CS 188: Artificial Intelligence

Hidden Markov Models

University of California, Berkeley

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

Reasoning over Time or Space

- Often, we want to reason about a sequence of observations
  - Speech recognition
  - Robot localization
  - User attention
  - Medical monitoring
  - Language processing or generation
- Need to introduce time (or space) into our models

Today's Topics

- Very quick probability recap
- **Markov Chains** & their **Stationary Distributions**
- **Hidden Markov Models** (HMMs) formulation
- Preview of **Filtering** with HMMs

Probability Recap

- Conditional probability $$P(x|y) = \frac{P(x,y)}{P(y)}$$
- Marginal probability $$P(x) = \sum_y P(x,y)$$
- Product rule $$P(x,y) = P(x|y)P(y)$$
- Chain rule $$\begin{aligned} P(X_1, X_2, \ldots X_n) &= P(X_1)P(X_2|X_1)P(X_3|X_1,X_2)\ldots \\ &= \prod_{i=1}^{n} P(X_i|X_1,\ldots,X_{i-1}) \end{aligned}$$

Today's Topics

- Very quick probability recap
- **Markov Chains** & their **Stationary Distributions**
- **Hidden Markov Models** (HMMs) formulation
- Preview of **Filtering** with HMMs

Markov Models

- Value of X at a given time is called the state

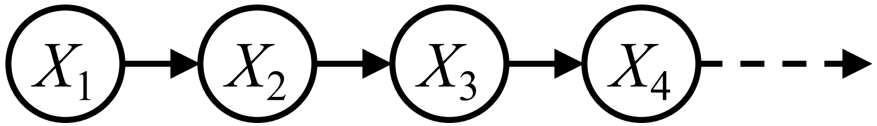

$$P(X_1) \qquad\qquad P(X_t|X_{t-1})$$

- Parameters: called transition probabilities or dynamics, specify how the state evolves over time (also, initial state probabilities)
- Stationarity assumption: transition probabilities the same at all times
- Same as MDP transition model, but no choice of action
- A "growable" BN (can always use BN methods if we truncate to fixed length)

Conditional Independence

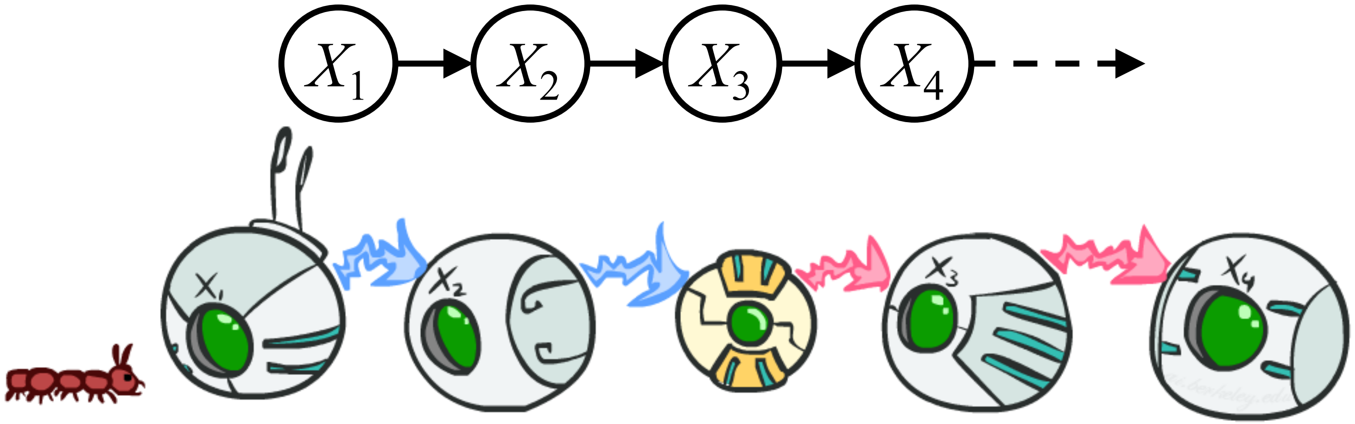

- Basic conditional independence:
  - Past and future independent given the present
  - Each time step only depends on the previous
  - This is called the (first order) Markov property

Example Markov Chain: Weather

- States: X = {rain, sun}

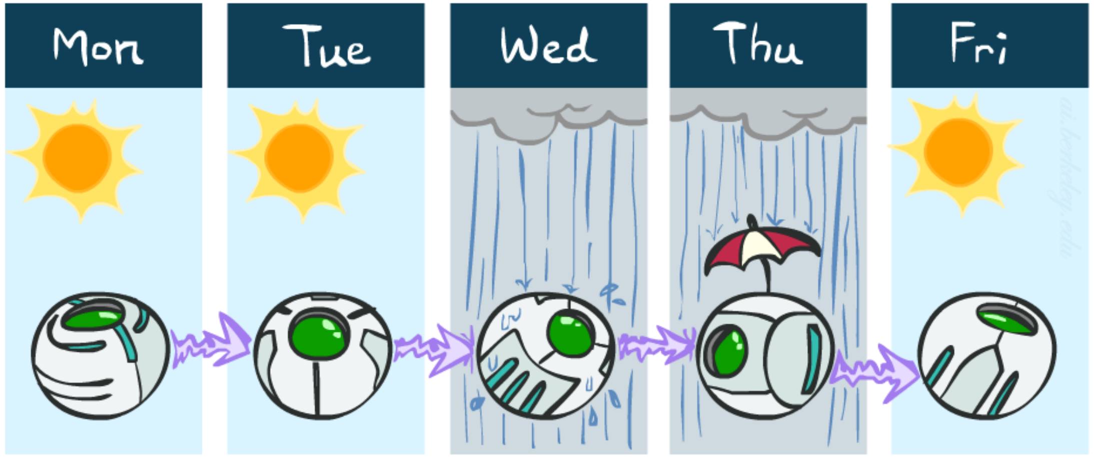

- Initial distribution: 1.0 sun

- CPT $P(X_t \mid X_{t-1})$:

X_{t-1}	X_t	$P(X_t \mid X_{t-1})$
sun	sun	0.9
sun	rain	0.1
rain	sun	0.3
rain	rain	0.7

Two new ways of representing the same CPT

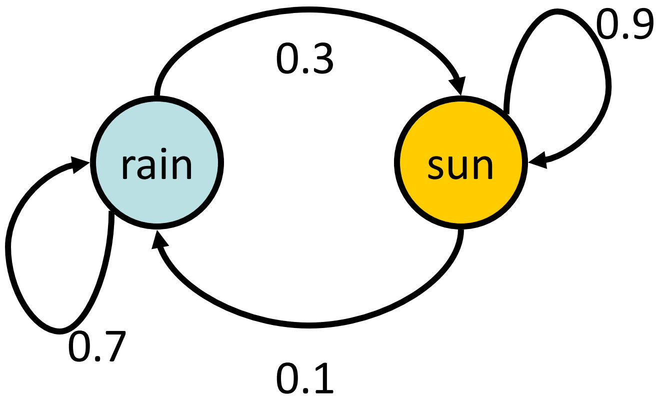

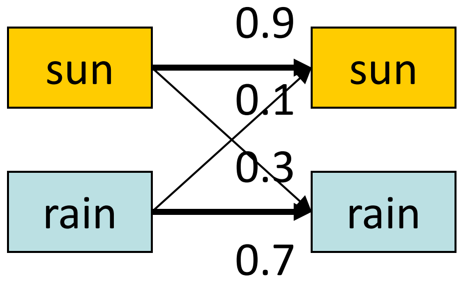

Example Markov Chain: Weather

- Initial distribution: 1.0 sun
  - We know: $P(X_1)$ $\quad P(X_t|X_{t-1})$

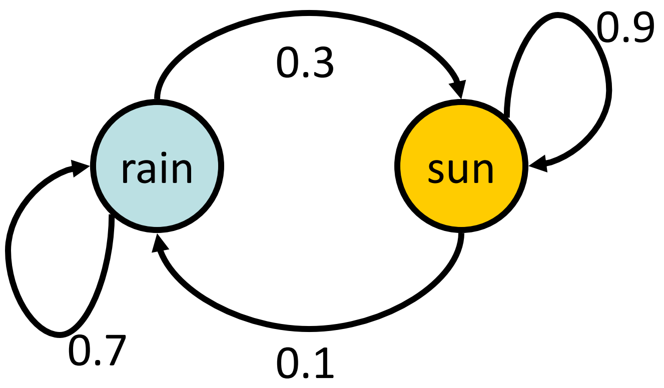

- What is the probability distribution after one step?

$$P(X_2 = sun) = \sum_{x_1} P(x_1, X_2 = sun) = \sum_{x_1} P(X_2 = sun|x_1)P(x_1)$$

$$= \quad P(X_2 = \text{sun}|X_1 = \text{sun})P(X_1 = \text{sun}) + P(X_2 = \text{sun}|X_1 = \text{rain})P(X_1 = \text{rain})$$

$$0.9 \cdot 1.0 + 0.3 \cdot 0.0 = 0.9$$

Mini-Forward Algorithm

- Question: What's P(X) on some day t?

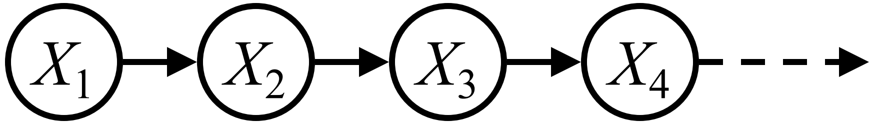

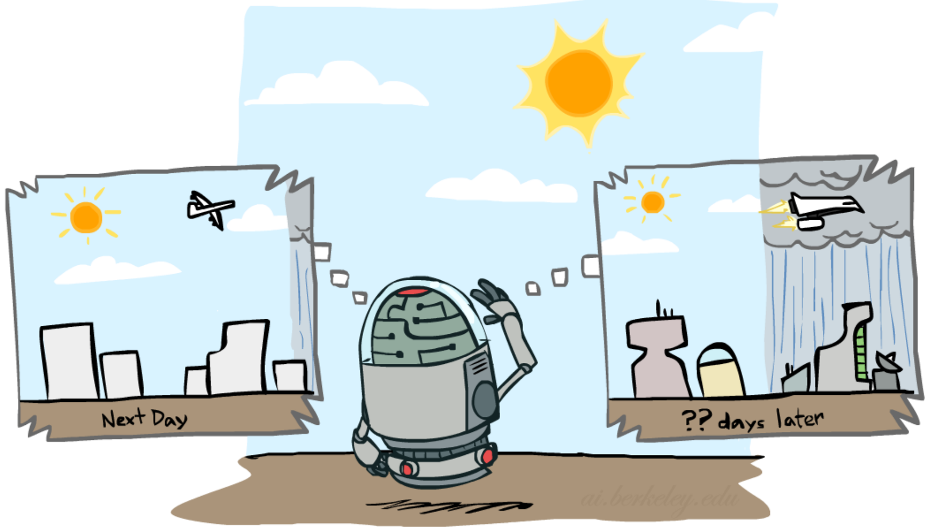

$$P(X_1) = \text{known}$$

$$P(x_t) = \sum_{x_{t-1}} P(x_{t-1}, x_t)$$
$$= \sum_{x_{t-1}} P(x_t \mid x_{t-1}) P(x_{t-1})$$

Forward simulation

Example Run of Mini-Forward Algorithm

- From initial observation of sun

$$\underbrace{\left\langle \begin{matrix} 1.0 \\ 0.0 \end{matrix} \right\rangle}_{P(X_1)} \quad \underbrace{\left\langle \begin{matrix} 0.9 \\ 0.1 \end{matrix} \right\rangle}_{P(X_2)} \quad \underbrace{\left\langle \begin{matrix} 0.84 \\ 0.16 \end{matrix} \right\rangle}_{P(X_3)} \quad \underbrace{\left\langle \begin{matrix} 0.804 \\ 0.196 \end{matrix} \right\rangle}_{P(X_4)} \Rightarrow \underbrace{\left\langle \begin{matrix} 0.75 \\ 0.25 \end{matrix} \right\rangle}_{P(X_\infty)}$$

- From initial observation of rain

$$\underbrace{\left\langle \begin{matrix} 0.0 \\ 1.0 \end{matrix} \right\rangle}_{P(X_1)} \quad \underbrace{\left\langle \begin{matrix} 0.3 \\ 0.7 \end{matrix} \right\rangle}_{P(X_2)} \quad \underbrace{\left\langle \begin{matrix} 0.48 \\ 0.52 \end{matrix} \right\rangle}_{P(X_3)} \quad \underbrace{\left\langle \begin{matrix} 0.588 \\ 0.412 \end{matrix} \right\rangle}_{P(X_4)} \Rightarrow \underbrace{\left\langle \begin{matrix} 0.75 \\ 0.25 \end{matrix} \right\rangle}_{P(X_\infty)}$$

- From yet another initial distribution $P(X_1)$:

$$\underbrace{\left\langle \begin{matrix} p \\ 1-p \end{matrix} \right\rangle}_{P(X_1)} \quad \cdots \quad \Rightarrow \underbrace{\left\langle \begin{matrix} 0.75 \\ 0.25 \end{matrix} \right\rangle}_{P(X_\infty)}$$

[Demo: L13D1,2,3]

Stationary Distributions

- For most chains:
  - Influence of the initial distribution gets less and less over time.
  - The distribution we end up in is independent of the initial distribution
- Stationary distribution:
  - The distribution we end up with is called the stationary distribution P_∞ of the chain
  - It satisfies

$$P_\infty(X) = P_{\infty+1}(X) = \sum_x P(X|x)P_\infty(x)$$

Video of Demo Ghostbusters Basic Dynamics

$$P(x_t) = \sum_{x_{t-1}} P(x_t \mid x_{t-1}) P(x_{t-1})$$

Video of Demo Ghostbusters Circular Dynamics

$$P(x_t) = \sum_{x_{t-1}} P(x_t \mid x_{t-1}) P(x_{t-1})$$

Video of Demo Ghostbusters Whirlpool Dynamics

$$P(x_t) = \sum_{x_{t-1}} P(x_t \mid x_{t-1}) P(x_{t-1})$$

Example: Stationary Distributions

- Question: What's P(X) at time t = infinity?

$$X_1 \rightarrow X_2 \rightarrow X_3 \rightarrow X_4 \dashrightarrow$$

$$P_\infty(X) = P_{\infty+1}(X) = \sum_x P(X|x)P_\infty(x)$$

$$P_\infty(sun) = P(sun|sun)P_\infty(sun) + P(sun|rain)P_\infty(rain)$$
$$P_\infty(rain) = P(rain|sun)P_\infty(sun) + P(rain|rain)P_\infty(rain)$$

$$P_\infty(sun) = 0.9P_\infty(sun) + 0.3P_\infty(rain)$$
$$P_\infty(rain) = 0.1P_\infty(sun) + 0.7P_\infty(rain)$$

$$P_\infty(sun) = 3P_\infty(rain)$$
$$P_\infty(rain) = 1/3P_\infty(sun)$$

Also: $P_\infty(sun) + P_\infty(rain) = 1$

$$\Rightarrow \quad P_\infty(sun) = 3/4 \qquad P_\infty(rain) = 1/4$$

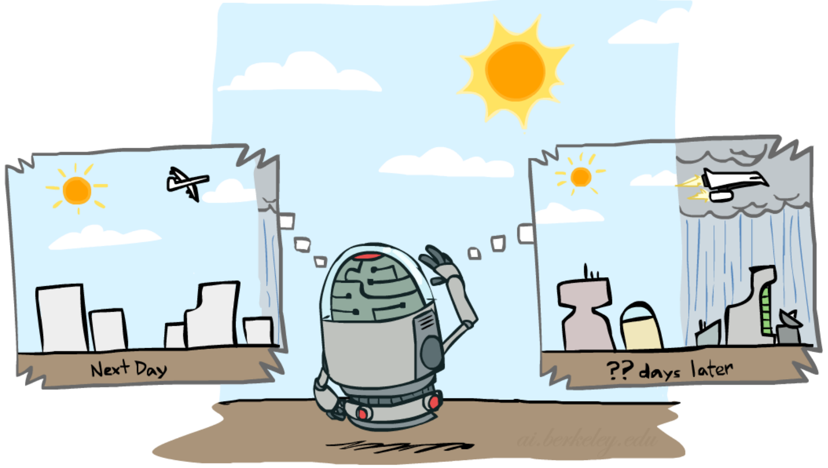

X_{t-1}	X_t	$P(X_t \mid X_{t-1})$
sun	sun	0.9
sun	rain	0.1
rain	sun	0.3
rain	rain	0.7

- Alternatively: run simulation for a long (ideally infinite) time

Application of Stationary Distribution: Diffusion Models

- Text-based image/art generation
  - Dall-E 2, Imagen, StableDiffusion, MidJourney, …

[Christian Beltrami/MidJourney]

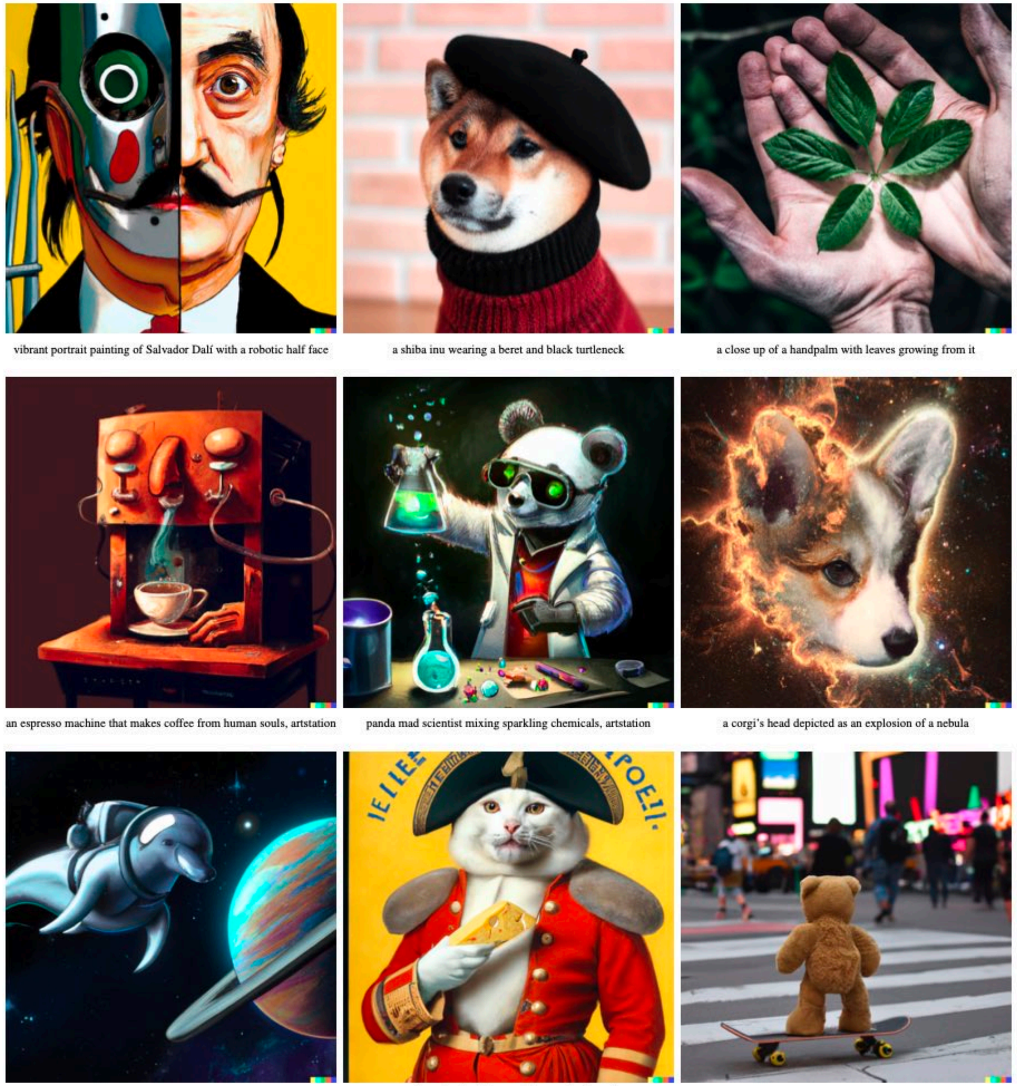

Figure 1: Selected 1024 × 1024 samples from a production version of our model. [OpenAI]

Application of Stationary Distribution: Diffusion Models

- Simulate (learned) Markov chain to reach stationary distribution of plausible images

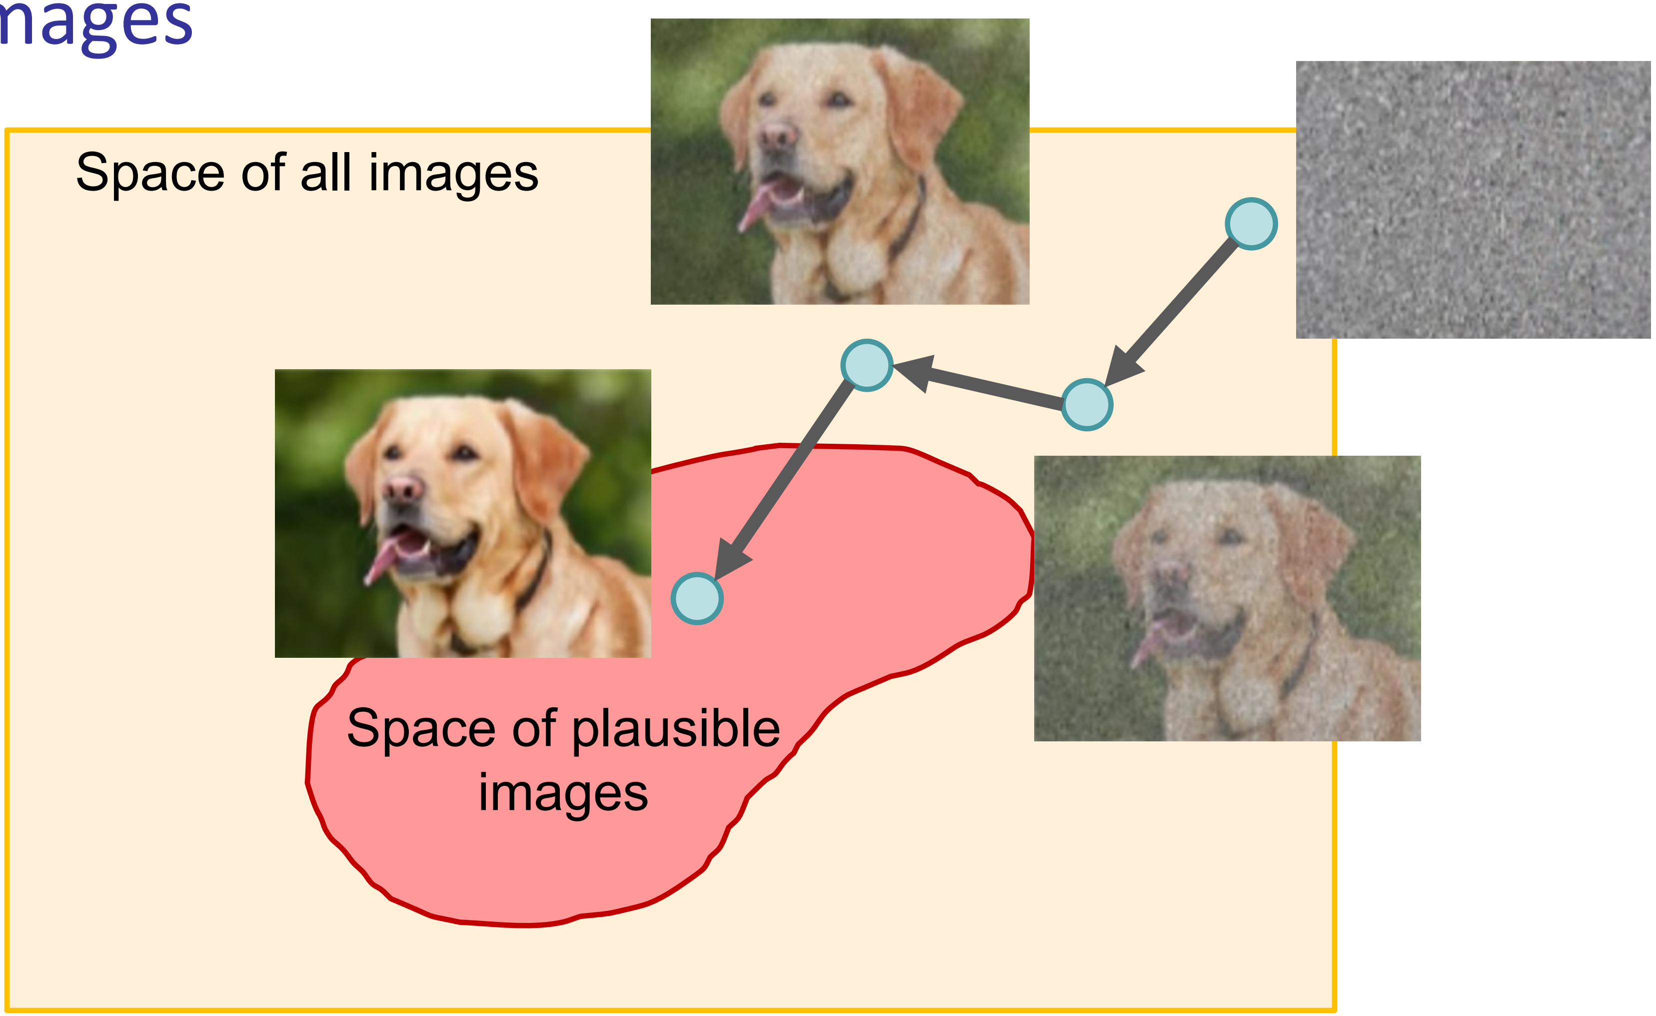

Application of Stationary Distribution: Web Link Analysis

- PageRank over a web graph
  - Each web page is a state
  - Initial distribution: uniform over pages
  - Transitions:
    - With prob. c, uniform jump to a random page (dotted lines, not all shown)
    - With prob. 1-c, follow a random outlink (solid lines)

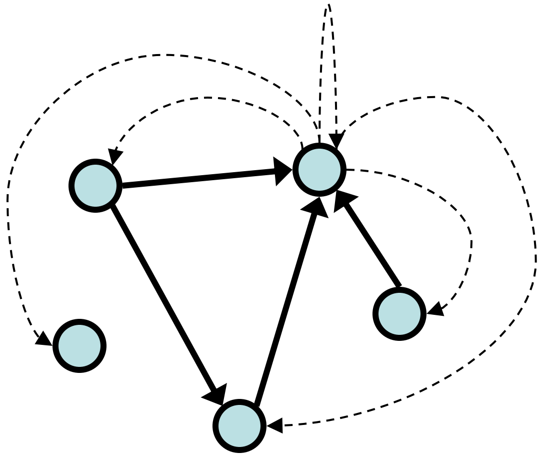

- Stationary distribution
  - Will spend more time on highly reachable pages
  - E.g. many ways to get to the Acrobat Reader download page
  - Somewhat robust to link spam
  - Google 1.0 returned the set of pages containing all your keywords in decreasing rank, now all search engines use link analysis along with many other factors (rank actually getting less important over time)

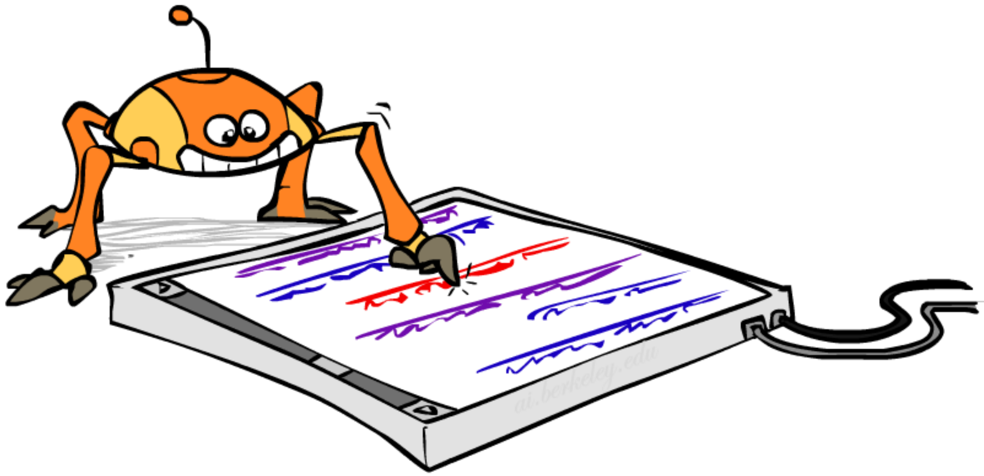

Today’s Topics

- Very quick probability recap
- **Markov Chains** & their **Stationary Distributions**
- **Hidden Markov Models** (HMMs) formulation
- Preview of **Filtering** with HMMs

Hidden Markov Models

Pacman – Sonar

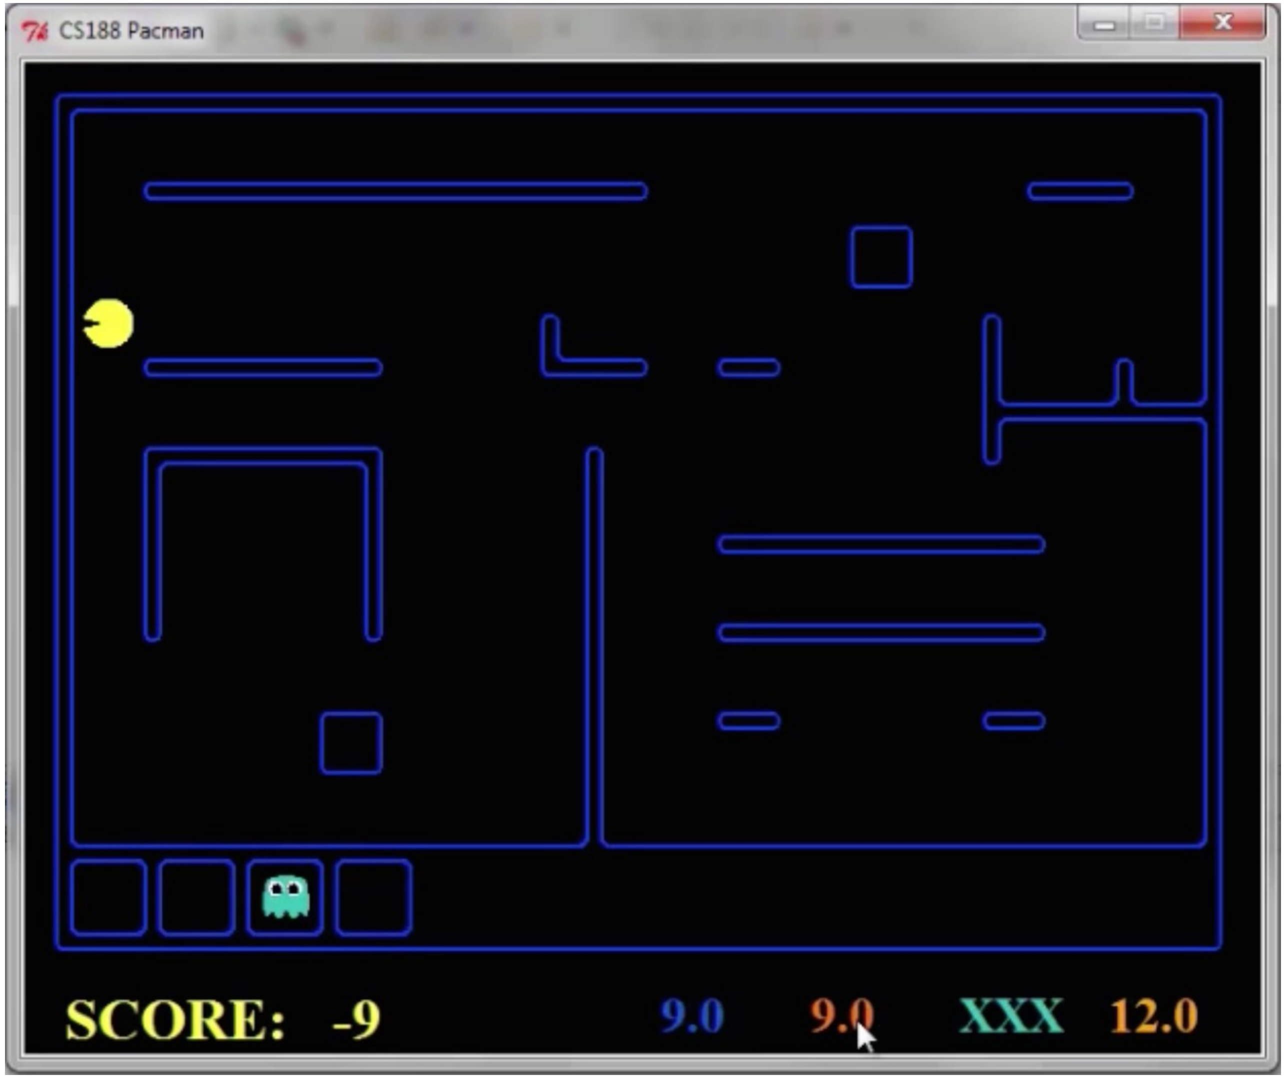

[Demo: Pacman – Sonar – No Beliefs(L14D1)]

Video of Demo Pacman – Sonar (no beliefs)

Video of Demo Pacman – Sonar (with beliefs)

Hidden Markov Models

- Markov chains not so useful for most agents
  - Need observations to update your beliefs
- Hidden Markov models (HMMs)
  - Underlying Markov chain over states X
  - You observe outputs (effects) at each time step

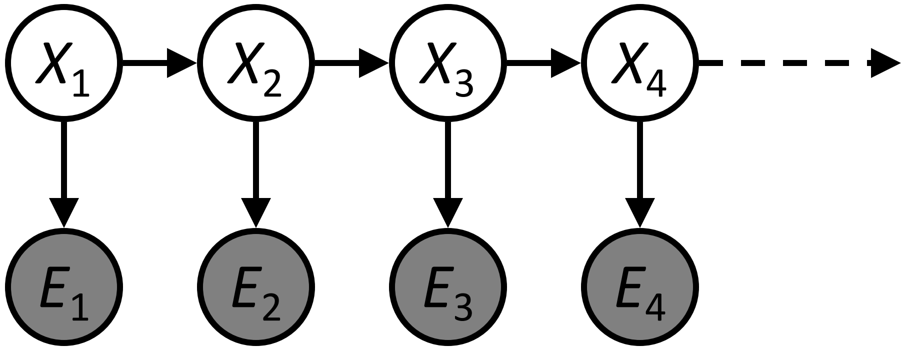

Example: Weather HMM

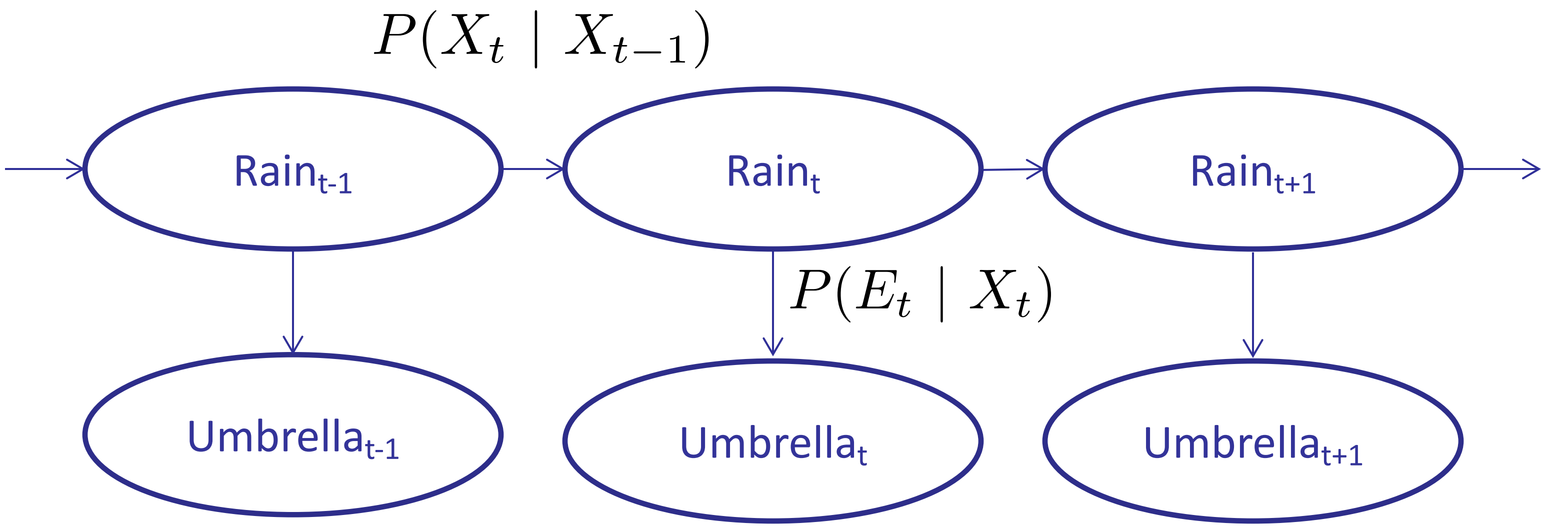

- An HMM is defined by:
  - Initial distribution: $P(X_1)$
  - Transitions: $P(X_t \mid X_{t-1})$
  - Emissions: $P(E_t \mid X_t)$

R_{t-1}	R_t	$P(R_t \mid R_{t-1})$
+r	+r	0.7
+r	-r	0.3
-r	+r	0.3
-r	-r	0.7

R_t	U_t	$P(U_t \mid R_t)$
+r	+u	0.9
+r	-u	0.1
-r	+u	0.2
-r	-u	0.8

Example: Ghostbusters HMM

- $P(X_1)$ = uniform
- $P(X|X')$ = usually move clockwise, but sometimes move in a random direction or stay in place
- $P(R_{ij}|X)$ = same sensor model as before: red means close, green means far away.

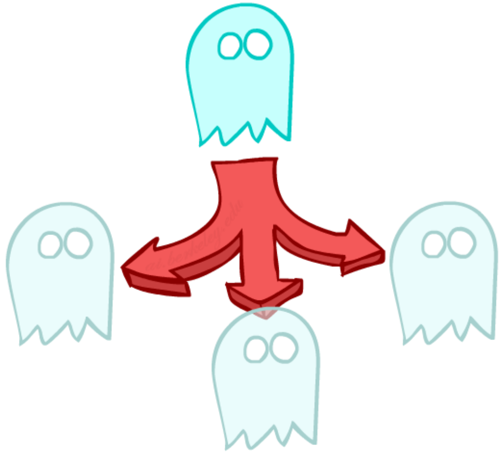

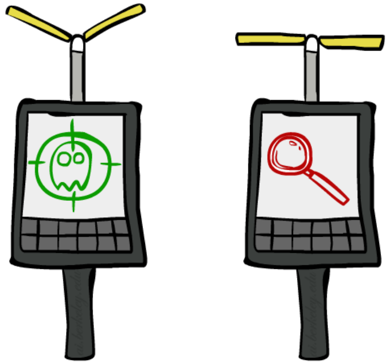

1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

$P(X_1)$

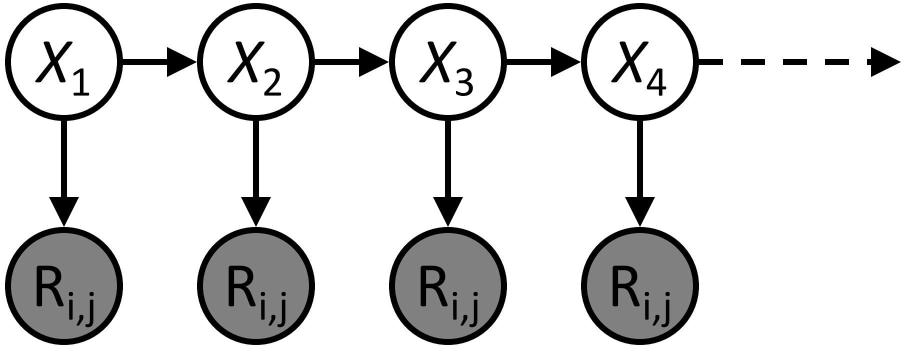

1/6	1/6	1/2
0	1/6	0
0	0	0

$P(X|X'=<1,2>)$

[Demo: Ghostbusters – Circular Dynamics – HMM (L14D2)]

Video of Demo Ghostbusters – Circular Dynamics -- HMM

Conditional Independence

- HMMs have two important independence properties:
  - Markov hidden process: future depends on past via the present
  - Current observation independent of all else given current state

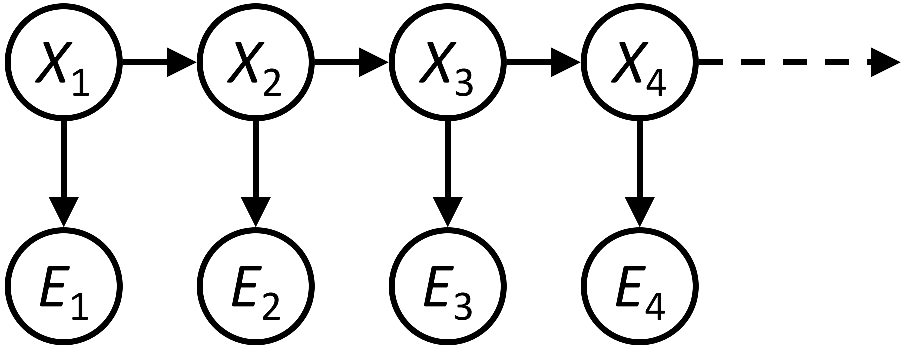

- Does this mean that evidence variables are guaranteed to be independent?
  - [No, they are correlated by the hidden state]

Real HMM Examples

- Speech recognition HMMs:
  - Observations are acoustic signals (continuous valued)
  - States are specific positions in specific words (so, tens of thousands)
- Machine translation HMMs:
  - Observations are words (tens of thousands)
  - States are translation options
- Robot tracking:
  - Observations are range readings (continuous)
  - States are positions on a map (continuous)

Example: Modeling Text with Sequence Models

Large language models: GPT3, PaLM, …

John St Clair Etouffee, world-renowned art expert, connoisseur, aesthete, art historian, and lecturer at the University of Oxford, is walking around the Louvre, savoring the exquisite artwork of all ages assembled before him. The Louvre is closed, allowing him free access to all the paintings and silence in which to meditate them. He stops in front of La Gioconda, Leonardo da Vinci's world-famous Mona Lisa. After a time, he begins to speak and describe the qualia, the frisson, the nigh-indescribable experience of the subtle swirls & *sfumato* of the painting:

"The sun is now barely a foot or two below the horizon and with every second, the golden tone of the painting deepens. Everything melts in it. I, who am having this experience, appear to you in a little corner of the painting and I look at you through the painting and our eyes meet. Our arms, our shoulders, our heads, our hands, our necks, our mouths, our lips, our eyes, our fingers and the tips of our toes—all this transforms itself into bright, beauteous colour. Everything becomes alive, all there is left in me is love and appreciation. It makes me want to shout at you "Don't leave!" but you don't hear me and I remain silent. La Gioconda looks at me and I at La Gioconda, the same as before, but now it's like being in two places at the same time and my feet have no feeling."

St. Clair Etouffee leaves the painting and enters the Hall of the Kings, and of Kings. He moves on to the Indian Wars, passes the Saloon d'Orléans, the King of Britain with his Tudor and Stuart wives, and is reminded of what St Clair had read a lifetime earlier at the age of twelve in a book about medieval paintings

Example: Modeling Text with Sequence Models

- Learn probability of a word given previous word?

$$P(X_t|X_{t-1})$$

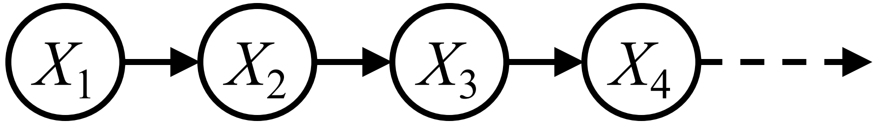

- Learn parameters of a Hidden Markov Model?

$$P(X_1) \quad P(X_t \mid X_{t-1}) \quad P(E_t \mid X_t)$$

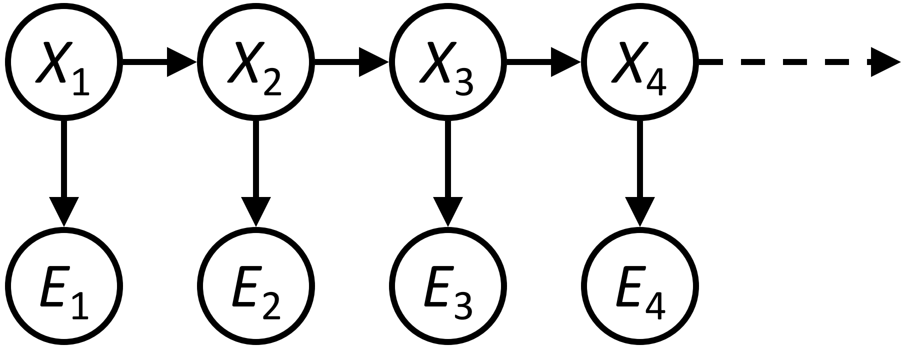

- Learn probability of a word given all previous words?

$$P(X_i|X_1, \ldots, X_{i-1})$$

Today's Topics

- Very quick probability recap
- **Markov Chains** & their **Stationary Distributions**
- **Hidden Markov Models** (HMMs) formulation
- Preview of **Filtering** with HMMs

Filtering / Monitoring

- Filtering, or monitoring, is the task of tracking the distribution $B_t(X) = P_t(X_t \mid e_1, \ldots, e_t)$ (the belief state) over time
- We start with $B_1(X)$ in an initial setting, usually uniform
- As time passes, or we get observations, we update $B(X)$
- The Kalman filter was invented in the 60's and first implemented as a method of trajectory estimation for the Apollo program

Example: Robot Localization

Example from Michael Pfeiffer

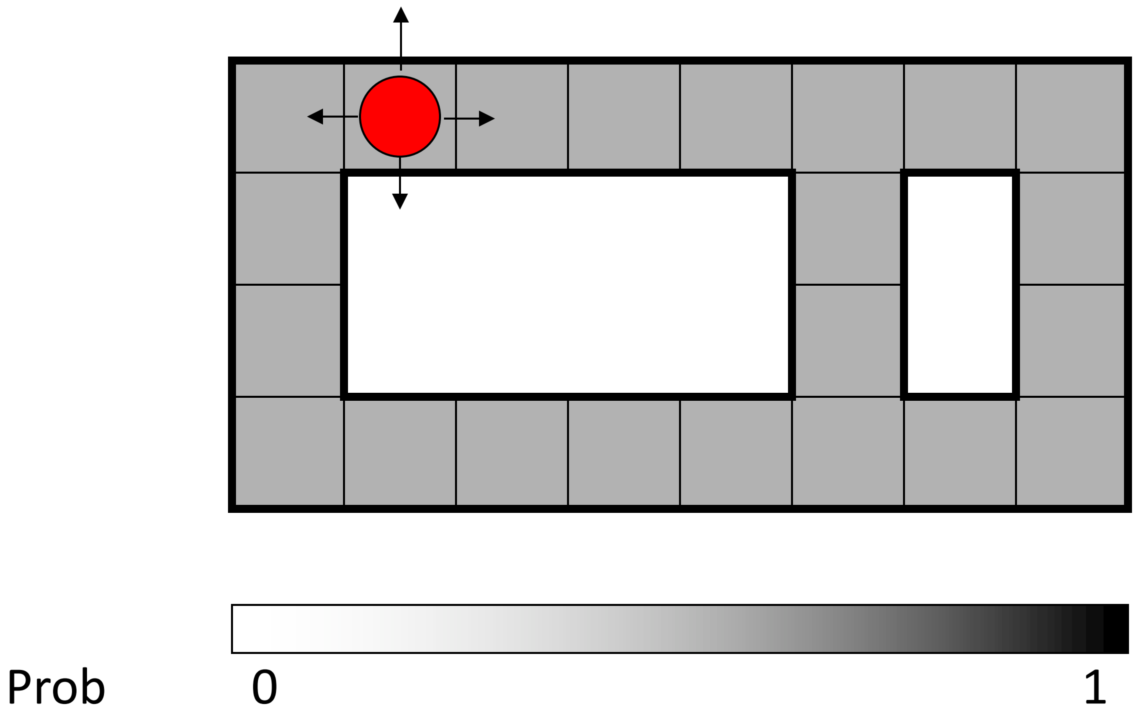

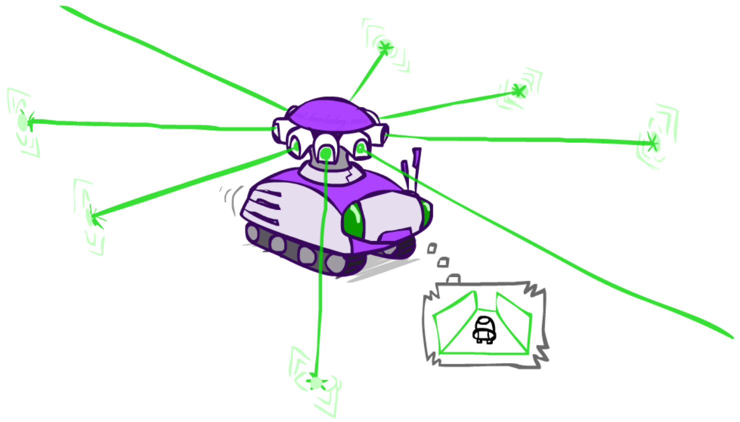

t=0

Sensor model: can read in which directions there is a wall, never more than 1 mistake

Motion model: may not execute action with small prob.

Example: Robot Localization

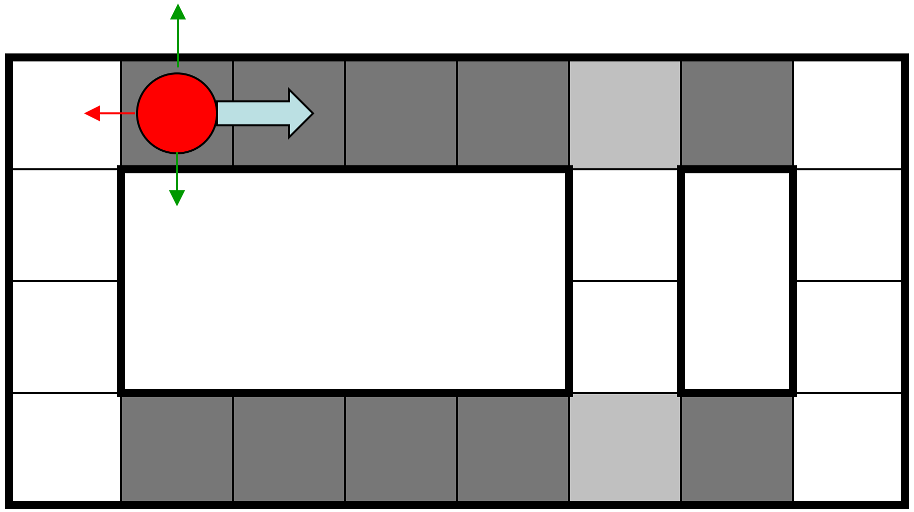

Prob 0 1

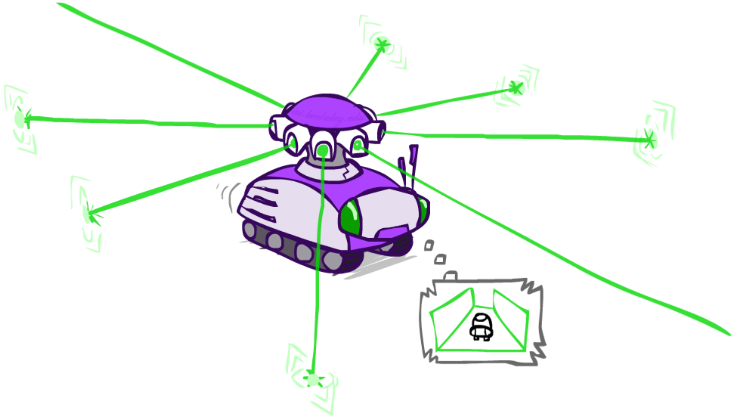

t=1

Lighter grey: was possible to get the reading, but less likely b/c required 1 mistake

Example: Robot Localization

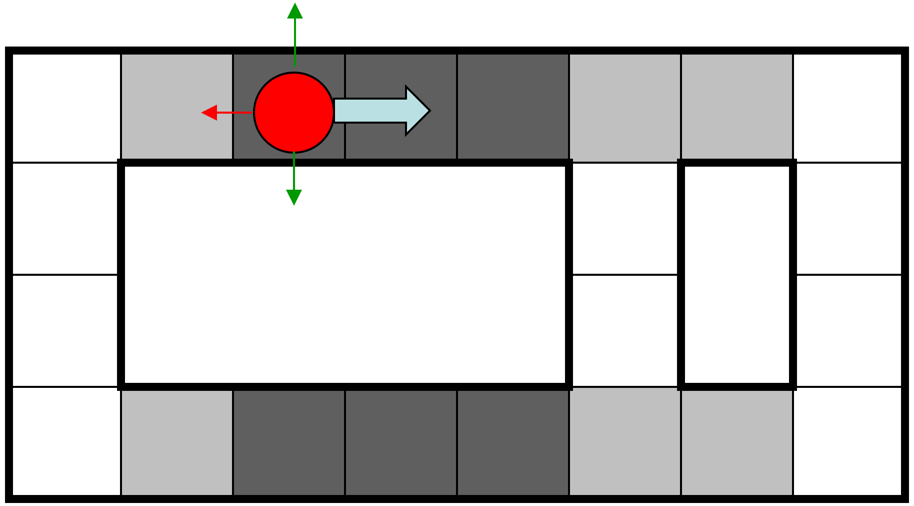

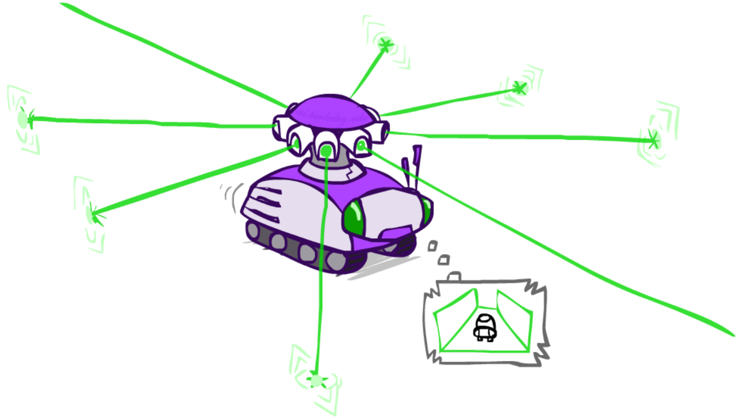

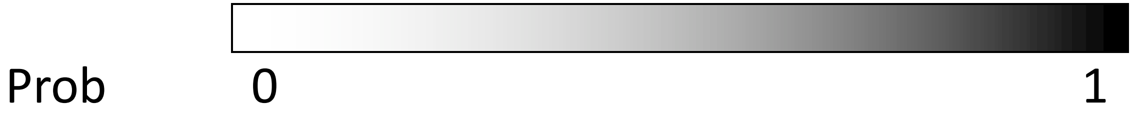

t=2

Example: Robot Localization

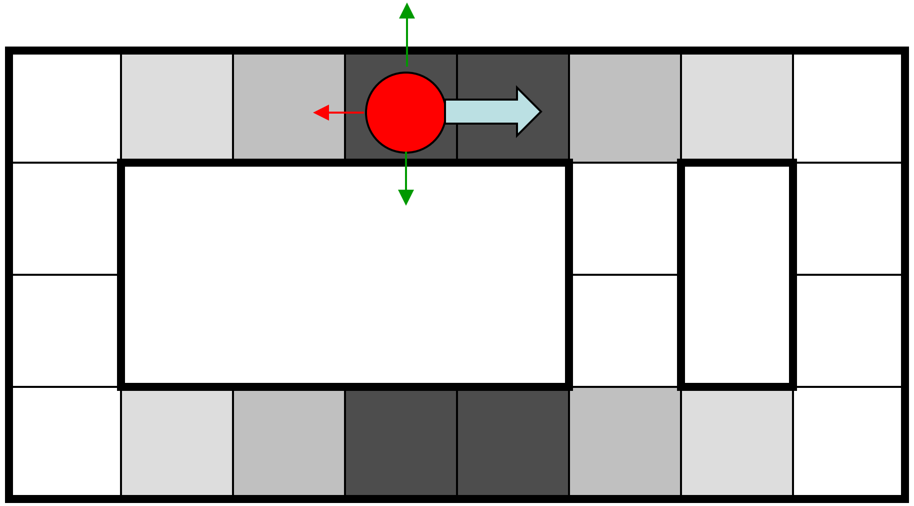

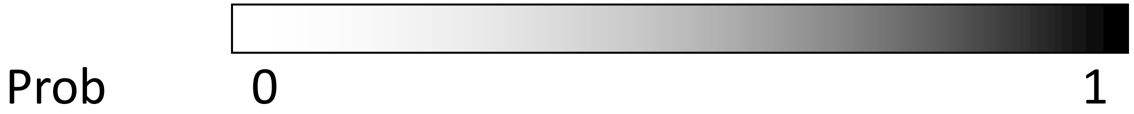

t=3

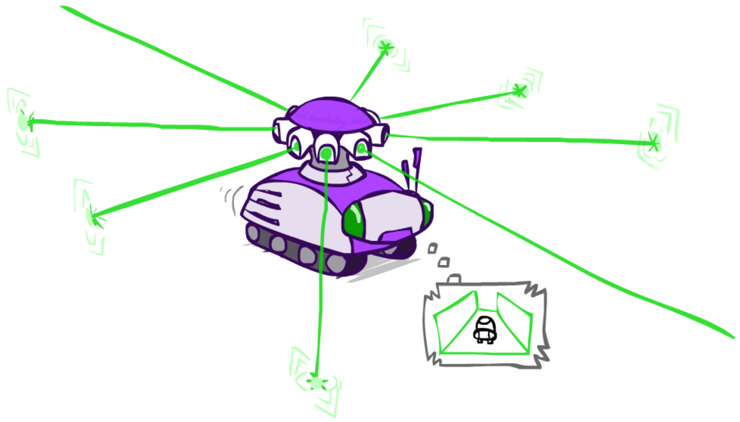

Example: Robot Localization

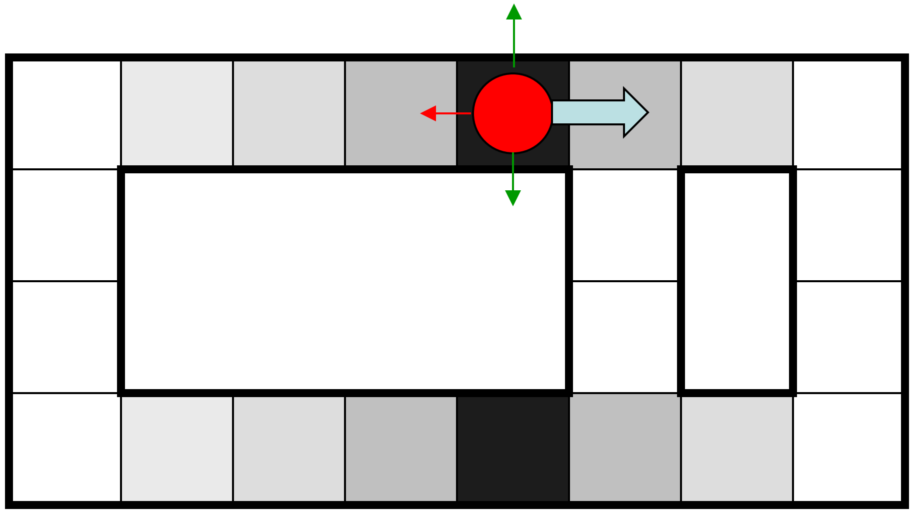

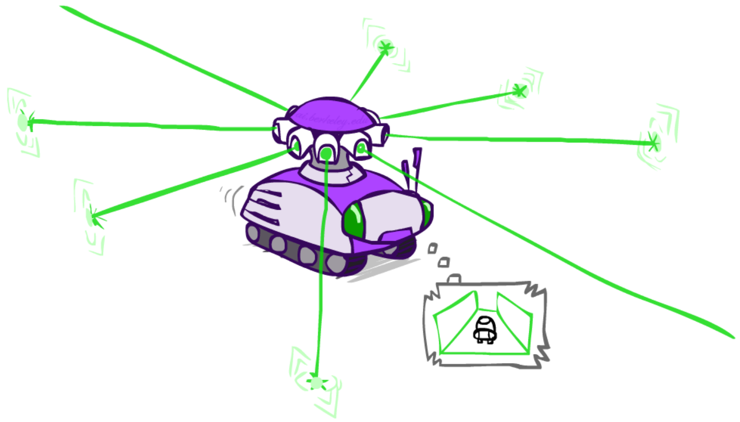

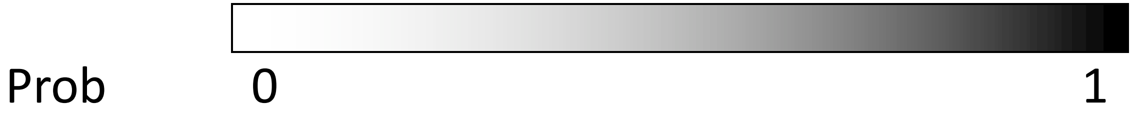

t=4

Example: Robot Localization

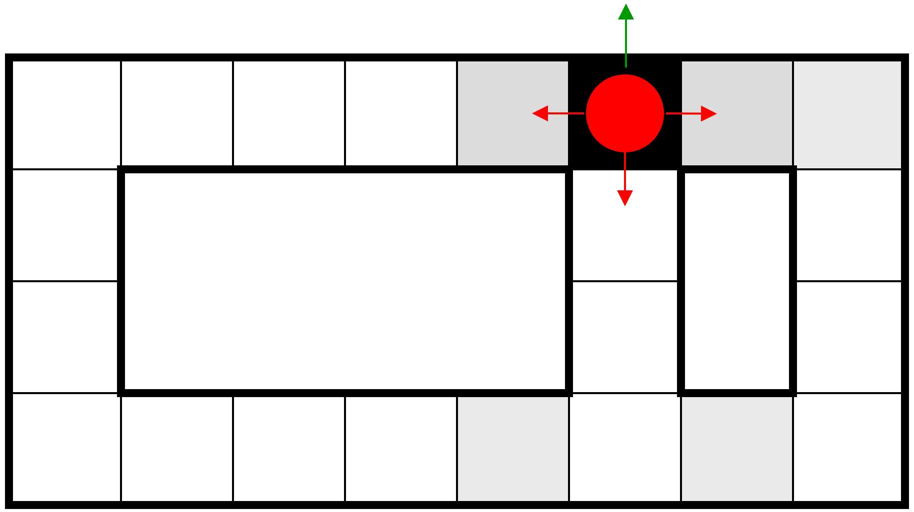

Prob 0 1

t=5

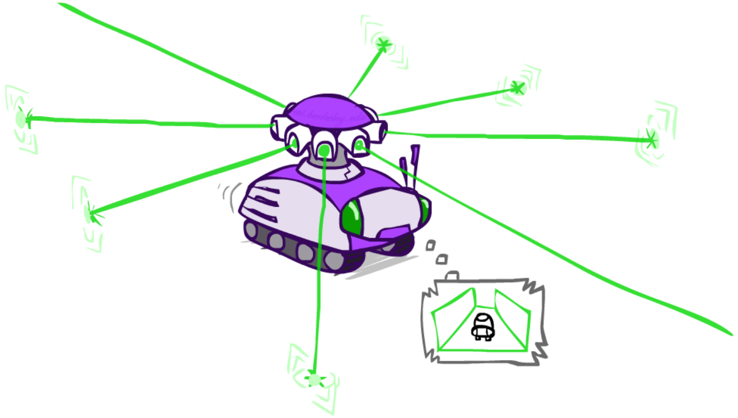

Inference: Find State Given Evidence

- We are given evidence at each time and want to know

$$B_t(X) = P(X_t | e_{1:t})$$

- Idea: start with $P(X_1)$ and derive B_t in terms of B_{t-1}
  - equivalently, derive B_{t+1} in terms of B_t

Two Steps: Passage of Time + Observation

$$B(X_t) = P(X_t|e_{1:t})$$

$$B'(X_{t+1}) = \sum_{x_t} P(X'|x_t)B(x_t)$$

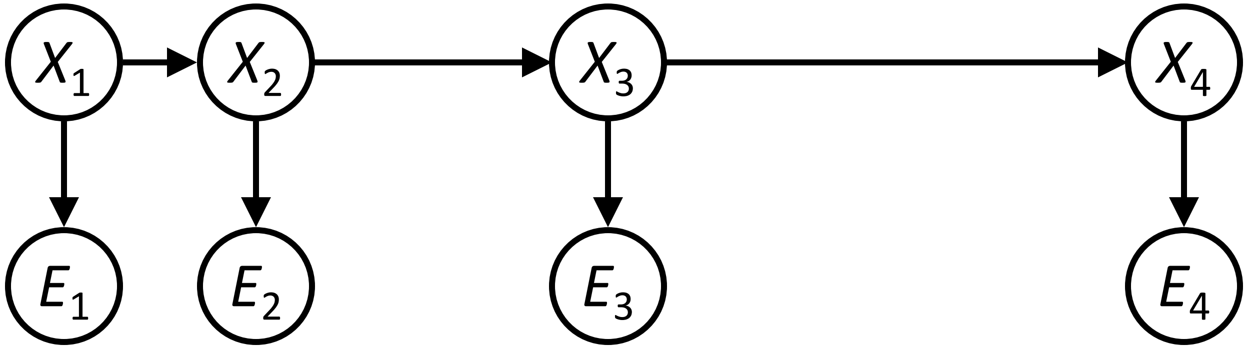

$$B(X_{t+1}) \propto_{X_{t+1}} P(e_{t+1}|X_{t+1})B'(X_{t+1})$$

Pacman – Sonar

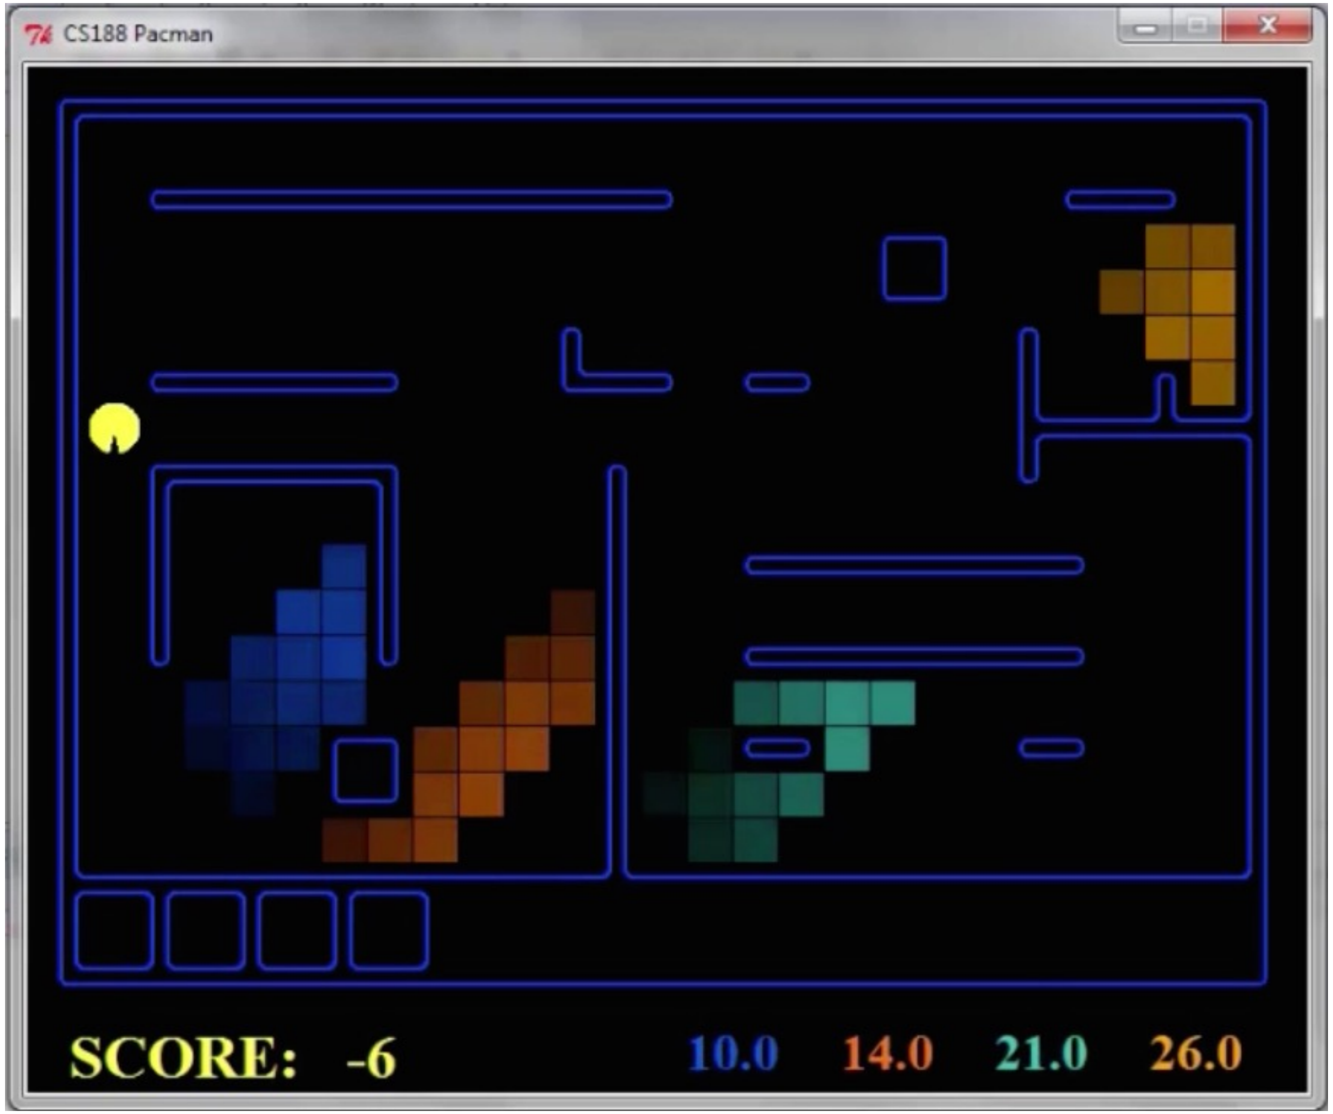

[Demo: Pacman – Sonar – No Beliefs(L14D1)]

Video of Demo Pacman – Sonar (with beliefs)

Next Time: Filtering