


Artificial Intelligence

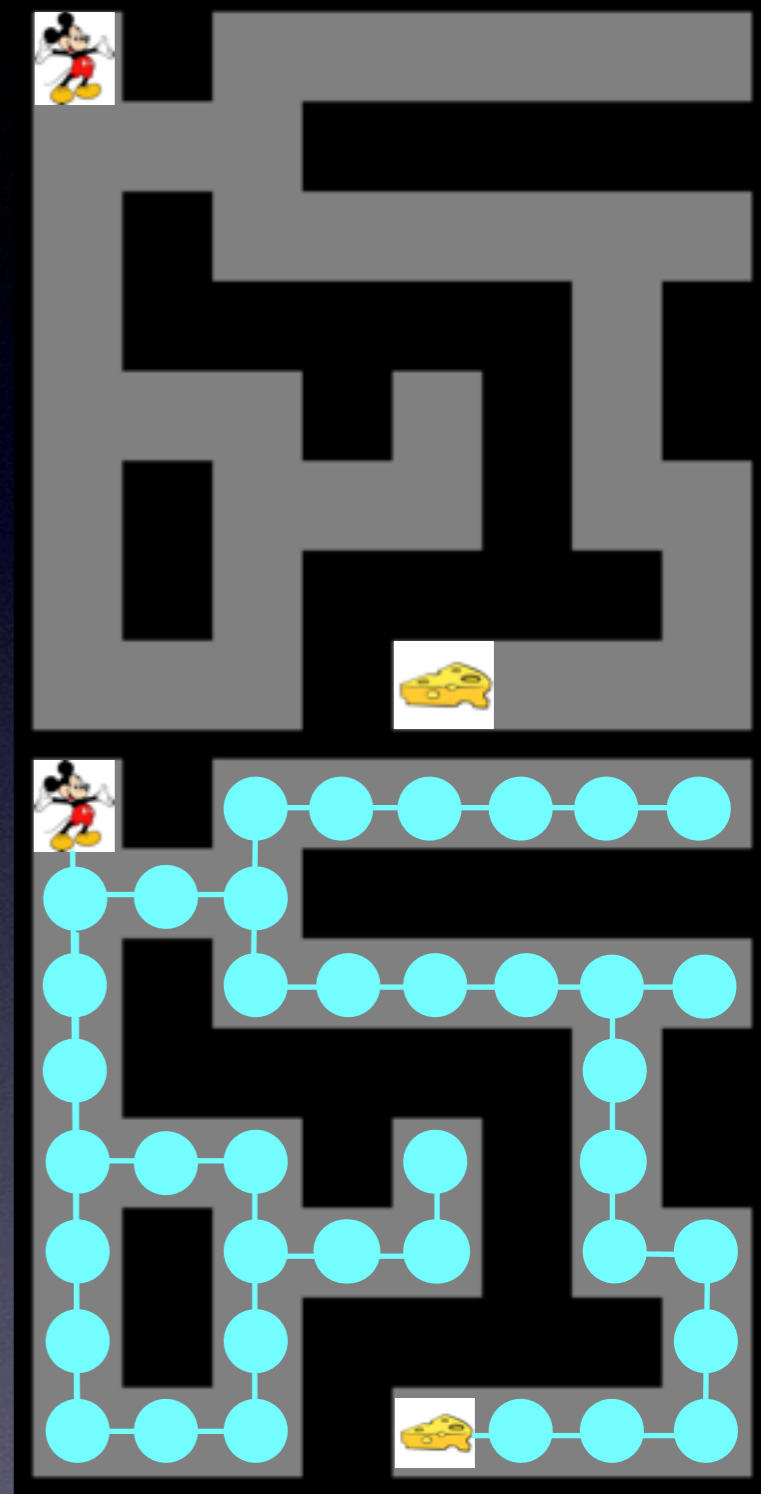



Jeff Clune

Assistant Professor
Evolving Artificial Intelligence Laboratory




AI Challenge One




AI Challenge One

- Transform to graph
- Explore the graph, looking for?
 - dust
 - unexplored squares



AI Challenge One, Question 2

- Some good solutions
 - 2013 champ: 0.97
 - Move greedily towards dust or unexplored tiles in sensor range.
 - Otherwise start a BFS toward nearest unexplored tile.
 - 2014 champ: 0.98
 - BFS with a maximum depth of 2 towards either dust or unexplored tile
 - If nothing found: run simple reflex agent (move randomly for the most part)
 - 2015: 0.998 (!)
 - Uniform Cost Search (effectively BFS) toward dust first, if no known dust UCS toward unexplored tile instead
 - 2016: 0.9992 (!!!) & 0.9988 (!!)



AI Challenge Three!

- Due: Sept. 25th
- This Sunday!



Evaluation Functions

- alpha-beta still needs to find leaves of the tree
 - too deep in many cases
- workaround: don't go to leaves, but instead estimate the expected value of an intermediate state



Value of these states?

Evaluation Functions

- humans use them
 - no one can “see ahead” to terminal states in chess
- effect of evaluation functions is large
 - a bad one will lead to bad play and vice versa
- must be fast
 - (that’s the point...to save computation)
- Example from chess:
 - Sum of: Pawns (1), knight/bishop (3), rook (5), queen (9).
 - Possibly add “good pawn structure” (0.5), castle (0.5), etc.
 - Called “**features**”

Evaluation Functions

- Must decide what to conflate
 - learn value of each board state
 - vs. counting pieces
 - assumes layout doesn't matter
- Often a **weighted linear sum** is used:
 - E.g. $\text{value} = 9 * \text{numQueen} + 5 * \text{numRook} + 1 * \text{numPawn} \dots$
 - assumes contributions are independent/non-interacting/
non-epistatic
 - to include interactions a non-linear function can be used

Evaluation Function

- Note: They are not part of the rules, must be learned
- Can be learned!
 - How would you do it?
 - In groups come up with as many ways as you can (and pick the one you'd recommend)

Evaluation Function

- Note: They are not part of the rules, must be learned
- Can be learned!
 - How would you do it?
 - In groups come up with as many ways as you can (and pick the one you'd recommend)
 - Ideas we won't talk about in detail
 - Evolve the weights in the linear weighted sum
 - Deep learning

Evaluation Function


- Monte Carlo (“rollouts”)
 - random play to the end repeated N times to estimate state value
 - works pretty well with random play, though would be better with intelligent play
- UCT: more intelligent play
 - increasingly focuses search on promising areas discovered during random play

Evaluation Functions + Alpha Beta

- Can use Alpha Beta out of the box with evaluation functions
 - just pick a max-depth or other stopping criterion
 - or pick `maxTimeAllowed` and run iterative deepening until you run out of time

Lookup

- Silly to search millions of nodes to pick the opening move
- Can just lookup what to do in common situations
 - openings and endgames
 - read book for fascinating discussion of how much better AI is than humans at endgames
 - one series requires 517 moves but leads to a guaranteed checkmate!
- Usually after 10 moves the board state is rare enough that AI has to switch from lookup to search



Deep Blue

- regularly got to 14 ply
 - some forcing sequences went to 40ply
- 30 billion positions per move
- evaluation function had 8000 features!
- 4000 position opening book
- 700,000-game library of games to learn from
- all 5, and most 6-piece endgames solved
- nowadays better algos mean standard PCs can play well

Humans Can No Longer Win At...

- Humans can no longer win at...
 - Checkers (solved)
 - Chess
 - Othello
 - Scrabble
- Tie
 - backgammon
- Humans better at
 - Go
 - Hopscotch

Is this out of date?
If so, email me.

Humans Can No Longer Win At...

- Humans can no longer win at...
 - Checkers (solved)
 - Chess
 - Othello
 - Scrabble
 - Go
- Tie
 - backgammon
- Humans better at
 - Hopscotch

Is this out of date?
If so, email me.

Go

- Branching factor too large: ~250 to ~361 (depending on source)
 - and games go for ~350 moves
- Evaluations too hard (so far!):
 - UCT
 - with extra tricks to suggest which plays to explore
 - (similar to killer move heuristic)
- Current programs can only compete on smaller boards

Wrong! Time to rewrite the textbooks!

Adversarial Search: Key concepts

- Pruning
- Evaluation function
 - evaluate intermediate game states since optimal search is impossible
- Minimax
- Alpha-beta pruning
 - saves time, without any cost in game performance
 - killer heuristic



Stochastic Games

- Examples?

Stochastic Games

- Instead of a minimax value, we calculate expected value
 - (value of each node) * (chance of that node occurring)
 - Which has higher expected utility/value?
 - Option 1: 50% chance of payoff=10, 50% chance of payoff=1
 - Option 2: 90% chance of payoff=6, 10% chance of payoff=0

Partially Observable Games

- E.g. war, bridge, etc.
 - my favorite is Stratego
- Gathering info becomes a move in some games
 - scouts/spies
- Bluffing is important

Ch. 13: Uncertainty

- Uncertainty is pervasive in the world
 - e.g. diagnosing an illness
- Goal: maximize expected utility/value
- Probability Theory is our best tool
 - Lots of help with basic equations & notation in the book

Bayesian Statistics

- Very important in AI
- Allow you to
 - have prior knowledge about the world
 - e.g. phones don't have cameras
 - update your knowledge of the world
 - e.g. now they do!
- Most of the important info is in Ch. 13.

Bayesian Statistics

- Priors
 - aka “unconditional probabilities” or “prior probabilities”
 - belief before seeing evidence
 - e.g. most phones don't have cameras (belief in 2000)
 - $P(\text{mostPhonesHaveCameras})$
- Posterior
 - aka “conditional probabilities” or “posterior

Bayesian Statistics

- Prior
 - $P(\text{two dice sum to } 12) = ??$
- Posterior
 - $P(\text{two dice sum to } 12 \mid \text{Die}_1=6) = ??$

Reminder About Probability

Mathematically speaking, conditional probabilities are defined in terms of unconditional probabilities as follows: for any propositions a and b , we have

$$P(a | b) = \frac{P(a \wedge b)}{P(b)}, \quad \text{Note: } \wedge \text{ means "and"} \quad (13.3)$$

which holds whenever $P(b) > 0$. For example,

$$P(\text{doubles} | \text{Die}_1 = 5) = \frac{P(\text{doubles} \wedge \text{Die}_1 = 5)}{P(\text{Die}_1 = 5)}.$$

The definition of conditional probability, Equation (13.3), can be written in a different form called the **product rule**:

$$P(a \wedge b) = P(a | b)P(b),$$

The product rule is perhaps easier to remember: it comes from the fact that, for a and b to be true, we need b to be true, and we also need a to be true given b .