

CSE 473: Artificial Intelligence Spring 2012

Adversarial Search: Expectimax Dan Weld

Based on slides from
Dan Klein, Stuart Russell, Andrew Moore and Luke Zettlemoyer

Space of Search Strategies

- **Blind Search**
 - DFS, BFS, IDS
- **Informed Search**
 - Systematic: Uniform cost, greedy, A*, IDA*
 - Stochastic: Hill climbing w/ random walk & restarts
- **Constraint Satisfaction**
 - Backtracking=DFS, FC, k-consistency, exploiting structure
- **Adversary Search**
 - Mini-max
 - Alpha-beta
 - Evaluation functions
 - Expecti-max

2

Overview


- Introduction & Agents
- Search, Heuristics & CSPs
- Adversarial Search
- Logical Knowledge Representation
- Planning & MDPs
- Reinforcement Learning
- Uncertainty & Bayesian Networks
- Machine Learning
- NLP & Special Topics

Types of Games

	deterministic	chance
perfect information	chess, checkers, go, othello	backgammon, monopoly
imperfect information	stratego	bridge, poker, scrabble, nuclear war

Number of Players? 1, 2, ...?

Tic-tac-toe Game Tree




Mini-Max

- **Assumptions**
 - High scoring leaf == good for you (bad for opp)
 - Opponent is super-smart, rational; never errs
 - Will play optimally against you
 - **Idea**
 - Exhaustive search
 - Alternate: best move for you; best for opponent
- Max ↗ ↖ Min
- **Guarantee**
 - Will find best move for you (given assumptions)


α-β Pruning General Case

- Add α, β bounds to each node
- α is the best value that MAX can get at any choice point along the current path
- If value of n becomes worse than α , MAX will avoid it, so can stop considering n 's other children
- Define β similarly for MIN



Heuristic Evaluation Function

- Function which scores non-terminals



$$Eval(s) = w_1 f_1(s) + w_2 f_2(s) + \dots + w_n f_n(s)$$


- Ideal function: returns the utility of the position
- In practice: typically weighted linear sum of features:
 - e.g. $f_1(s) = (\text{num white queens} - \text{num black queens})$, etc.

Modeling Opponent

- So far assumed
Opponent = rational, infinitely smart
- What if
Opponent = random?
2 player w/ random opponent = 1 player stochastic
- Later...
Sorta smart
Infinitely smart, but actions have chance

Stochastic Single-Player

- What if we don't know what the result of an action will be? E.g.,
 - In solitaire, shuffle is unknown
 - In minesweeper, mine locations
- Can do **expectimax search**
 - Chance nodes, like actions except the environment controls the action chosen
 - Max nodes as before
 - Chance nodes take average (expectation) of value of children



Maximum Expected Utility

- Why should we average utilities? Why not minimax?
- Principle of maximum expected utility: an agent should choose the action which maximizes its expected utility, given its knowledge
 - General principle for decision making
 - Often taken as the definition of rationality
 - We'll see this idea over and over in this course!
- Let's decompress this definition...

Reminder: Probabilities

- A **random variable** represents an event whose outcome is unknown
- A **probability distribution** is an assignment of weights to outcomes
- Example: traffic on freeway?
 - Random variable: T = whether there's traffic
 - Outcomes: T in {none, light, heavy}
 - Distribution: $P(T=\text{none}) = 0.25$, $P(T=\text{light}) = 0.55$, $P(T=\text{heavy}) = 0.20$
- Some laws of probability (more later):
 - Probabilities are always $\in [0, 1]$
 - Probabilities (over all possible outcomes) sum to one
- As we get more evidence, probabilities may change:
 - $P(T=\text{heavy}) = 0.20$, $P(T=\text{heavy} | \text{Hour}=8\text{am}) = 0.60$
 - We'll talk about methods for reasoning and updating probabilities later

What are Probabilities?

- **Objectivist / frequentist answer:**
 - Averages over repeated *experiments*
 - E.g. empirically estimating P(rain) from historical observation
 - E.g. pacman's estimate of what the ghost will do, given what it has done in the past
 - Assertion about how future experiments will go (in the limit)
 - Makes one think of *inherently random* events, like rolling dice
- **Subjectivist / Bayesian answer:**
 - Degrees of belief about unobserved variables
 - E.g. an agent's belief that it's raining, given the temperature
 - E.g. pacman's belief that the ghost will turn left, given the state
 - Often *learn* probabilities from past experiences (more later)
 - New evidence *updates beliefs* (more later)

Uncertainty Everywhere

- **Not just for games of chance!**
 - I'm sick: will I sneeze this minute?
 - Email contains "FREE!": is it spam?
 - Tooth hurts: have cavity?
 - 60 min enough to get to the airport?
 - Robot rotated wheel three times, how far did it advance?
 - Safe to cross street? (Look both ways!)
- **Sources of uncertainty in random variables:**
 - Inherently random process (dice, etc)
 - Insufficient or weak evidence
 - Ignorance of underlying processes
 - Unmodeled variables
 - The world's just noisy – it doesn't behave according to plan!

Reminder: Expectations

- We can define function $f(X)$ of a random variable X
- **The expected value of a function is its average value, weighted by the probability distribution over inputs**
- **Example: How long to get to the airport?**
 - Length of driving time as a function of traffic: L(none) = 20, L(light) = 30, L(heavy) = 60
 - What is my expected driving time?
 - Notation: $E_{P(T)}[L(T)]$
 - Remember, $P(T) = \{\text{none: 0.25, light: 0.5, heavy: 0.25}\}$
 - $E[L(T)] = L(\text{none}) * P(\text{none}) + L(\text{light}) * P(\text{light}) + L(\text{heavy}) * P(\text{heavy})$
 - $E[L(T)] = (20 * 0.25) + (30 * 0.5) + (60 * 0.25) = 35$

Expectations II

- Real valued functions of random variables:

$$f : X \rightarrow R$$
- Expectation of a function of a random variable

$$E_{P(X)}[f(X)] = \sum_x f(x)P(x)$$
- **Example: Expected value of a fair die roll**

X	P	f
1	1/6	1
2	1/6	2
3	1/6	3
4	1/6	4
5	1/6	5
6	1/6	6

$$1 \times \frac{1}{6} + 2 \times \frac{1}{6} + 3 \times \frac{1}{6} + 4 \times \frac{1}{6} + 5 \times \frac{1}{6} + 6 \times \frac{1}{6} = 3.5$$

Utilities

- **Utilities are functions from outcomes (states of the world) to real numbers that describe an agent's preferences**
- **Where do utilities come from?**
 - In a game, may be simple (+1/-1)
 - Utilities summarize the agent's goals
 - Theorem: any set of preferences between outcomes can be summarized as a utility function (provided the preferences meet certain conditions)
- In general, we hard-wire utilities and let actions emerge (why don't we let agents decide their own utilities?)
- More on utilities soon...


Expectimax Search Trees

- What if we don't know what the result of an action will be? E.g.,
 - In solitaire, next card is unknown
 - In minesweeper, mine locations
 - In pacman, the ghosts act randomly
- Can do **expectimax search**
 - Chance nodes, like min nodes, except the outcome is uncertain
 - Calculate **expected utilities**
 - Max nodes as in minimax search
 - Chance nodes take average (expectation) of value of children

Later, we'll learn how to formalize the underlying problem as a **Markov Decision Process**

Expectimax Search

- In expectimax search, we have a probabilistic model of how the opponent (or environment) will behave in any state
 - Model can be a simple uniform distribution (roll a die)
 - Model can be sophisticated and require a great deal of computation
 - We have a node for every outcome out of our control: opponent or environment
 - The model might say that adversarial actions are likely!
- For now, assume for any state we magically have a distribution to assign probabilities to opponent actions / environment outcomes




Expectimax Pseudocode

```

def value(s)
  if s is a max node return maxValue(s)
  if s is an exp node return expValue(s)
  if s is a terminal node return evaluation(s)

def maxValue(s)
  values = [value(s') for s' in successors(s)]
  return max(values)

def expValue(s)
  values = [value(s') for s' in successors(s)]
  weights = [probability(s, s') for s' in successors(s)]
  return expectation(values, weights)
    
```



Which Algorithm?

Minimax: no point in trying

QuickTime™ and a GIF-decompressor are needed to see this picture.

3 ply look ahead, ghosts move randomly

Which Algorithm?

Expectimax: wins some of the time

QuickTime™ and a GIF-decompressor are needed to see this picture.

3 ply look ahead, ghosts move randomly


Expectimax for Pacman

- Notice that we've stopped thinking that the ghosts are trying to minimize pacman's score
- Instead, they are now a part of the environment
- Pacman has a belief (distribution) over how they will act
- Quiz: Can we see minimax as a special case of expectimax?
- Quiz: what would pacman's computation look like if we assumed that the ghosts were doing 1-ply minimax and taking the result 80% of the time, otherwise moving randomly?

Expectimax for Pacman

Results from playing 5 games

	Minimizing Ghost	Random Ghost
Minimax Pacman	Won 5/5 Avg. Score: 493	Won 5/5 Avg. Score: 483
Expectimax Pacman	Won 1/5 Avg. Score: -303	Won 5/5 Avg. Score: 503



Pacman does depth 4 search with an eval function that avoids trouble
Minimizing ghost does depth 2 search with an eval function that seeks Pacman

Expectimax Pruning?

- Not easy
 - exact: need bounds on possible values
 - approximate: sample high-probability branches

Expectimax Evaluation

- Evaluation functions quickly return an estimate for a node's true value (which value, expectimax or minimax?)
- For minimax, evaluation function scale doesn't matter
 - We just want better states to have higher evaluations (get the ordering right)
 - We call this **insensitivity to monotonic transformations**
- For expectimax, we need *magnitudes* to be meaningful

Stochastic Two-Player

- E.g. backgammon
- Expectiminimax (!)
 - Environment is an extra player that moves after each agent
 - Chance nodes take expectations, otherwise like minimax

```

if state is a MAX node then
    return the highest EXPECTIMINIMAX-VALUE of SUCCESSORS(state)
if state is a MIN node then
    return the lowest EXPECTIMINIMAX-VALUE of SUCCESSORS(state)
if state is a chance node then
    return average of EXPECTIMINIMAX-VALUE of SUCCESSORS(state)
    
```

Stochastic Two-Player

- Dice rolls increase b : 21 possible rolls with 2 dice
 - Backgammon H 20 legal moves
 - Depth 4 = $20 \times (21 \times 20)^3 = 1.2 \times 10^9$
- As depth increases, probability of reaching a given node shrinks
 - So value of lookahead is diminished
 - So limiting depth is less damaging
 - But pruning is less possible...
- TDGammon uses depth-2 search + very good eval function + reinforcement learning → world-champion level play

Multi-player Non-Zero-Sum Games

	B = silent (cooperates)	B = confesses (defects)
A = silent (cooperates)	A: 1 month B: 1 month	A: 10 years B: 0
A = confesses (defects)	A: 0 B: 3 years	A: 3 years B: 3 years

Multi-player Non-Zero-Sum Games

- Similar to minimax:
 - Utilities are now tuples
 - Each player maximizes their own entry at each node
 - Propagate (or back up) nodes from children
 - Can give rise to cooperation and competition dynamically...

Types of Games

	deterministic	chance
perfect information	chess, checkers, go, othello	backgammon, monopoly
imperfect information	stratego	bridge, poker, scrabble, nuclear war

Number of Players? 1, 2, ...?