# CSE 573: Artificial Intelligence Autumn 2010 

## Lecture 9: RL / Probability Review 10/28/2010

## Luke Zettlemoyer

Many slides over the course adapted from either Dan Klein, Stuart Russell or Andrew Moore

## Outline

- Reinforcement Learning
- (review) Q-learning
- (finish) Linear function approximation
- Policy Iteration (optional)
- POMDPs (definition only)
- Probability review
- Random Variables and Events
- Joint / Marginal / Conditional Distributions
- Product Rule, Chain Rule, Bayes’ Rule
- Probabilistic Inference


## Recap: Reinforcement Learning

- Reinforcement learning:
- Still have an MDP:
- A set of states $s \in S$
- A set of actions (per state) A
- A model T(s,a,s')
- A reward function R(s,a,s')
- Still looking for a policy $\pi(\mathrm{s})$
- New twist: don't know T or R
- I.e. don't know which states are good or what the actions do
- Must actually try actions and states out to learn


## Recap: Q-Value Iteration

- Value iteration: find successive approx optimal values
- Start with $\mathrm{V}_{0}{ }^{*}(\mathrm{~s})=0$
- Given $\mathrm{V}_{\mathrm{i}}^{*}$, calculate the values for all states for depth $\mathrm{i}+1$ :

$$
V_{i+1}(s) \leftarrow \max _{a} \sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V_{i}\left(s^{\prime}\right)\right]
$$

- But Q-values are more useful!
- Start with $\mathrm{Q}_{0}{ }^{*}(\mathrm{~s}, \mathrm{a})=0$
- Given $\mathrm{Q}_{\mathrm{i}}{ }^{*}$, calculate the q -values for all $q$-states for depth $\mathrm{i}+1$ :

$$
Q_{i+1}(s, a) \leftarrow \sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \max _{a^{\prime}} Q_{i}\left(s^{\prime}, a^{\prime}\right)\right]
$$

## Recap: Q-Learning Update

- Q-Learning: sample-based Q-value iteration

$$
Q^{*}(s, a)=\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \max _{a^{\prime}} Q^{*}\left(s^{\prime}, a^{\prime}\right)\right]
$$

- Learn Q*(s,a) values
- Receive a sample (s,a,s'r)
- Consider your old estimate: $Q(s, a)$
- Consider your new sample estimate:

$$
\text { sample }=r+\gamma \max _{a^{\prime}} Q\left(s^{\prime}, a^{\prime}\right)
$$

- Incorporate the new estimate into a running average:

$$
Q(s, a) \leftarrow(1-\alpha) Q(s, a)+(\alpha)[\text { sample }]
$$

## Recap: Exploration / Exploitation

- Several schemes for action selection
- Simplest: random actions ( $\varepsilon$ greedy)
- Every time step, flip a coin
- With probability $\varepsilon$, act randomly
- With probability 1-ع, act according to current policy
- Problems with random actions?
- You do explore the space, but keep thrashing around once learning is done
- One solution: lower $\varepsilon$ over time
- Another solution: exploration functions


## Q-Learning: $\varepsilon$ Greedy



CURRENT Q-VALUES

## Q-Learning Final Solution

- Q-learning produces tables of q-values:


Q-VALUES AFTER 1000 EPISODES

## Q-Learning

- In realistic situations, we cannot possibly learn about every single state!
- Too many states to visit them all in training
- Too many states to hold the q-tables in memory
- Instead, we want to generalize:
- Learn about some small number of training states from experience
- Generalize that experience to new, similar states
- This is a fundamental idea in machine learning, and we'll see it over and over again


## Example: Pacman

- Let's say we discover through experience that this state is bad:
- In naïve q learning, we know nothing about related states and their q values:



## Feature-Based Representations

- Solution: describe a state using a vector of features (properties)
- Features are functions from states to real numbers (often 0/1) that capture important properties of the state
- Example features:
- Distance to closest ghost

- Distance to closest dot
- Number of ghosts
- 1 / (dist to dot) ${ }^{2}$
- Is Pacman in a tunnel? (0/1)
- ...... etc.
- Is it the exact state on this slide?
- Can also describe a q-state (s, a) with features (e.g. action moves closer to food)


## Function Approximation

$$
Q(s, a)=w_{1} f_{1}(s, a)+w_{2} f_{2}(s, a)+\ldots+w_{n} f_{n}(s, a)
$$

- Q-learning with linear q-functions:

$$
\begin{aligned}
& \text { transition }=\left(s, a, r, s^{\prime}\right) \\
& \text { difference }=\left[r+\gamma \max _{a^{\prime}} Q\left(s^{\prime}, a^{\prime}\right)\right]-Q(s, a)
\end{aligned}
$$

$$
Q(s, a) \leftarrow Q(s, a)+\alpha[\text { difference }] \quad \text { Exact Q's }
$$

$$
w_{i} \leftarrow w_{i}+\alpha \text { [difference] } f_{i}(s, a) \quad \text { Approximate Q's }
$$

- Intuitive interpretation:
- Adjust weights of active features
- E.g. if something unexpectedly bad happens, disprefer all states with that state's features
- Formal justification: online least squares


## Example: Q-Pacman

$$
\begin{gathered}
Q(s, a)=4.0 f_{D O T}(s, a)-1.0 f_{G S T}(s, a) \\
f_{D O T}(s, \text { NORTH })=0.5 \\
f_{G S T}(s, \text { NORTH })=1.0 \\
Q(s, a)=+1
\end{gathered}
$$



$$
\begin{gathered}
a=\mathrm{NORTH} \\
r=-500
\end{gathered}
$$

$$
\begin{gathered}
\text { correction }=-501 \\
w_{D O T} \leftarrow 4.0+\alpha[-501] 0.5 \\
w_{G S T} \leftarrow-1.0+\alpha[-501] 1.0
\end{gathered}
$$

$$
Q(s, a)=3.0 f_{D O T}(s, a)-3.0 f_{G S T}(s, a)
$$

## Linear Regression




## Prediction

$\widehat{y}=w_{0}+w_{1} f_{1}(x)$

## Prediction

$\widehat{y}_{i}=w_{0}+w_{1} f_{1}(x)+w_{2} f_{2}(x)$

## Ordinary Least Squares (OLS)

total error $=\sum_{i}\left(y_{i}-\widehat{y_{i}}\right)^{2}=\sum_{i}\left(y_{i}-\sum_{k} w_{k} f_{k}\left(x_{i}\right)\right)^{2}$


## Minimizing Error

Imagine we had only one point $x$ with features $f(x)$ :

$$
\begin{aligned}
\operatorname{error}(w) & =\frac{1}{2}\left(y-\sum_{k} w_{k} f_{k}(x)\right)^{2} \\
\frac{\partial \operatorname{error}(w)}{\partial w_{m}} & =-\left(y-\sum_{k} w_{k} f_{k}(x)\right) f_{m}(x) \\
w_{m} \leftarrow w_{m} & +\alpha\left(y-\sum_{k} w_{k} f_{k}(x)\right) f_{m}(x)
\end{aligned}
$$

Approximate q update:

$$
w_{m} \leftarrow w_{m}+\alpha\left[r+\gamma \max _{a} Q\left(s^{\prime}, a^{\prime}\right)-Q(s, a)\right] f_{m}(s, a)
$$

## Overfitting



## Which Algorithm?

Q-learning, no features, 50 learning trials:


## Which Algorithm?

Q-learning, no features, 1000 learning trials:


## Which Algorithm?

Q-learning, simple features, 50 learning trials:


## Policy Search*

- Problem: often the feature-based policies that work well aren't the ones that approximate $\mathrm{V} / \mathrm{Q}$ best
- E.g. your value functions from project 2 were probably horrible estimates of future rewards, but they still produced good decisions
- We'll see this distinction between modeling and prediction again later in the course
- Solution: learn the policy that maximizes rewards rather than the value that predicts rewards
- This is the idea behind policy search, such as what controlled the upside-down helicopter


## Policy Search*

- Simplest policy search:
- Start with an initial linear value function or q-function
- Nudge each feature weight up and down and see if your policy is better than before
- Problems:
- How do we tell the policy got better?
- Need to run many sample episodes!
- If there are a lot of features, this can be impractical


## Policy Search*

- Advanced policy search:
- Write a stochastic (soft) policy:

$$
\pi_{w}(s) \propto e^{\sum_{i} w_{i} f_{i}(s, a)}
$$

- Turns out you can efficiently approximate the derivative of the returns with respect to the parameters w (details in the book, optional material)
- Take uphill steps, recalculate derivatives, etc.


## Review: MDPs

- Markov decision processes:
- States S
- Actions A
- Transitions P(s'|s,a) (or T(s,a,s'))
- Rewards R(s,a,s') (and discount $\gamma$ )
- Start state dist. $b_{0}$



## Partially observable MDPs

- Markov decision processes:
- States S
- Actions A
- Transitions P(s'|s,a) (or T(s,a,s'))
- Rewards R(s,a,s') (and discount $\gamma$ )
- Start state distribution $b_{0}=P\left(s_{0}\right)$
- POMDPs, just add:

- Observations O
- Observation model P(o|s,a) (or O(s,a,o))


## A POMDP: Ghost Hunter



## POMDP Computations

- Sufficient statistic: belief states
- $b_{0}=P\left(s_{0}\right)$
- $b\left(s^{\prime}\right)=P\left(s^{\prime} \mid o, a, b\right)$
$=\frac{O\left(s^{\prime}, a, o\right) \sum_{s \in S} T\left(s, a, s^{\prime}\right) b(s)}{P(o \mid a, b)}$
- POMDPs search trees

- max nodes are belief states
- expectation nodes branch on possible observations
- (this is motivational; we will not discuss in detail)


## Probability Review

- Probability
- Random Variables
- Joint and Marginal Distributions
- Conditional Distribution
- Product Rule, Chain Rule, Bayes' Rule
- Inference
- You'll need all this stuff A LOT for the next few weeks, so make sure you go over it now!


## Inference in Ghostbusters

- A ghost is in the grid somewhere
- Sensor readings tell how close a square is to the ghost
- On the ghost: red
- 1 or 2 away: orange
- 3 or 4 away: yellow
- 5+ away: green
- Sensors are noisy, but we know P(Color | Distance)

| $P($ red \| 3) | $P($ orange \| 3) | $P($ yellow \| 3) | $P($ green \| 3) |
| :---: | :---: | :---: | :---: |
| 0.05 | 0.15 | 0.5 | 0.3 |

## Uncertainty

- General situation:
- Evidence: Agent knows certain things about the state of the world (e.g., sensor readings or symptoms)
- Hidden variables: Agent needs to reason about other aspects (e.g. where an object is or what disease is present)
- Model: Agent knows something about how the known variables relate to the unknown variables
- Probabilistic reasoning gives us a framework for managing our beliefs and knowledge

| 0.11 | 0.11 | 0.11 |
| :--- | :--- | :--- |
| 0.11 | 0.11 | 0.11 |
| 0.11 | 0.11 | 0.11 |


| 0.17 | 0.10 | 0.10 |
| :---: | :---: | :---: |
| 0.09 | 0.17 | 0.10 |
| 0.01 | 0.09 | 0.17 |



## Random Variables

- A random variable is some aspect of the world about which we (may) have uncertainty
- $\mathrm{R}=\mathrm{Is}$ it raining?
- $\mathrm{D}=$ How long will it take to drive to work?
- L = Where am I?
- We denote random variables with capital letters
- Random variables have domains
- R in \{true, false\}
- D in $[0, \infty)$
- Lin possible locations, maybe $\{(0,0),(0,1), \ldots\}$


## Probability Distributions

- Unobserved random variables have distributions

$P(T)$| $T$ | $P$ |
| :---: | :---: |
| warm | 0.5 |
| cold | 0.5 |

$P(W)$

| $W$ | $P$ |
| :---: | :---: |
| sun | 0.6 |
| rain | 0.1 |
| fog | 0.3 |
| meteor | 0.0 |

- A distribution is a TABLE of probabilities of values
- A probability (lower case value) is a single number

$$
P(W=\text { rain })=0.1 \quad P(\text { rain })=0.1
$$

- Must have:

$$
\forall x P(x) \geq 0 \quad \sum_{x} P(x)=1
$$

## Joint Distributions

- A joint distribution over a set of random variables: $X_{1}, X_{2}, \ldots X_{n}$ specifies a real number for each assignment (or outcome):

$$
\begin{aligned}
& P\left(X_{1}=x_{1}, X_{2}=x_{2}, \ldots X_{n}=x_{n}\right) \\
& P\left(x_{1}, x_{2}, \ldots x_{n}\right)
\end{aligned}
$$

$P(T, W)$

- Size of distribution if n variables with domain sizes d ?
- Must obey:

$$
P\left(x_{1}, x_{2}, \ldots x_{n}\right) \geq 0
$$

| T | W | P |
| :---: | :---: | :---: |
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

$$
\sum_{\left(x_{1}, x_{2}, \ldots x_{n}\right)} P\left(x_{1}, x_{2}, \ldots x_{n}\right)=1
$$

- A probabilistic model is a joint distribution over variables of interest
- For all but the smallest distributions, impractical to write out


## Events

- An event is a set E of outcomes

$$
P(E)=\sum_{\left(x_{1} \ldots x_{n}\right) \in E} P\left(x_{1} \ldots x_{n}\right)
$$

- From a joint distribution, we can calculate the probability of any event

| $T$ | $W$ | $P$ |
| :---: | :---: | :---: |
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

- Probability that it's hot AND sunny?
- Probability that it's hot?
- Probability that it's hot OR sunny?
- Typically, the events we care about are partial assignments, like $\mathrm{P}(\mathrm{T}=$ hot $)$


## Marginal Distributions

- Marginal distributions are sub-tables which eliminate variables
- Marginalization (summing out): Combine collapsed rows by adding

$$
P\left(X_{1}=x_{1}\right)=\sum_{x_{2}} P\left(X_{1}=x_{1}, X_{2}=x_{2}\right)
$$

$P(T)$

| $P(T, W)$ |  |  |
| :---: | :---: | :---: |
| T | W | P |
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

$$
\begin{aligned}
& P(t)=\sum_{w} P(t, w) \\
& P(w)=\sum_{t} P(t, w)
\end{aligned}
$$

| T | P |
| :---: | :---: |
| hot | 0.5 |
| cold | 0.5 |

$P(W)$

| $W$ | $P$ |
| :---: | :---: |
| sun | 0.6 |
| rain | 0.4 |

## Conditional Probabilities

- A simple relation between joint and conditional probabilities
- In fact, this is taken as the definition of a conditional probability

$$
P(a \mid b)=\frac{P(a, b)}{P(b)}
$$

$$
P(T, W)
$$



| $T$ | $W$ | $P$ |
| :---: | :---: | :---: |
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

$$
P(W=r \mid T=c)=? ? ?
$$

## Conditional Distributions

- Conditional distributions are probability distributions over some variables given fixed values of others

Conditional Distributions
Joint Distribution


| $P(T, W)$ |  |  |
| :---: | :---: | :---: |
| T | W | P |
| hot | sun | 0.4 |
| hot | rain | 0.1 |
| cold | sun | 0.2 |
| cold | rain | 0.3 |

## Normalization Trick

- A trick to get a whole conditional distribution at once:
- Select the joint probabilities matching the evidence
- Normalize the selection (make it sum to one)

| $P(T, W)$ |  |  | $\xrightarrow{\text { Select }}$ | $P(T, r)$ | Normalize |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| T | W | P |  |  |  |  |  |
| hot | sun | 0.4 | T | R | P | T | P |
| hot | rain | 0.1 | hot | rain | 0.1 | hot | 0.25 |
| cold | sun | 0.2 | cold | rain | 0.3 | cold | 0.75 |
| cold | rain | 0.3 |  |  |  |  |  |

- Why does this work? Sum of selection is $\mathrm{P}($ evidence)! ( $\mathrm{P}(\mathrm{r})$, here)

$$
P\left(x_{1} \mid x_{2}\right)=\frac{P\left(x_{1}, x_{2}\right)}{P\left(x_{2}\right)}=\frac{P\left(x_{1}, x_{2}\right)}{\sum_{x_{1}} P\left(x_{1}, x_{2}\right)}
$$

## Probabilistic Inference

- Probabilistic inference: compute a desired probability from other known probabilities (e.g. conditional from joint)
- We generally compute conditional probabilities
- P (on time | no reported accidents) $=0.90$
- These represent the agent's beliefs given the evidence
- Probabilities change with new evidence:
- $P$ (on time $\mid$ no accidents, 5 a.m.) $=0.95$
- $P$ (on time | no accidents, 5 a.m., raining) $=0.80$
- Observing new evidence causes beliefs to be updated


## Inference by Enumeration

- $\mathrm{P}($ sun $)$ ?
- $\mathrm{P}($ sun | winter)?
- $\mathrm{P}($ sun | winter, warm)?

| S | T | W | P |
| :---: | :---: | :---: | :---: |
| summer | hot | sun | 0.30 |
| summer | hot | rain | 0.05 |
| summer | cold | sun | 0.10 |
| summer | cold | rain | 0.05 |
| winter | hot | sun | 0.10 |
| winter | hot | rain | 0.05 |
| winter | cold | sun | 0.15 |
| winter | cold | rain | 0.20 |

## Inference by Enumeration

- General case:
- Evidence variables: $E_{1} \ldots E_{k}=e_{1} \ldots e_{k}$
- Query* variable: $Q$
- Hidden variables: $H_{1} \ldots H_{r}$
$X_{1}, X_{2}, \ldots X_{n}$
All variables
- We want: $P\left(Q \mid e_{1} \ldots e_{k}\right)$
- First, select the entries consistent with the evidence
- Second, sum out H to get joint of Query and evidence:

$$
P\left(Q, e_{1} \ldots e_{k}\right)=\sum_{h_{1} \ldots h_{r}} \underbrace{P\left(Q, h_{1} \ldots h_{r}, e_{1} \ldots e_{k}\right)}_{X_{1}, X_{2}, \ldots X_{n}}
$$

- Finally, normalize the remaining entries to conditionalize
- Obvious problems:
- Worst-case time complexity O(dn)
- Space complexity $O\left(d^{n}\right)$ to store the joint distribution


## The Product Rule

- Sometimes have conditional distributions but want the joint

$$
P(x \mid y)=\frac{P(x, y)}{P(y)} \Longleftrightarrow P(x, y)=P(x \mid y) P(y)
$$

- Example:

$$
P(D \mid W)
$$

$P(D, W)$

| $P(W)$ |  |  |  |  | $\rangle$ | D | W | P |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | D |  | P |  |  |  |  |
|  |  | wet |  |  |  | sun | 0.08 |  |
| R | P |  | wet | sun |  | 0.9 | dry | sun | 0.72 |
| sun | 0.8 | dr | rain | 0.9 |  | wet | rain | 0.14 |
| rain | 0.2 | wet | rain | 0.7 |  | dry | rain | 0.06 |

## The Chain Rule

- More generally, can always write any joint distribution as an incremental product of conditional distributions

$$
\begin{aligned}
& P\left(x_{1}, x_{2}, x_{3}\right)=P\left(x_{1}\right) P\left(x_{2} \mid x_{1}\right) P\left(x_{3} \mid x_{1}, x_{2}\right) \\
& P\left(x_{1}, x_{2}, \ldots x_{n}\right)=\prod_{i} P\left(x_{i} \mid x_{1} \ldots x_{i-1}\right)
\end{aligned}
$$

- Why is this always true?


## Bayes' Rule

- Two ways to factor a joint distribution over two variables:

$$
P(x, y)=P(x \mid y) P(y)=P(y \mid x) P(x)
$$

- Dividing, we get:

$$
P(x \mid y)=\frac{P(y \mid x)}{P(y)} P(x)
$$

- Why is this at all helpful?
- Lets us build one conditional from its reverse
- Often one conditional is tricky but the other one is simple
- Foundation of many systems we'll see later (e.g. ASR, MT)
- In the running for most important Al equation!


## Inference with Bayes' Rule

- Example: Diagnostic probability from causal probability:

$$
P(\text { Cause } \mid \text { Effect })=\frac{P(\text { Effect } \mid \text { Cause }) P(\text { Cause })}{P(\text { Effect })}
$$

- Example:
- m is meningitis, s is stiff neck $\left.\begin{array}{l}P(s \mid m)=0.8 \\ \\ P(m)=0.0001 \\ P(s)=0.1\end{array}\right] \begin{aligned} & \text { Example } \\ & \text { givens }\end{aligned}$

$$
P(m \mid s)=\frac{P(s \mid m) P(m)}{P(s)}=\frac{0.8 \times 0.0001}{0.1}=0.0008
$$

- Note: posterior probability of meningitis still very small
- Note: you should still get stiff necks checked out! Why?


## Ghostbusters, Revisited

- Let's say we have two distributions:
- Prior distribution over ghost location: P(G)
- Let's say this is uniform
- Sensor reading model: $P(R \mid G)$
- Given: we know what our sensors do
- $\mathrm{R}=$ reading color measured at $(1,1)$
- E.g. $P(R=$ yellow $\mid G=(1,1))=0.1$
- We can calculate the posterior distribution $\mathrm{P}(\mathrm{G} \mid \mathrm{r})$ over ghost locations given a reading using Bayes' rule:

$$
P(g \mid r) \propto P(r \mid g) P(g)
$$

| 0.11 | 0.11 | 0.11 |
| :--- | :--- | :--- |
| 0.11 | 0.11 | 0.11 |
| 0.11 | 0.11 | 0.11 |


| 0.17 | 0.10 | 0.10 |
| :--- | :--- | :--- |
| 0.09 | 0.17 | 0.10 |
| $<0.01$ | 0.09 | 0.17 |

