

PCA

Machine Learning – 10701/15781

Carlos Guestrin

Carnegie Mellon University

November 28th, 2007

©2005-2007 Carlos Guestrin

1

Lower dimensional projections

- Rather than picking a subset of the features, we can new features that are combinations of existing features

e.g., feature selection:
use x_1, x_7, x_{11}

low. dim. proj.
 $\tilde{x} = 0.1x_1 + 0.7x_2 - 0.35x_3 \dots$

- Let's see this in the unsupervised setting
 - just X, but no Y

©2005-2007 Carlos Guestrin

2

PCA finds projection that minimizes reconstruction error

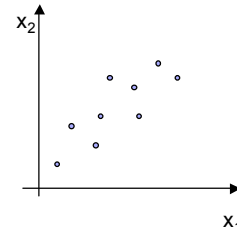
- Given m data points: $\mathbf{x}^i = (x_1^i, \dots, x_n^i)$, $i=1 \dots m$
- Will represent each point as a projection:

$$\square \hat{\mathbf{x}}^i = \bar{\mathbf{x}} + \sum_{j=1}^k z_j^i \mathbf{u}_j \quad \text{where: } \bar{\mathbf{x}} = \frac{1}{m} \sum_{i=1}^m \mathbf{x}^i \quad \text{and} \quad z_j^i = (\mathbf{x}^i - \bar{\mathbf{x}}) \cdot \mathbf{u}_j$$

- PCA:

- Given $k \ll n$, find $(\mathbf{u}_1, \dots, \mathbf{u}_k)$ minimizing reconstruction error:

$$error_k = \sum_{i=1}^m (\mathbf{x}^i - \hat{\mathbf{x}}^i)^2$$



Understanding the reconstruction error

- Note that \mathbf{x}^i can be represented exactly by n -dimensional projection:

$$\mathbf{x}^i = \bar{\mathbf{x}} + \sum_{j=1}^n z_j^i \mathbf{u}_j$$

$$\hat{\mathbf{x}}^i = \bar{\mathbf{x}} + \sum_{j=1}^k z_j^i \mathbf{u}_j$$

$$z_j^i = (\mathbf{x}^i - \bar{\mathbf{x}}) \cdot \mathbf{u}_j$$

- Given $k \ll n$, find $(\mathbf{u}_1, \dots, \mathbf{u}_k)$ minimizing reconstruction error:

$$error_k = \sum_{i=1}^m (\mathbf{x}^i - \hat{\mathbf{x}}^i)^2$$

- Rewriting error:

Reconstruction error and covariance matrix

$$error_k = \sum_{i=1}^m \sum_{j=k+1}^n [\mathbf{u}_j \cdot (\mathbf{x}^i - \bar{\mathbf{x}})]^2$$

$$\Sigma = \frac{1}{m} \sum_{i=1}^m (\mathbf{x}^i - \bar{\mathbf{x}})(\mathbf{x}^i - \bar{\mathbf{x}})^T$$

Minimizing reconstruction error and eigen vectors

- Minimizing reconstruction error equivalent to picking orthonormal basis $(\mathbf{u}_1, \dots, \mathbf{u}_n)$ minimizing:

$$error_k = \sum_{j=k+1}^n \mathbf{u}_j^T \Sigma \mathbf{u}_j$$

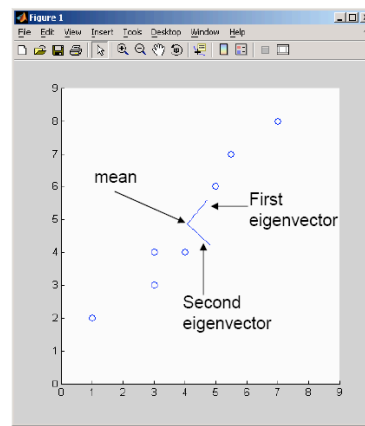
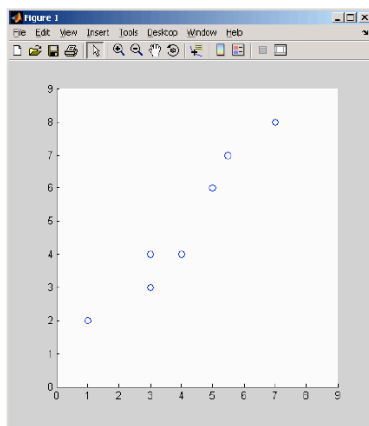
- Eigen vector:
- Minimizing reconstruction error equivalent to picking $(\mathbf{u}_{k+1}, \dots, \mathbf{u}_n)$ to be eigen vectors with smallest eigen values

Basic PCA algorithm

- Start from m by n data matrix \mathbf{X}
- **Recenter**: subtract mean from each row of \mathbf{X}
 - $\mathbf{X}_c \leftarrow \mathbf{X} - \bar{\mathbf{X}}$
- **Compute covariance matrix**:
 - $\Sigma \leftarrow 1/m \mathbf{X}_c^T \mathbf{X}_c$
- Find **eigen vectors and values** of Σ
- **Principal components**: k eigen vectors with highest eigen values

PCA example

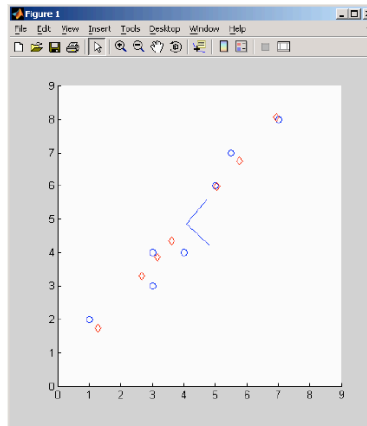
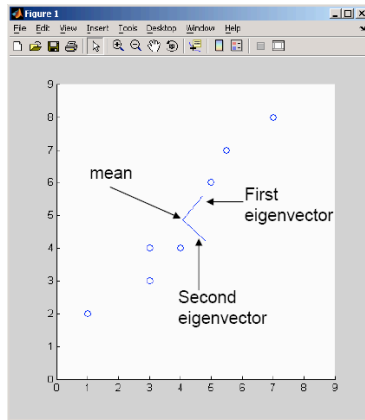
$$\hat{\mathbf{x}}^i = \bar{\mathbf{x}} + \sum_{j=1}^k z_j^i \mathbf{u}_j$$



PCA example – reconstruction

$$\hat{\mathbf{x}}^i = \bar{\mathbf{x}} + \sum_{j=1}^k z_j^i \mathbf{u}_j$$

only used first principal component



Eigenfaces [Turk, Pentland '91]

■ Input images:



■ Principal components:



Eigenfaces reconstruction

- Each image corresponds to adding 8 principal components:



Scaling up

- Covariance matrix can be really big!
 - Σ is n by n
 - 10000 features ! $|\Sigma|$
 - finding eigenvectors is very slow...
- Use singular value decomposition (SVD)
 - finds to k eigenvectors
 - great implementations available, e.g., Matlab svd

SVD

- Write $\mathbf{X} = \mathbf{W} \mathbf{S} \mathbf{V}^T$
 - $\mathbf{X} \leftarrow$ data matrix, one row per datapoint
 - $\mathbf{W} \leftarrow$ weight matrix, one row per datapoint – coordinate of \mathbf{x}^i in eigenspace
 - $\mathbf{S} \leftarrow$ singular value matrix, diagonal matrix
 - in our setting each entry is eigenvalue λ_j
 - $\mathbf{V}^T \leftarrow$ singular vector matrix
 - in our setting each row is eigenvector \mathbf{v}_j

PCA using SVD algorithm


- Start from m by n data matrix \mathbf{X}
- **Recenter**: subtract mean from each row of \mathbf{X}
 - $\mathbf{X}_c \leftarrow \mathbf{X} - \bar{\mathbf{X}}$
- Call SVD algorithm on \mathbf{X}_c – ask for k singular vectors
- **Principal components**: k singular vectors with highest singular values (rows of \mathbf{V}^T)
 - **Coefficients** become:

What you need to know

- Dimensionality reduction
 - why and when it's important
- Simple feature selection
- Principal component analysis
 - minimizing reconstruction error
 - relationship to covariance matrix and eigenvectors
 - using SVD

Announcements

- University Course Assessments
 - Please, please, please, please, please, please, please, please, please, please, please, please, please, please...
- Last lecture:
 - Thursday, 11/29, 4:40-6:30pm, Wean 7500



Markov Decision Processes (MDPs)

Machine Learning – 10701/15781

Carlos Guestrin

Carnegie Mellon University

November 28th, 2007

19

Thus far this semester



- Regression:
- Classification:
- Density estimation:

Learning to act



[Ng et al. '05]

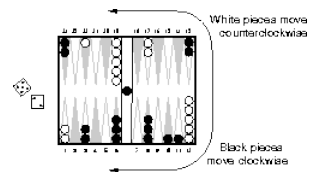
- Reinforcement learning
- An agent
 - Makes sensor observations
 - Must select action
 - Receives rewards
 - positive for “good” states
 - negative for “bad” states

Learning to play backgammon

[Tesauro '95]



- Combines reinforcement learning with neural networks
- Played 300,000 games against itself
- Achieved grandmaster level!



Roadmap to learning about reinforcement learning

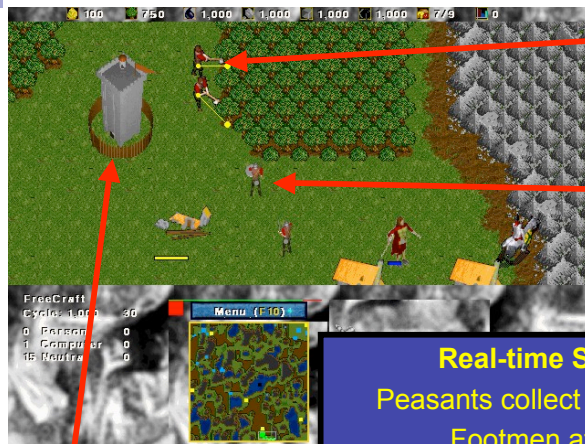
■ When we learned about Bayes nets:

- First talked about formal framework:
 - representation
 - inference
- Then learning for BNs

■ For reinforcement learning:

- Formal framework
 - Markov decision processes
- Then learning

FreeCraft



peasant

footman

building

Real-time Strategy Game

Peasants collect resources and build

Footmen attack enemies

Buildings train peasants and footmen

States and actions

- State space:
 - Joint state \mathbf{x} of entire system
- Action space:
 - Joint action $\mathbf{a} = \{a_1, \dots, a_n\}$ for all agents



States change over time

- Like an HMM, state changes over time
- Next state depends on current state and action selected
 - e.g., action="build castle" likely to lead to a state where you have a castle
- Transition model:
 - Dynamics of the entire system $P(\mathbf{x}'|\mathbf{x},\mathbf{a})$



Some states and actions are better than others

- Each state \mathbf{x} is associated with a reward
 - positive reward for successful attack
 - negative for loss
- Reward function:
 - Total reward $R(\mathbf{x})$



Markov Decision Process (MDP) Representation

- State space:
 - Joint state \mathbf{x} of entire system
- Action space:
 - Joint action $\mathbf{a} = \{a_1, \dots, a_n\}$ for all agents
- Reward function:
 - Total reward $R(\mathbf{x}, \mathbf{a})$
 - sometimes reward can depend on action
- Transition model:
 - Dynamics of the entire system $P(\mathbf{x}' | \mathbf{x}, \mathbf{a})$

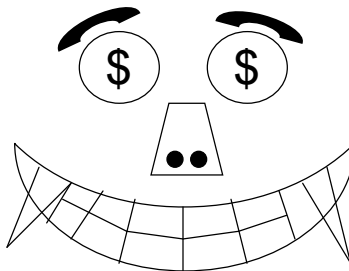


Discounted Rewards

An assistant professor gets paid, say, 20K per year.

How much, in total, will the A.P. earn in their life?

$20 + 20 + 20 + 20 + 20 + \dots = \text{Infinity}$



What's wrong with this argument?

Discounted Rewards

“A reward (payment) in the future is not worth quite as much as a reward now.”

- ☐ Because of chance of obliteration
- ☐ Because of inflation

Example:

Being promised \$10,000 next year is worth only 90% as much as receiving \$10,000 right now.

Assuming payment n years in future is worth only $(0.9)^n$ of payment now, what is the AP's **Future Discounted Sum of Rewards** ?

Discount Factors

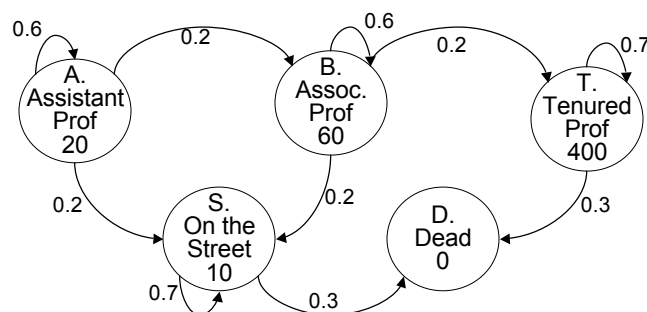
People in economics and probabilistic decision-making do this all the time.

The “Discounted sum of future rewards” using discount factor γ is

$$\begin{aligned}
 & (\text{reward now}) + \\
 & \gamma (\text{reward in 1 time step}) + \\
 & \gamma^2 (\text{reward in 2 time steps}) + \\
 & \gamma^3 (\text{reward in 3 time steps}) + \\
 & \vdots \\
 & \vdots \quad (\text{infinite sum})
 \end{aligned}$$

The Academic Life

Assume Discount Factor $\gamma = 0.9$



Define:

V_A = Expected discounted future rewards starting in state A

V_B = Expected discounted future rewards starting in state B

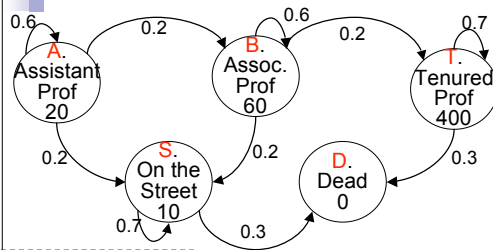
V_T = “ “ “ “ “ “ “ T

V_S = “ “ “ “ “ “ “ S

V_D = “ “ “ “ “ “ “ D

How do we compute V_A, V_B, V_T, V_S, V_D ?

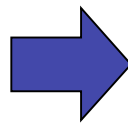
Computing the Future Rewards of an Academic



Assume Discount Factor $\gamma = 0.9$

Policy

Policy: $\pi(\mathbf{x}) = \mathbf{a}$



At state \mathbf{x} ,
action \mathbf{a} for all
agents



$\pi(\mathbf{x}_0) =$ both peasants get wood



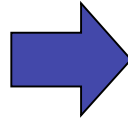
$\pi(\mathbf{x}_1) =$ one peasant builds
barrack, other gets gold



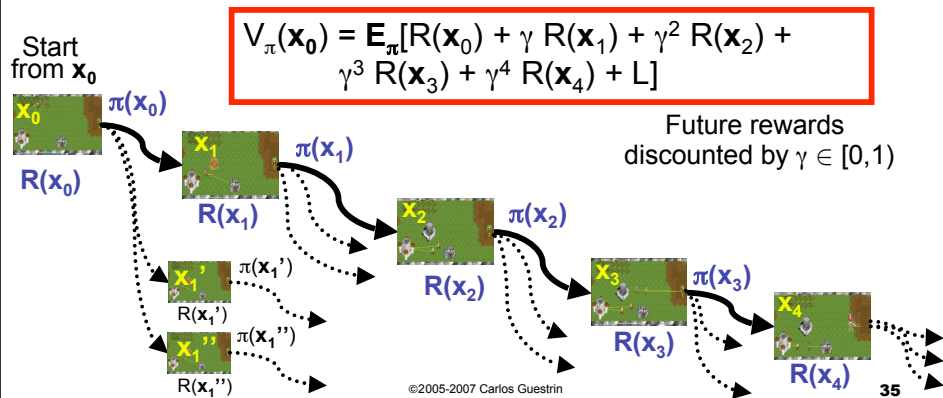
$\pi(\mathbf{x}_2) =$ peasants get gold,
footmen attack

Value of Policy

Value: $V_{\pi}(\mathbf{x})$



Expected long-term reward starting from \mathbf{x}



Computing the value of a policy

$$V_{\pi}(\mathbf{x}_0) = \mathbf{E}_{\pi}[R(\mathbf{x}_0) + \gamma R(\mathbf{x}_1) + \gamma^2 R(\mathbf{x}_2) + \gamma^3 R(\mathbf{x}_3) + \gamma^4 R(\mathbf{x}_4) + L]$$

- Discounted value of a state:

- value of starting from \mathbf{x}_0 and continuing with policy π from then on

$$\begin{aligned} V_{\pi}(x_0) &= E_{\pi}[R(x_0) + \gamma R(x_1) + \gamma^2 R(x_2) + \gamma^3 R(x_3) + \dots] \\ &= E_{\pi}\left[\sum_{t=0}^{\infty} \gamma^t R(x_t)\right] \end{aligned}$$

- A recursion!

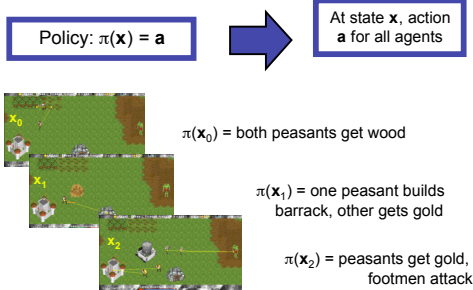
Simple approach for computing the value of a policy: Iteratively

$$V_{\pi}(x) = R(x) + \gamma \sum_{x'} P(x' | x, a = \pi(x)) V_{\pi}(x')$$

- Can solve using a simple convergent iterative approach: (a.k.a. dynamic programming)
 - Start with some guess V_0
 - Iteratively say:
 - $V_{t+1} = R + \gamma P_{\pi} V_t$
 - Stop when $\|V_{t+1} - V_t\|_{\infty} \leq \epsilon$
 - means that $\|V_{\pi} - V_{t+1}\|_{\infty} \leq \epsilon/(1-\gamma)$

But we want to learn a Policy

- So far, told you how good a policy is...
- But how can we choose the best policy???
- Suppose there was only one time step:
 - world is about to end!!!
 - select action that maximizes reward!



Unrolling the recursion

- Choose actions that lead to best value in the long run
 - Optimal value policy achieves optimal value V^*

$$V^*(x_0) = \max_{a_0} R(x_0, a_0) + \gamma E_{a_0} [\max_{a_1} R(x_1) + \gamma^2 E_{a_1} [\max_{a_2} R(x_2) + \dots]]$$

Bellman equation

- Evaluating policy π :

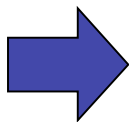
$$V_\pi(x) = R(x) + \gamma \sum_{x'} P(x' | x, a = \pi(x)) V_\pi(x')$$

- Computing the optimal value V^* - Bellman equation

$$V^*(\mathbf{x}) = \max_{\mathbf{a}} R(\mathbf{x}, \mathbf{a}) + \gamma \sum_{\mathbf{x}'} P(\mathbf{x}' | \mathbf{x}, \mathbf{a}) V^*(\mathbf{x}')$$

Optimal Long-term Plan

Optimal value
function $V^*(\mathbf{x})$



Optimal Policy: $\pi^*(\mathbf{x})$

Optimal policy:

$$\pi^*(\mathbf{x}) = \arg \max_a R(\mathbf{x}, \mathbf{a}) + \gamma \sum_{\mathbf{x}'} P(\mathbf{x}' | \mathbf{x}, \mathbf{a}) V^*(\mathbf{x}')$$

Interesting fact – Unique value

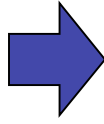
$$V^*(\mathbf{x}) = \max_a R(\mathbf{x}, \mathbf{a}) + \gamma \sum_{\mathbf{x}'} P(\mathbf{x}' | \mathbf{x}, \mathbf{a}) V^*(\mathbf{x}')$$

- *Slightly surprising fact:* There is only one V^* that solves Bellman equation!
 - there may be many optimal policies that achieve V^*
- *Surprising fact:* optimal policies are good everywhere!!!

$$V_{\pi^*}(x) \geq V_{\pi}(x), \quad \forall x, \quad \forall \pi$$

Solving an MDP

Solve
Bellman
equation



Optimal
value $V^*(\mathbf{x})$



Optimal
policy $\pi^*(\mathbf{x})$

$$V^*(\mathbf{x}) = \max_{\mathbf{a}} R(\mathbf{x}, \mathbf{a}) + \gamma \sum_{\mathbf{x}'} P(\mathbf{x}' | \mathbf{x}, \mathbf{a}) V^*(\mathbf{x}')$$

Bellman equation is non-linear!!!

Many algorithms solve the Bellman equations:

- Policy iteration [Howard '60, Bellman '57]
- Value iteration [Bellman '57]
- Linear programming [Mann '60]
- ...

Value iteration (a.k.a. dynamic programming) – the simplest of all

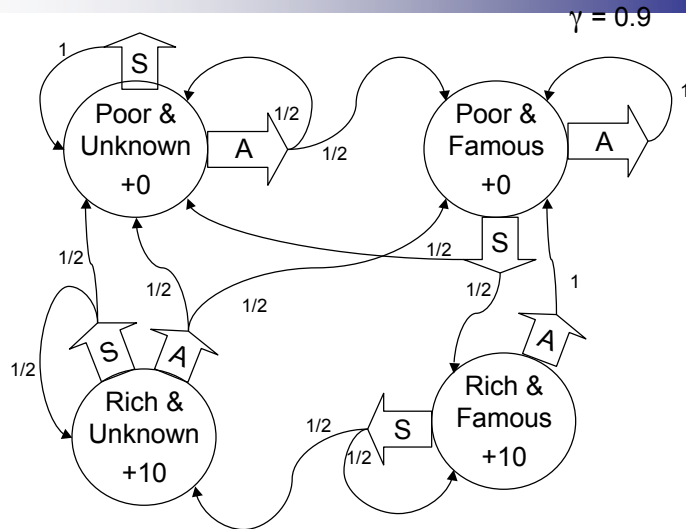
$$V^*(\mathbf{x}) = \max_{\mathbf{a}} R(\mathbf{x}, \mathbf{a}) + \gamma \sum_{\mathbf{x}'} P(\mathbf{x}' | \mathbf{x}, \mathbf{a}) V^*(\mathbf{x}')$$

- Start with some guess V_0
- Iteratively say:
 - $V_{t+1}(\mathbf{x}) = \max_{\mathbf{a}} R(\mathbf{x}, \mathbf{a}) + \gamma \sum_{\mathbf{x}'} P(\mathbf{x}' | \mathbf{x}, \mathbf{a}) V_t(\mathbf{x}')$
- Stop when $\|V_{t+1} - V_t\|_{\infty} \leq \epsilon$
 - means that $\|V^* - V_{t+1}\|_{\infty} \leq \epsilon / (1 - \gamma)$

A simple example

You run a startup company.

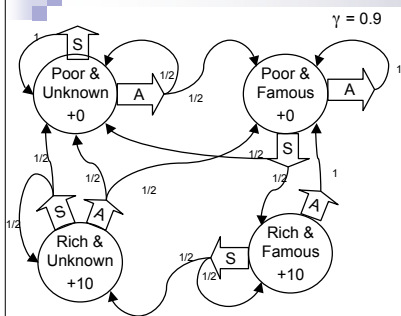
In every state you must choose between Saving money or Advertising.



©2005-2007 Carlos Guestrin

45

Let's compute $V_t(x)$ for our example



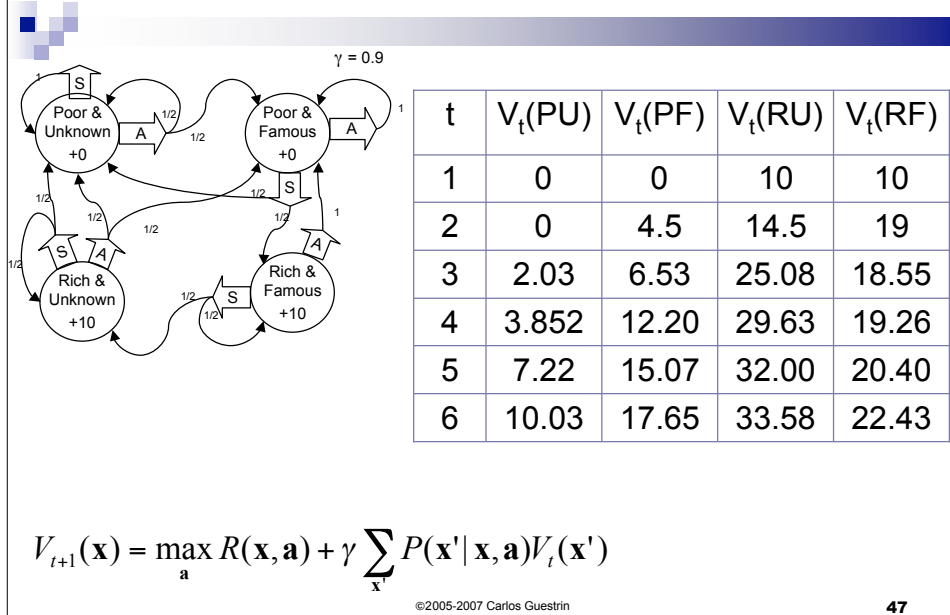
t	$V_t(\text{PU})$	$V_t(\text{PF})$	$V_t(\text{RU})$	$V_t(\text{RF})$
1				
2				
3				
4				
5				
6				

$$V_{t+1}(\mathbf{x}) = \max_{\mathbf{a}} R(\mathbf{x}, \mathbf{a}) + \gamma \sum_{\mathbf{x}'} P(\mathbf{x}' | \mathbf{x}, \mathbf{a}) V_t(\mathbf{x}')$$

©2005-2007 Carlos Guestrin

46

Let's compute $V_t(x)$ for our example



47

What you need to know

- What's a Markov decision process
 - state, actions, transitions, rewards
 - a policy
 - value function for a policy
 - computing V_{π}
- Optimal value function and optimal policy
 - Bellman equation
- Solving Bellman equation
 - with value iteration, policy iteration and linear programming

Acknowledgment



- This lecture contains some material from Andrew Moore's excellent collection of ML tutorials:

- <http://www.cs.cmu.edu/~awm/tutorials>