


Advanced Topics in AI

Probabilities



Instructor: Prof. Dr. techn. Wolfgang Nejdl
Leibniz University Hannover




Co-financed by the Connecting Europe
Facility of the European Union

[These slides were created by Dan Klein and Pieter Abbeel for CS188 Intro to AI at UC Berkeley. All materials are available at <http://ai.berkeley.edu>.]


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: $P(T = \text{none}) = 0.25$, $P(T = \text{light}) = 0.50$, $P(T = \text{heavy}) = 0.25$
- Some laws of probability (more later):
 - Probabilities are always non-negative
 - Probabilities over all possible outcomes sum to one
- As we get more evidence, probabilities may change:
 - $P(T = \text{heavy}) = 0.25$, $P(T = \text{heavy} | \text{Hour} = 8\text{am}) = 0.60$
 - We'll talk about methods for reasoning and updating probabilities later





Reminder: Expectations

- The expected value of a function of a random variable is the average, weighted by the probability distribution over outcomes
- Example: How long to get to the airport?

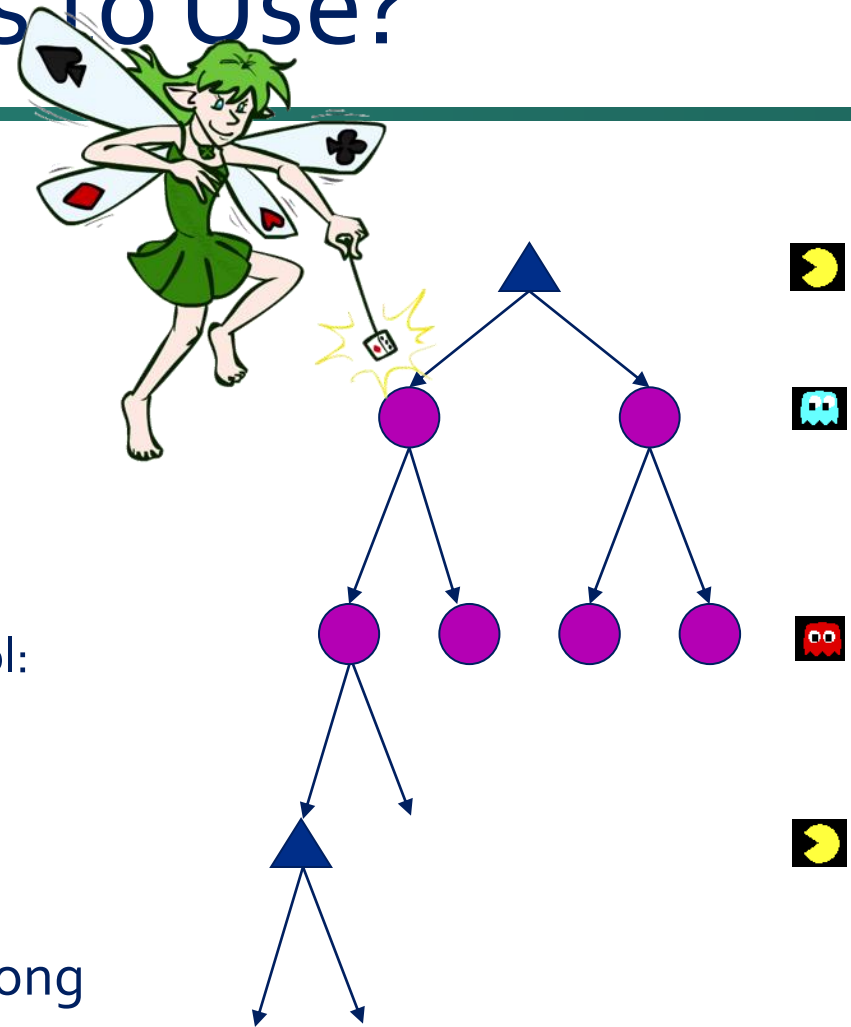


Time:	20 min	+	30 min	+	60 min	→	35 min
Probability:	0.25		0.50		0.25		



What Probabilities to Use?


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



Having a probabilistic belief about another agent's action does not mean that the agent is flipping any coins!

Quiz: Informed Probabilities


- Let's say you know that your opponent is actually running a depth 2 minimax, using the result 80% of the time, and moving randomly otherwise
- Question: What tree search should you use?



- Answer: **Expectimax!**
 - To figure out EACH chance node's probabilities, you have to run a simulation of your opponent
 - This kind of thing gets very slow very quickly
 - Even worse if you have to simulate your opponent simulating you...
 - ... except for minimax and maximax, which have the nice property that it all collapses into one game tree

Advanced Topics in AI

Next: Modeling Assumptions



Instructor: Prof. Dr. techn. Wolfgang Nejdl
Leibniz University Hannover



[These slides were created by Dan Klein and Pieter Abbeel for CS188 Intro to AI at UC Berkeley. All materials are available at <http://ai.berkeley.edu>.]