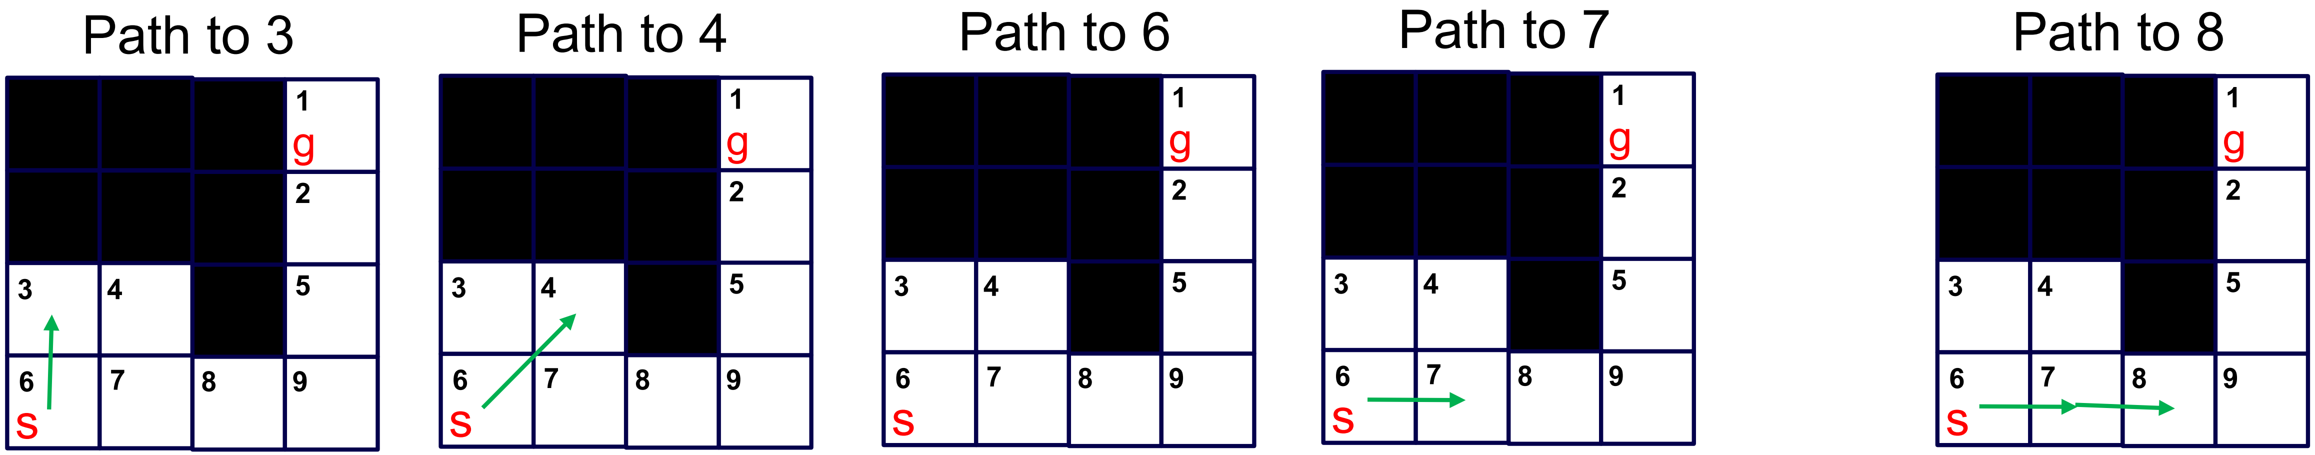

- ***Note, closed paths have no 'free' neighbors***
  - impassable or already visited from a shorter path

A* search

Key idea: H is a heuristic, and not the real distance:

h(p,q) = |(p.x – q.x)| + |(p.y – q.y)|

- Manhattan distance

h(p,q) = sqrt((p.x – q.x)^2 + (p.y – q.y)^2)

- Euclidean distance

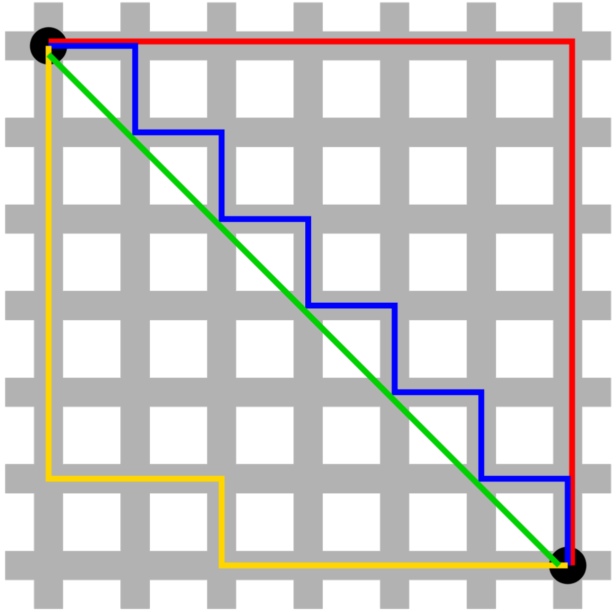

https://en.wikipedia.org/wiki/Taxicab_geometry

Conditions:

- a heuristic function is **admissible** if it never overestimates the cost of reaching the goal
- a heuristic function is said to be **consistent**, or **monotone**, if its estimate is always less than or equal to the estimated distance from any neighbouring vertex to the goal, plus the cost of reaching that neighbour

Variants

- ***Randomness***
- ***Make the AI dump/non-perfect***
  - *How?*
- ***Different terrain types?***

Two-player games

www.npr.org

Min-Max Trees

- Adversarial planning in a turn-taking environment
  - *Algorithm seeks to maximize our success* ***F***
  - *Adversary seeks to minimize* ***F***
  - $\boldsymbol{a_{we}} = \max_{\boldsymbol{we}} \min_{\boldsymbol{they}} \boldsymbol{F}(\boldsymbol{a_{we}}, \boldsymbol{a_{they}})$
- Key idea: at each step the algorithm selects the move that minimizes the highest (estimated) value of F the adversary can reach
  - *Assume the opponent does what is best*

Example

(from *uliana.lecturer.pens.ac.id/Kecerdasan%20Buatan/ppt/Game%20Playing/gametrees.ppt*)

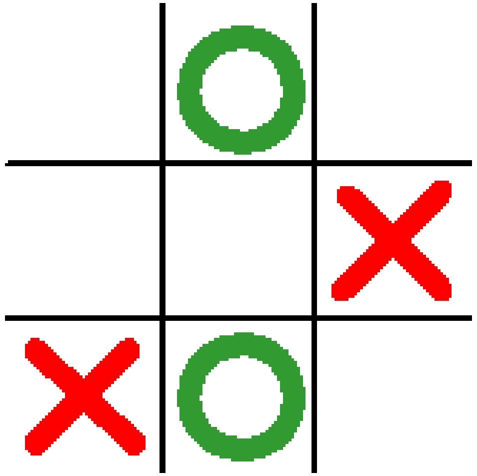

We are playing X, and it is now our turn.

Our options:

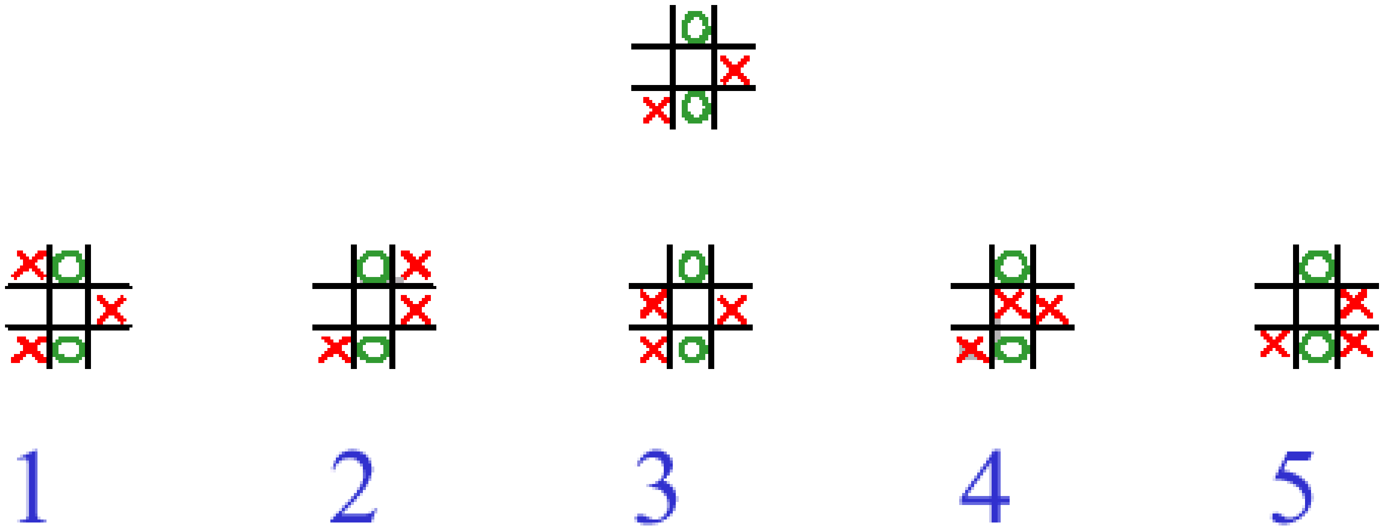

Number = position after each legal move

Opponent options

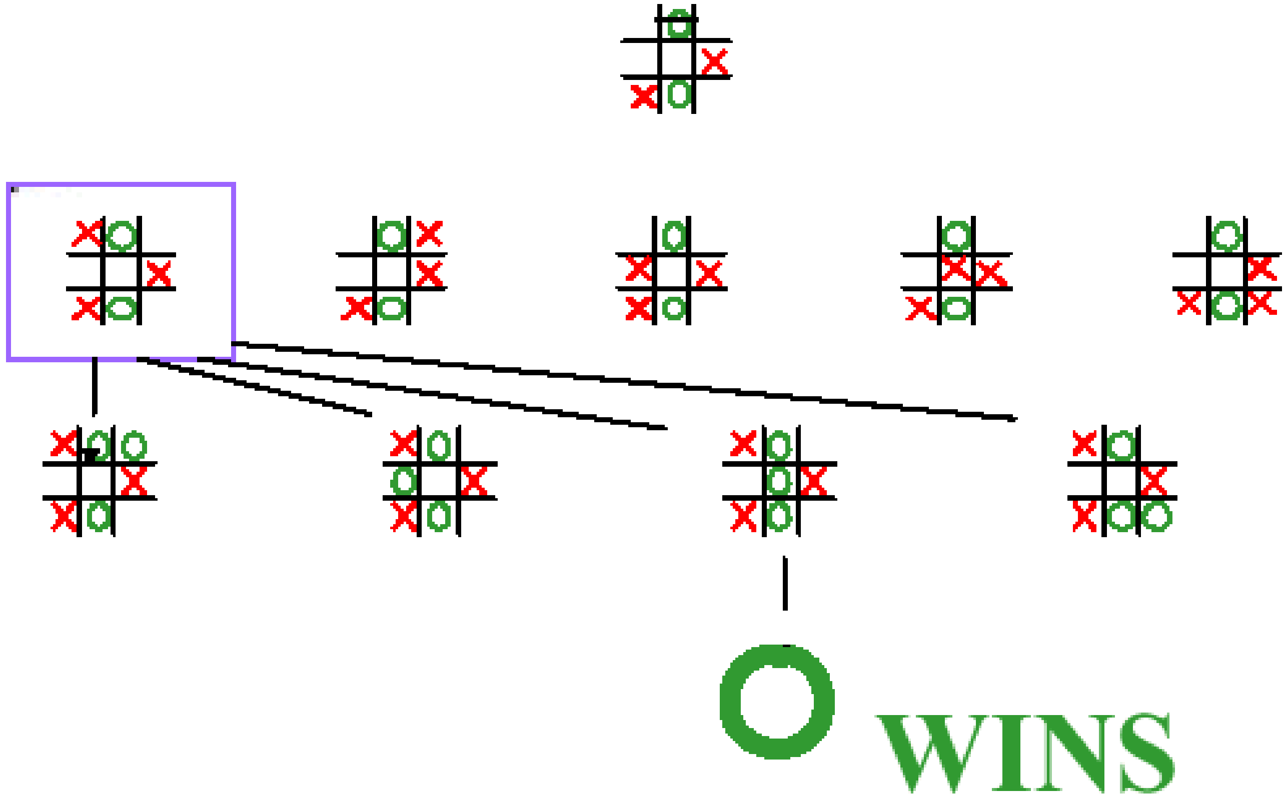

Here we are looking at all of the opponent responses to the first possible move we could make.

Opponent options

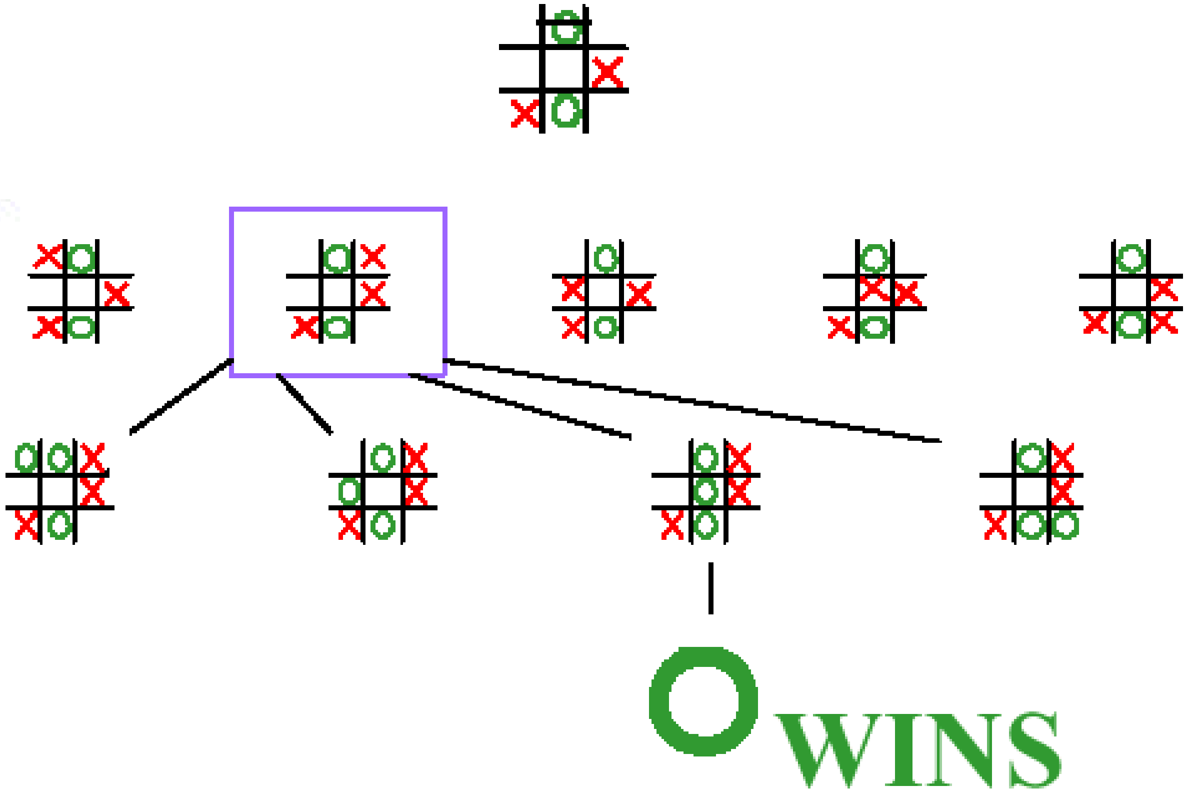

Opponent options after our second possibility. Not good again…

Opponent options

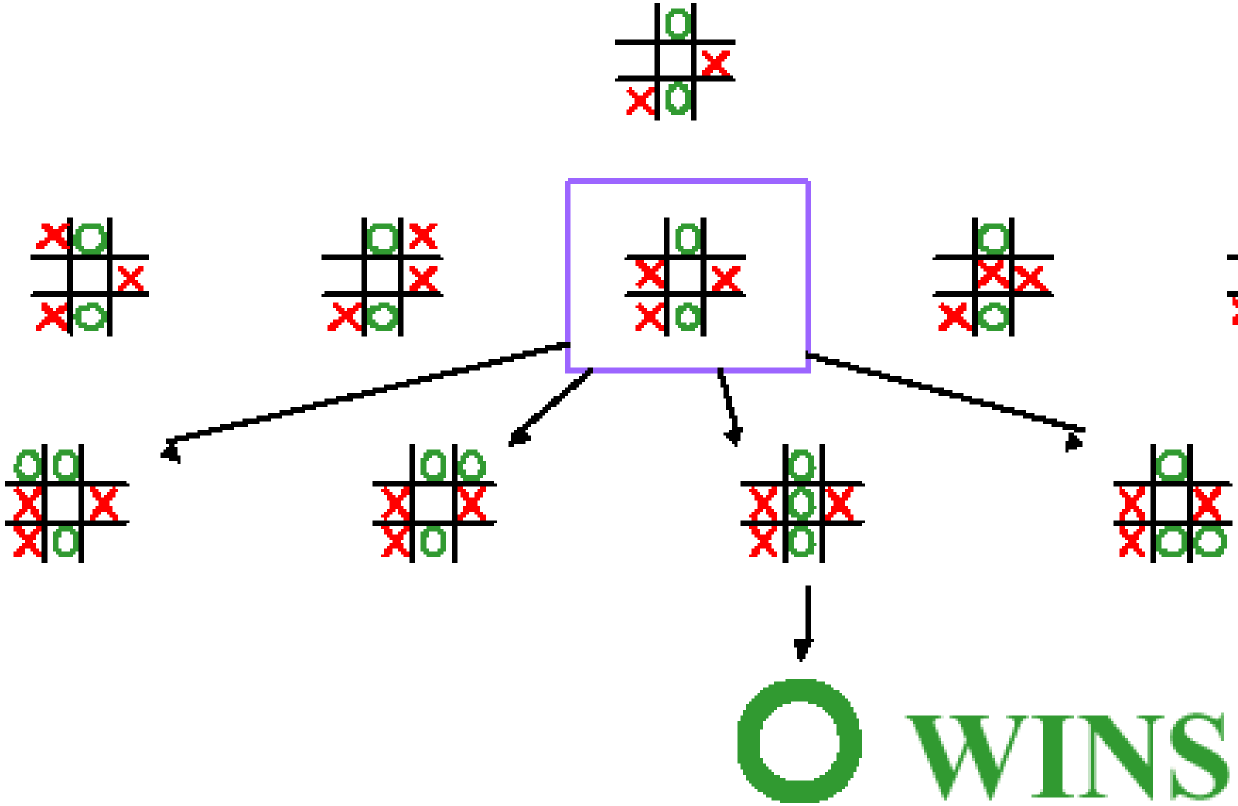

Opponent options => Our options

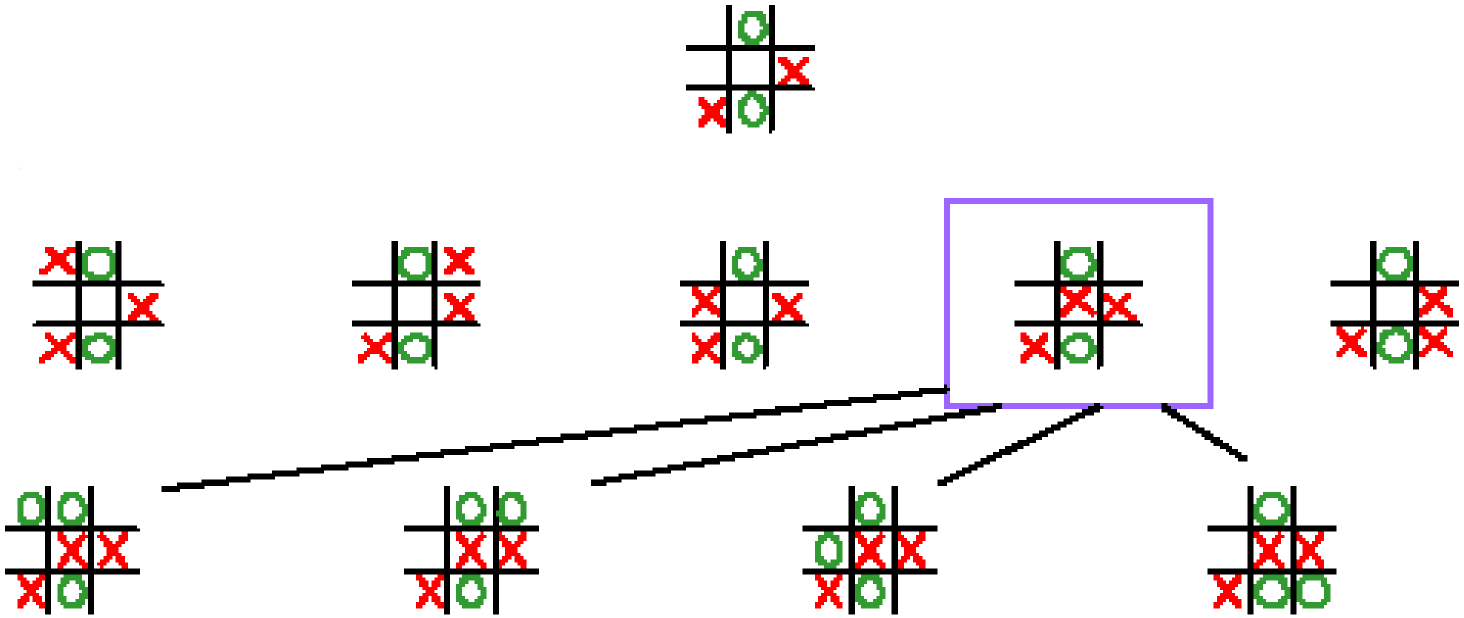

Now they don't have a way to win on their next move. So now we have to consider our responses to their responses.

Our options

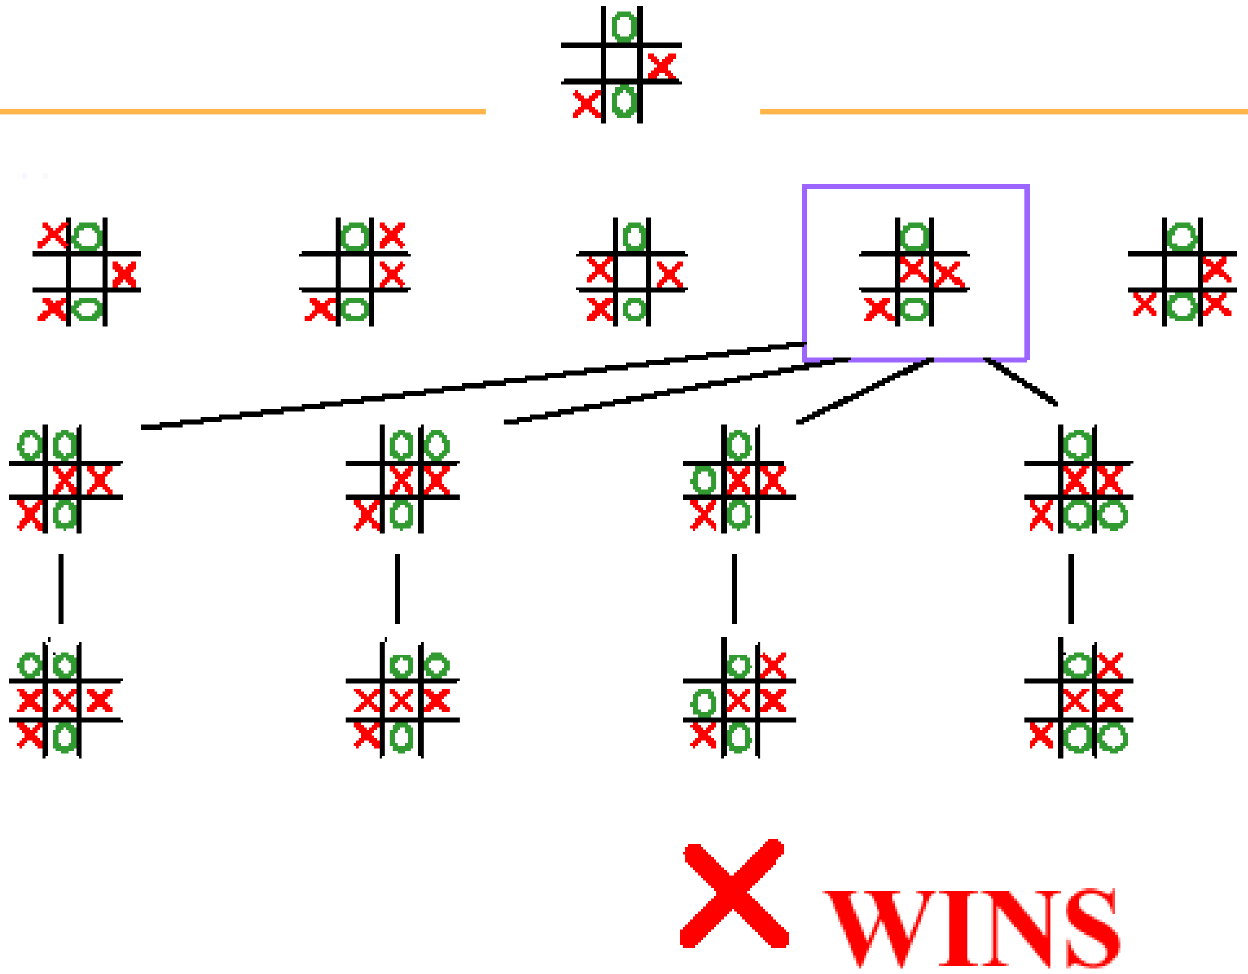

We have a win for any move they make.
Original position in purple is an X win.

Other options

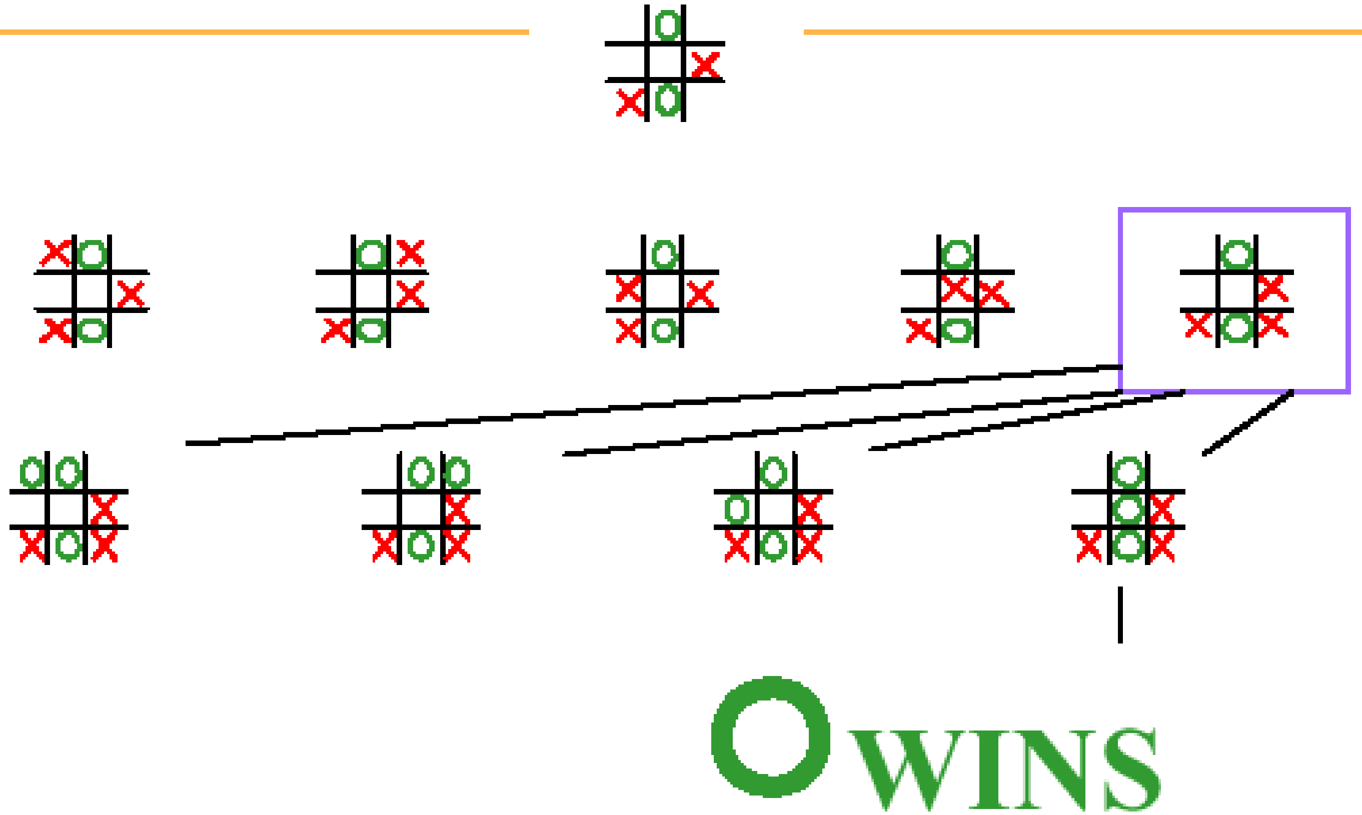

They win again if we take our fifth move.

Summary of the Analysis

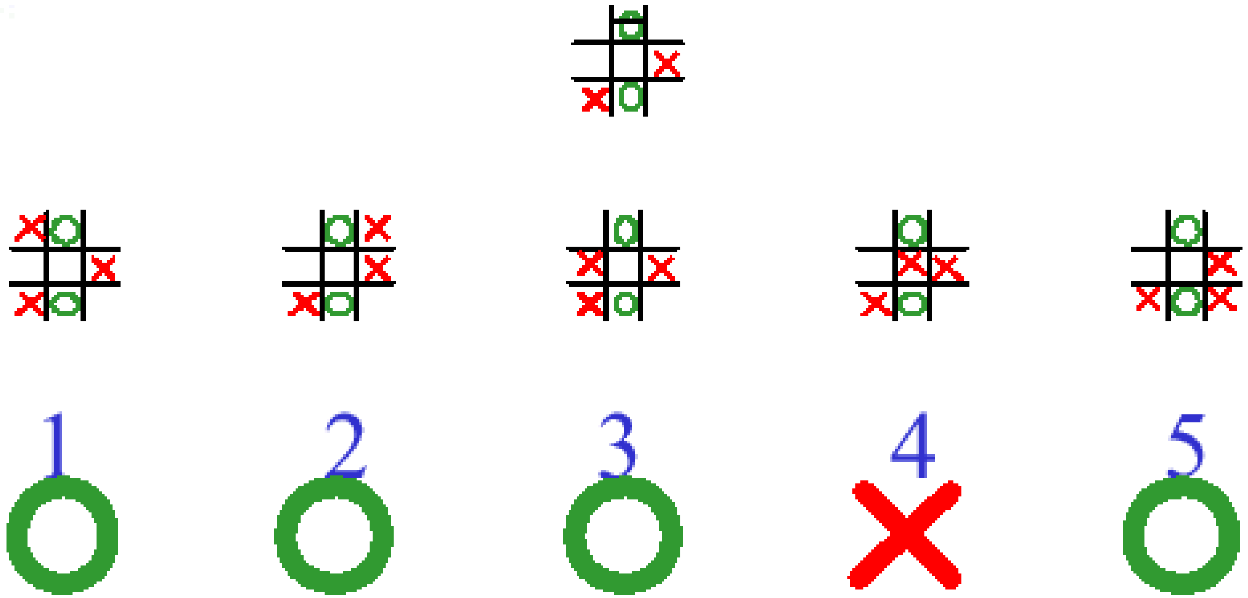

So which move should we make? ;-)

MinMax algorithm

- Traverse “game tree”:
  - *Enumerate all possible moves at each node.*
  - *The children of each node are the positions that result from making each move. A leaf is a position that is a draw or a win for some side.*

- Assume that we pick the best move for us, and the opponent picks the best move for them (causes most damage to us)

- Pick the move that maximizes the minimum amount of success for our side.

MinMax Algorithm

- Tic-Tac-Toe: three forms of success: Win, Tie, Lose.
  - *If you have a move that leads to a Win make it.*
  - *If you have no such move, then make the move that gives the tie.*
  - *If not even this exists, then it doesn’t matter what you do.*

Extensions

- Challenges: In practice
  - *Trees too deep/large to explore*
  - *Opponent not always makes the ‘best’ choice*
  - *Randomness*
- Solution - Heuristics
  - *Rate nodes based on local information.*
  - *For example, in Chess “rate” a position by examining difference in number of pieces*

Heuristics in MinMax

- Strategy that will let us cut off the game tree at fixed depth (layer)
- Apply heuristic scoring to bottom layer
  - *instead of just Win, Loss, Tie, we have a score.*
- For “our” level of the tree we want the move that yields the node (position) with highest score. For a “them” level “they” want the child with the lowest score.

Self stuy: Pseudocode

```
int Minimax(Board b, boolean myTurn, int depth) {
      if (depth==0)
            return b.Evaluate(); // Heuristic
      for(each possible move i)
            value[i] = Minimax(b.move(i), !myTurn,
depth-1);
      if (myTurn)
            return array_max(value);
      else
            return array_min(value);
}
```

Note: we don't use an explicit tree structure.
However, the pattern of recursive calls forms a tree on the call stack.

Real Minimax Example

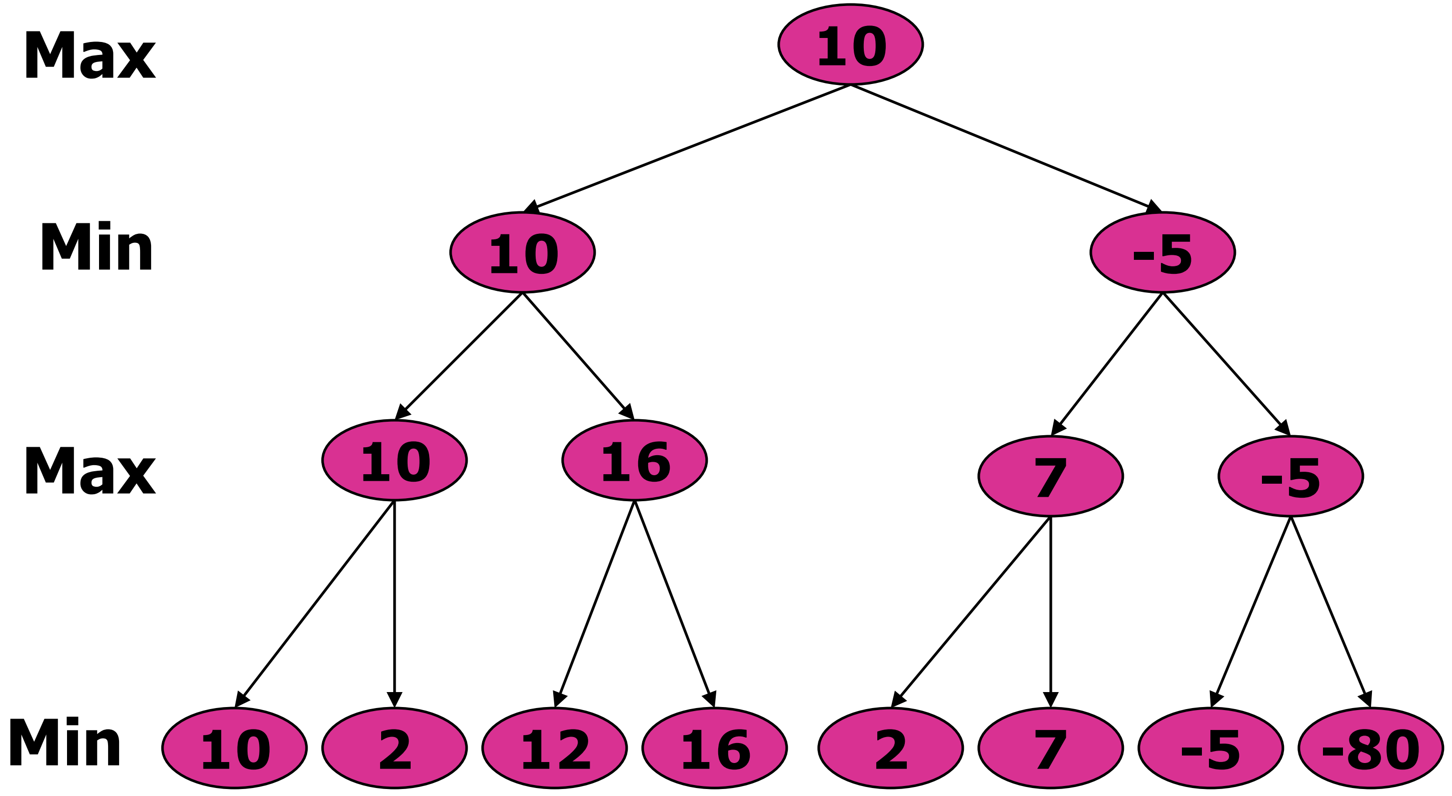

Evaluation function applied to the leaves!

Pruning Example

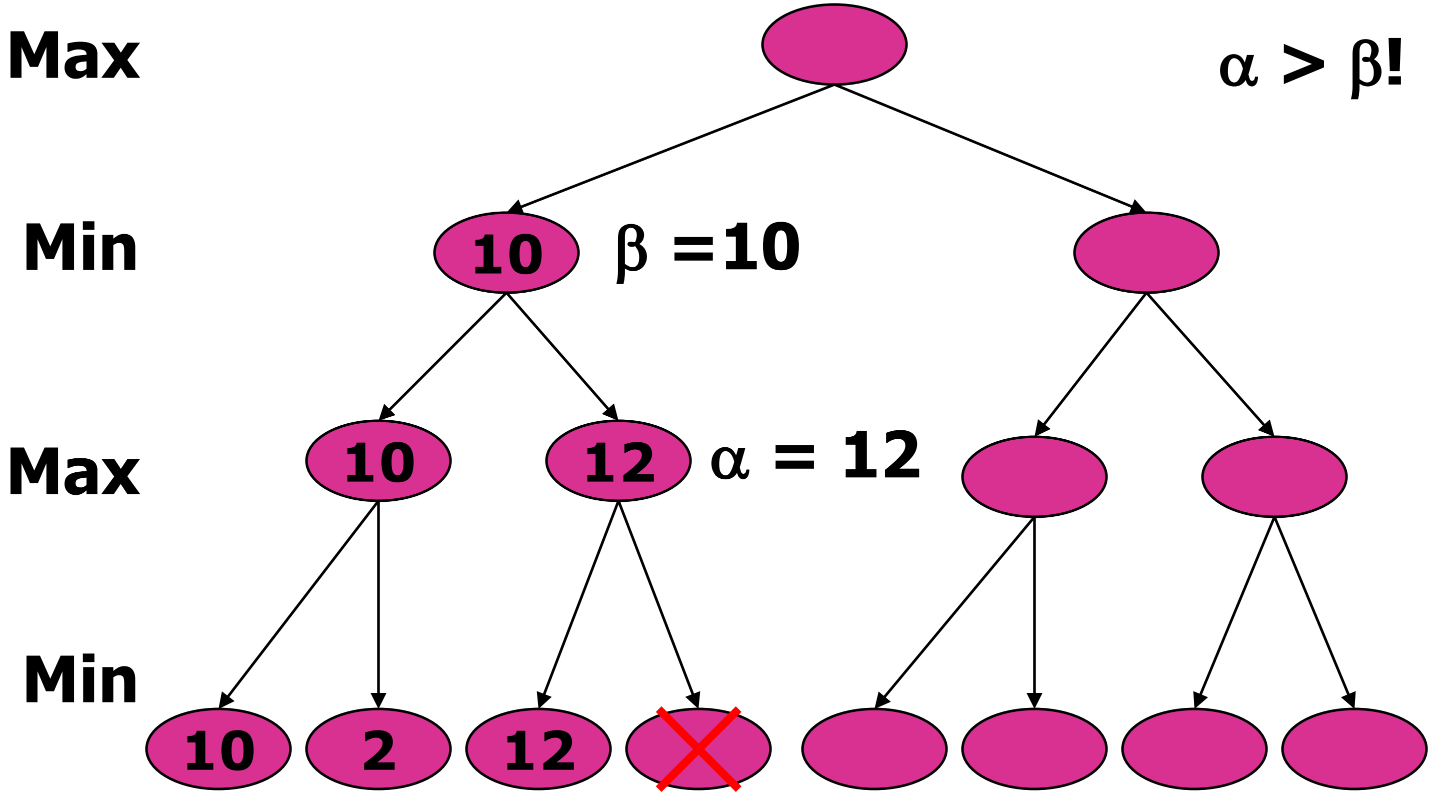

Self stuy: Alpha Beta Pruning

Idea: Track "window" of expectations.

Use two variables

- α – Best score so far at a **max** node ('our choice'): increases
  - *At a child **min** node:*
    - Parent wants **max**. To affect the parent's current α**,** our β cannot drop below α**.**
  - *If β ever gets less:*
    - Stop searching further subtrees *of that child*. They do not matter!
- β – Best score so far at a **min** node ('their choice'): decreases
  - *At a child **max** node.*
    - Parent wants **min**. To affect the parent's current β**,** our α cannot get above the parent's β**.**
  - *If α gets bigger than β:*
    - Stop searching further subtrees *of that child*. They do not matter!

Start with an infinite window ($\alpha = -\infty$, $\beta = \infty$)

Self stuy: Alpha Beta Example II

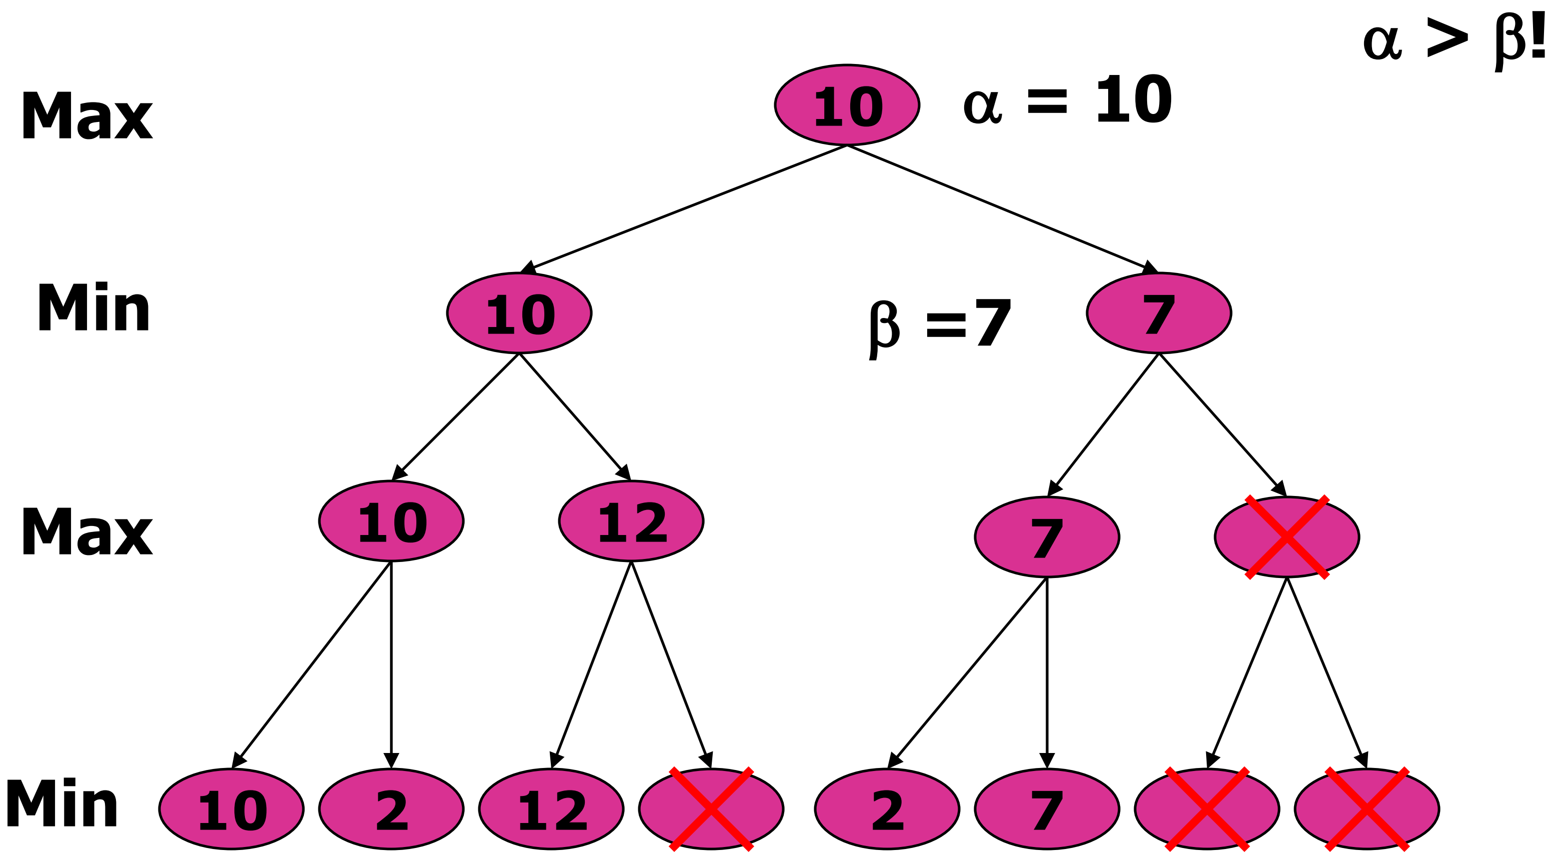

Self stuy: Pseudo Code

```
int AlphaBeta(Board b, boolean myTurn, int depth, int alpha, int beta) {
   if (depth==0)
      return b.Evaluate(); // Heuristic
   if (myTurn) {
      for(each possible move i && alpha < beta)
         alpha  = max(alpha,AlphaBeta(b.move(i),!myTurn,depth-1,alpha,beta));
      return alpha;
   }
   else {
      for(each possible move i && alpha < beta)
         beta  = min(beta,AlphaBeta(b.move(i), !myTurn, depth-1,alpha,beta));
      return beta;
   }
}
```

Variants

- ***More than two players?***

- ***More than two choices?***

- ***Opponent does not select best move?***