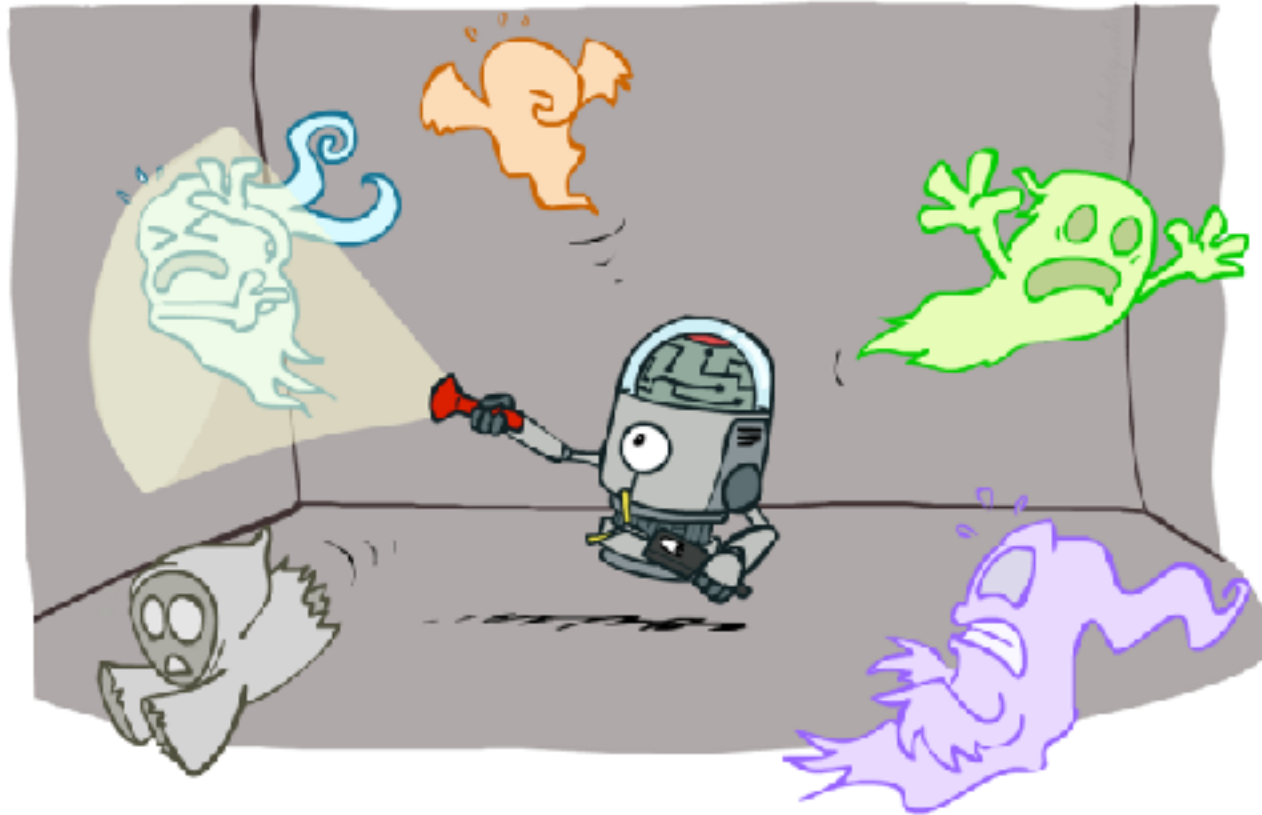


# CS 383: Artificial Intelligence

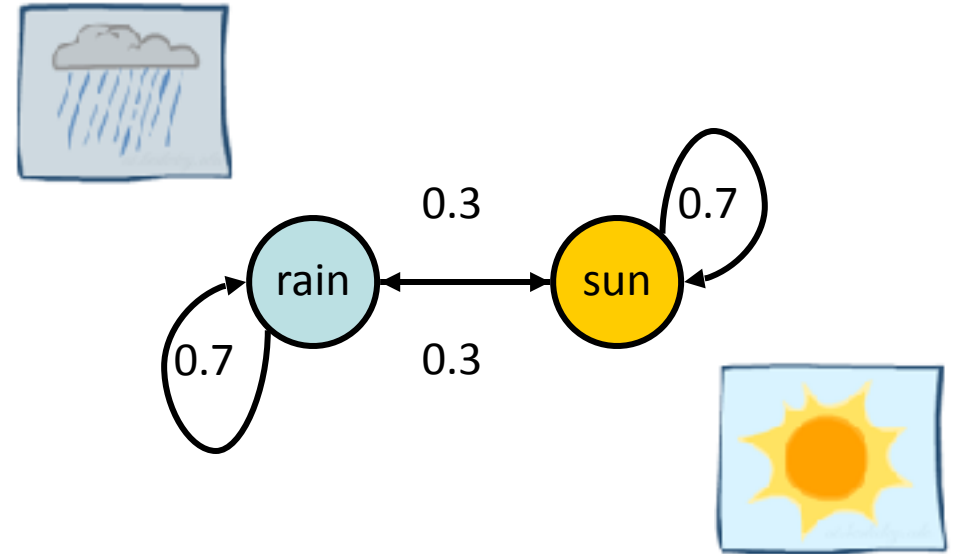
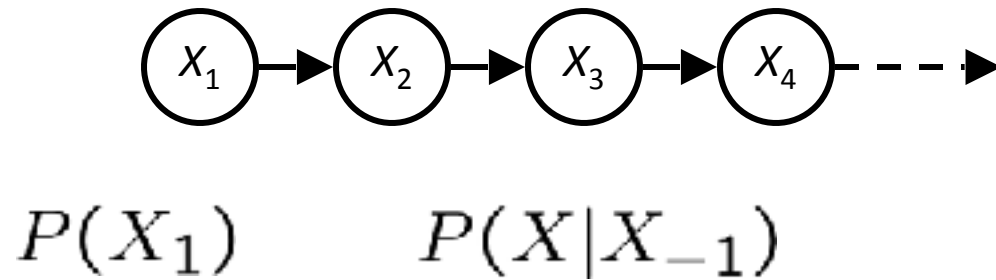
## Particle Filters and Applications of HMMs



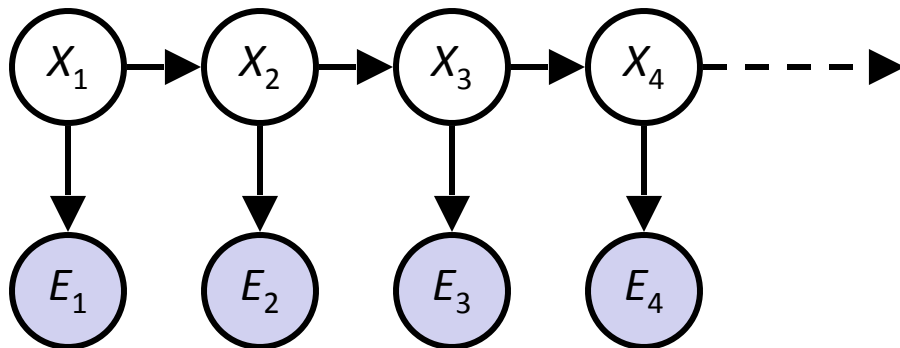
Prof. Scott Niekum — UMass Amherst

# Recap: Reasoning Over Time

## ■ Markov models



## ■ Hidden Markov models

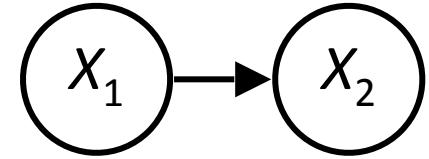


X	E	P
rain	umbrella	0.9
rain	no umbrella	0.1
sun	umbrella	0.2
sun	no umbrella	0.8

# Recap: Forward Algo - Passage of Time

- Assume we have current belief  $P(X | \text{evidence to date})$

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



- Then, after one time step passes:

$$\begin{aligned} P(X_{t+1} | e_{1:t}) &= \sum_{x_t} P(X_{t+1}, x_t | e_{1:t}) \\ &= \sum_{x_t} P(X_{t+1} | x_t, e_{1:t}) P(x_t | e_{1:t}) \\ &= \sum_{x_t} P(X_{t+1} | x_t) P(x_t | e_{1:t}) \end{aligned}$$

- Or compactly:

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

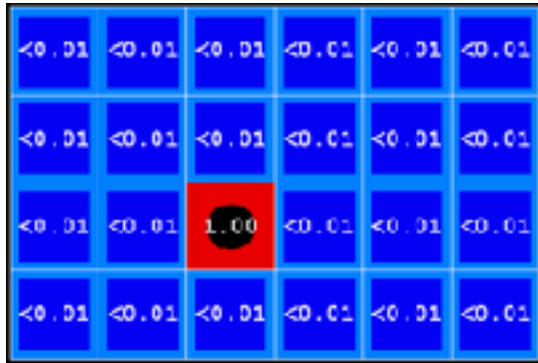
- Basic idea: beliefs get “pushed” through the transitions

- With the “B” notation, we have to be careful about what time step  $t$  the belief is about, and what evidence it includes

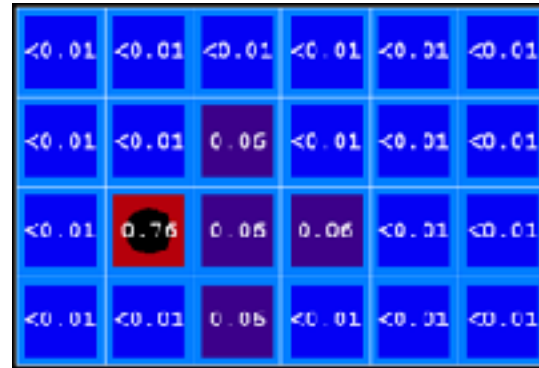
# Recap: Forward Algo - Passage of Time

- As time passes, uncertainty “accumulates”

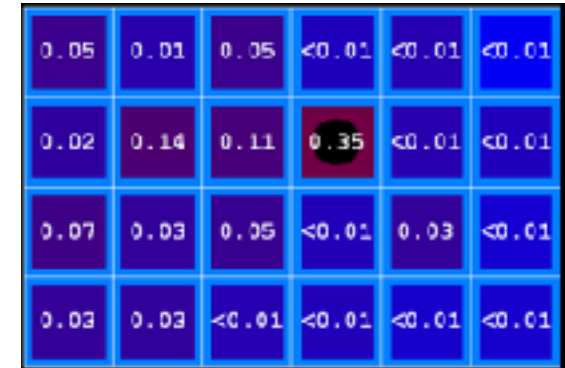
(Transition model: ghosts usually go clockwise)



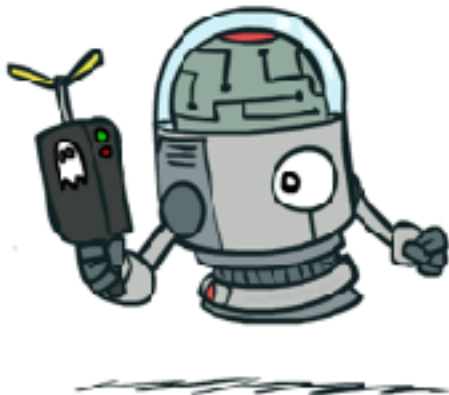
T = 1



T = 2



T = 5



# Recap: Forward Algo - Observation

- Assume we have current belief  $P(X \mid \text{previous evidence})$ :

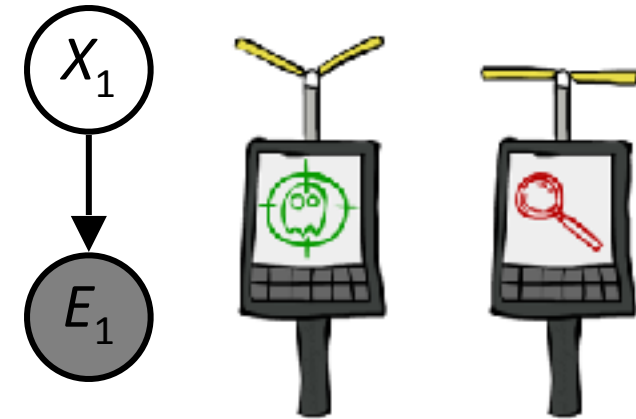
$$B'(X_{t+1}) = P(X_{t+1} \mid e_{1:t})$$

- Then, after evidence comes in:

$$\begin{aligned} P(X_{t+1} \mid e_{1:t+1}) &= P(X_{t+1}, e_{t+1} \mid e_{1:t}) / P(e_{t+1} \mid e_{1:t}) \\ &\propto_{X_{t+1}} P(X_{t+1}, e_{t+1} \mid e_{1:t}) \\ &= P(e_{t+1} \mid e_{1:t}, X_{t+1}) P(X_{t+1} \mid e_{1:t}) \\ &= P(e_{t+1} \mid X_{t+1}) P(X_{t+1} \mid e_{1:t}) \end{aligned}$$

- Or, compactly:

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



- Basic idea: beliefs “reweighted” by likelihood of evidence
- Unlike passage of time, we have to renormalize

# Recap: Forward Algo - Observation

- As we get observations, beliefs get reweighted, uncertainty “decreases”

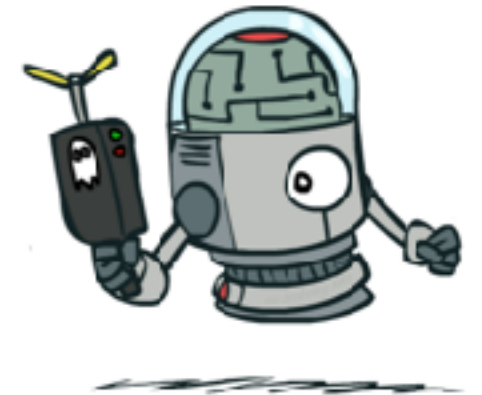
0.05	0.01	0.05	<0.01	<0.01	<0.01
0.02	0.14	0.11	0.35	<0.01	<0.01
0.07	0.02	0.05	<0.01	0.02	<0.01
0.03	0.03	<0.01	<0.01	<0.01	<0.01

Before observation

<0.01	<0.01	<0.01	<0.01	0.02	<0.01
<0.01	<0.01	<0.01	0.83	0.02	<0.01
<0.01	<0.01	0.11	<0.01	<0.01	<0.01
<0.01	<0.01	<0.01	<0.01	<0.01	<0.01

After observation

$$B(X) \propto P(e|X)B'(X)$$



# Recap: The Forward Algorithm

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

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

- We can derive the following updates

$$\begin{aligned} P(x_t | e_{1:t}) &\propto_X P(x_t, e_{1:t}) \\ &= \sum_{x_{t-1}} P(x_{t-1}, x_t, e_{1:t}) \\ &= \sum_{x_{t-1}} P(x_{t-1}, e_{1:t-1}) P(x_t | x_{t-1}) P(e_t | x_t) \\ &= P(e_t | x_t) \sum_{x_{t-1}} P(x_t | x_{t-1}) P(x_{t-1}, e_{1:t-1}) \end{aligned}$$

We can normalize as we go if we want to have  $P(x|e)$  at each time step, or just once at the end...

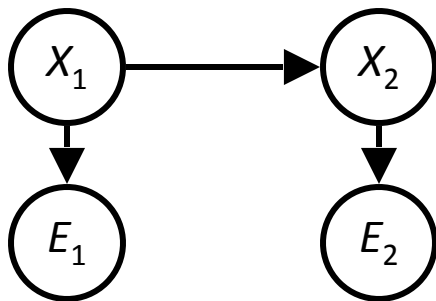
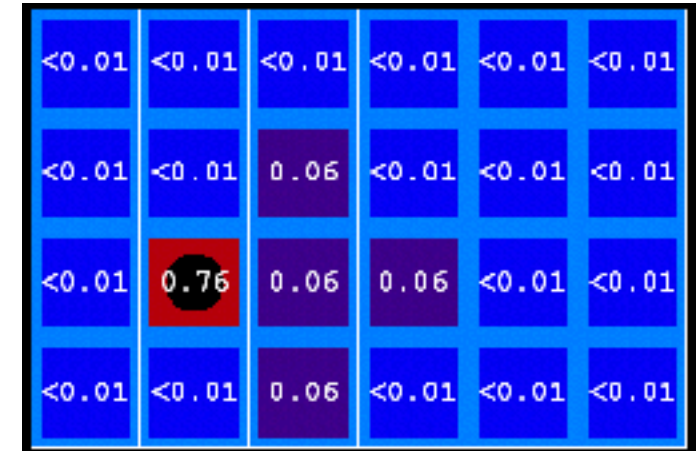
# Recap: Online Filtering w/ Forward Algo

**Elapse time:** compute  $P(X_t | e_{1:t-1})$

$$P(x_t | e_{1:t-1}) = \sum_{x_{t-1}} P(x_{t-1} | e_{1:t-1}) \cdot P(x_t | x_{t-1})$$

**Observe:** compute  $P(X_t | e_{1:t})$

$$P(x_t | e_{1:t}) \propto P(x_t | e_{1:t-1}) \cdot P(e_t | x_t)$$



**Belief:**  $\langle P(\text{rain}), P(\text{sun}) \rangle$

$P(X_1)$        $\langle 0.5, 0.5 \rangle$       *Prior on  $X_1$*

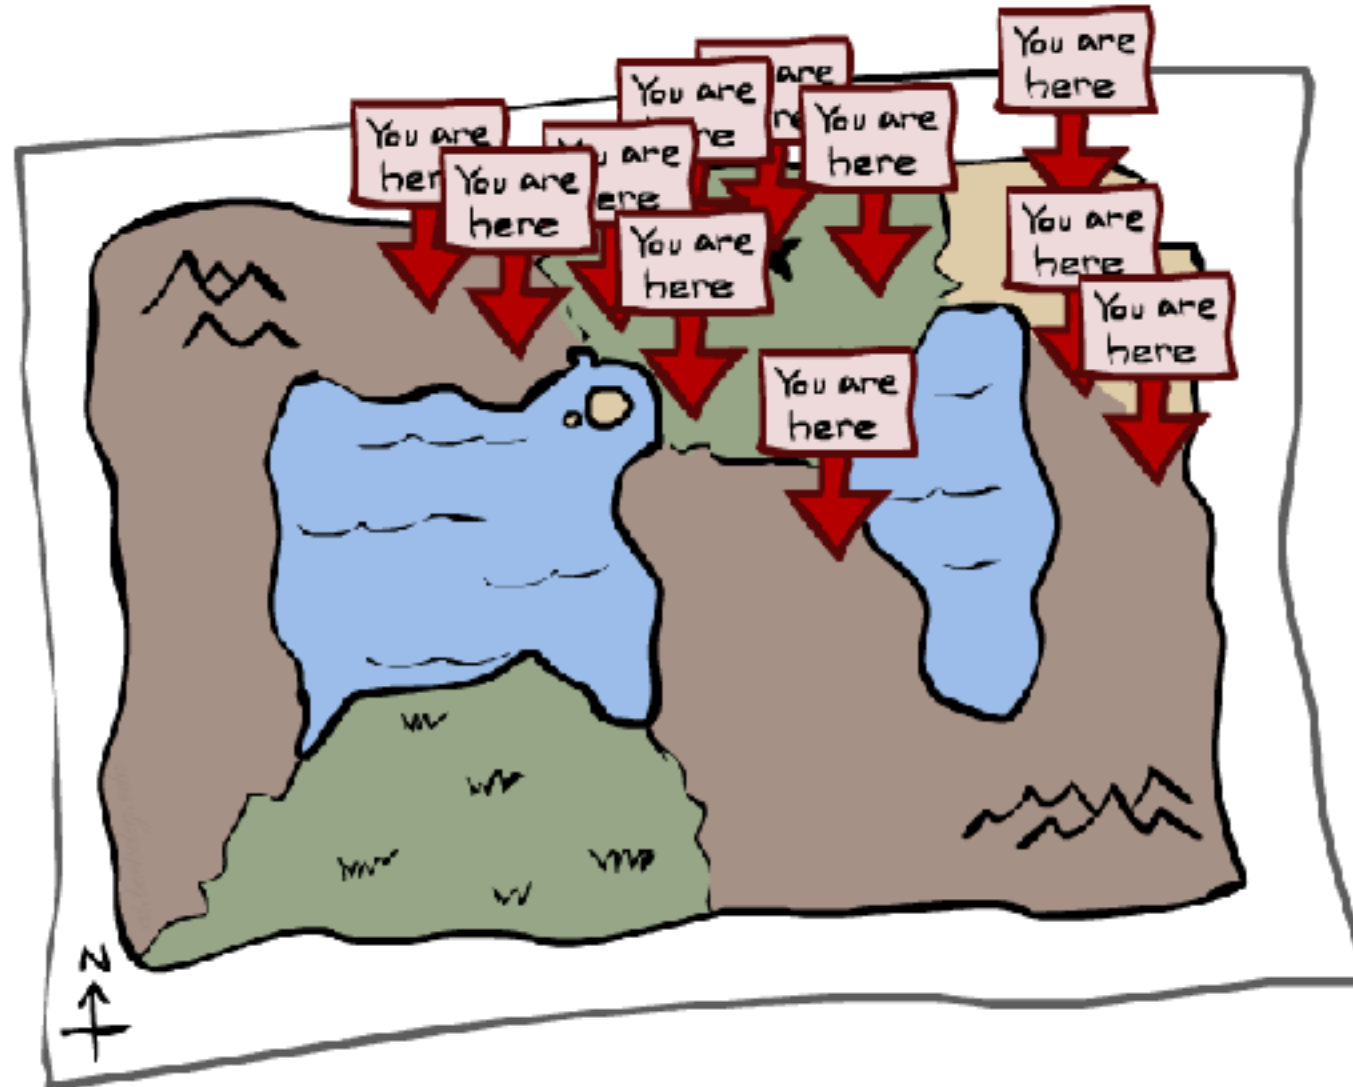
$P(X_1 | E_1 = \text{umbrella})$        $\langle 0.82, 0.18 \rangle$       *Observe*

$P(X_2 | E_1 = \text{umbrella})$        $\langle 0.63, 0.37 \rangle$       *Elapse time*

$P(X_2 | E_1 = \text{umb}, E_2 = \text{umb})$        $\langle 0.88, 0.12 \rangle$       *Observe*



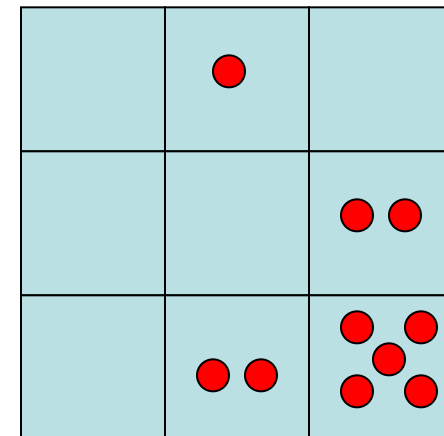
# Particle Filtering



# Particle Filtering

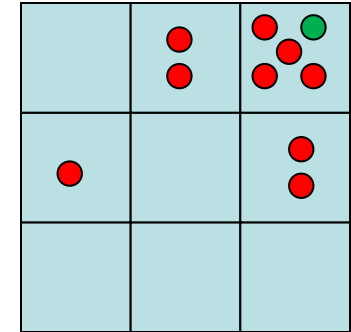
- Filtering: approximate solution
- Sometimes  $|X|$  is too big to use exact inference
  - $|X|$  may be too big to even store  $B(X)$
  - E.g.  $X$  is continuous
- Solution: approximate inference
  - Track samples of  $X$ , not all values
  - Samples are called particles
  - Time per step is linear in the number of samples
  - But: number needed may be large
  - In memory: list of particles, not states
- This is how robot localization works in practice
- Particle is just new name for sample

0.0	0.1	0.0
0.0	0.0	0.2
0.0	0.2	0.5



# Representation: Particles

- Our representation of  $P(X)$  is now a list of  $N$  particles (samples)
  - Generally,  $N \ll |X|$  (...but not in project 4)
  - Storing map from  $X$  to counts would defeat the point
- $P(x)$  approximated by number of particles with value  $x$ 
  - So, many  $x$  may have  $P(x) = 0$ !
  - More particles, more accuracy
- For now, all particles have a weight of 1
- Particle filtering uses three repeated steps:
  - Elapse time and observe (similar to exact filtering) and resample



Particles:

(3,3)

(2,3)

(3,3)

(3,2)

(3,3)

(3,2)

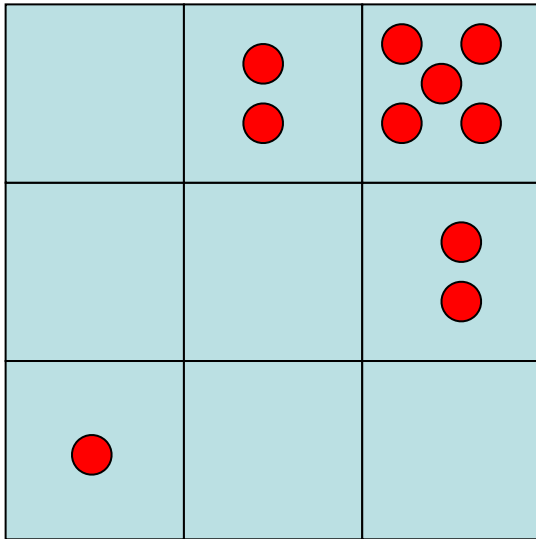
(1,2)

(3,3)

(3,3)

(2,3)

# Example: Elapse Time



Elapse Time



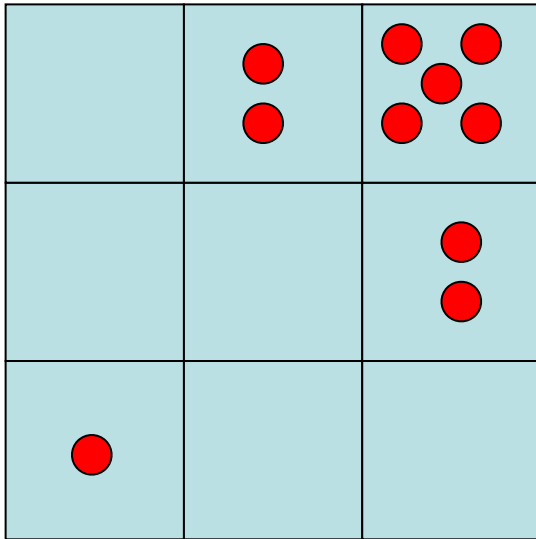
Policy: ghosts always move up  
(or stay in place if already at top)



Belief over possible  
ghost positions at time  $t$

New belief at  
time  $t+1$

# Example: Elapse Time

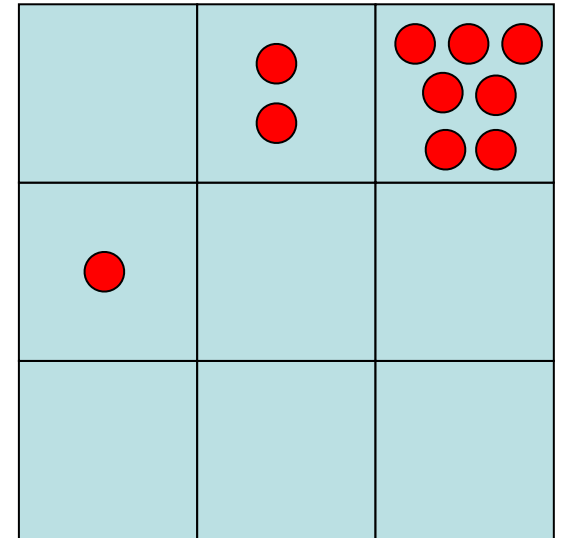


Belief over possible ghost positions at time  $t$

Elapse Time



Policy: ghosts always move up (or stay in place if already at top)



New belief at time  $t+1$

# Particle Filtering: Elapse Time

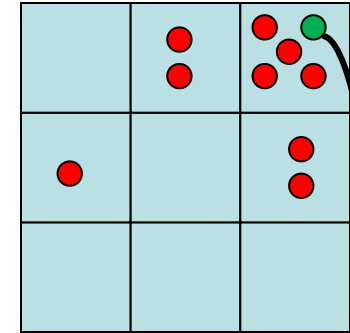
- Each particle is moved by sampling its next position from the transition model

$$x' = \text{sample}(P(X'|x))$$

- Sample frequencies reflect the transition probabilities
  - Here, most samples move clockwise, but some move in another direction or stay in place
- This captures the passage of time
    - If enough samples, close to exact values before and after (consistent)

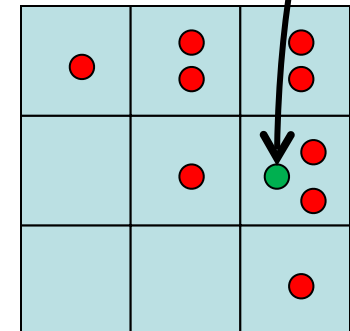
Particles:

(3,3)  
(2,3)  
(3,3)  
(3,2)  
(3,3)  
(3,2)  
(1,2)  
(3,3)  
(3,3)  
(2,3)

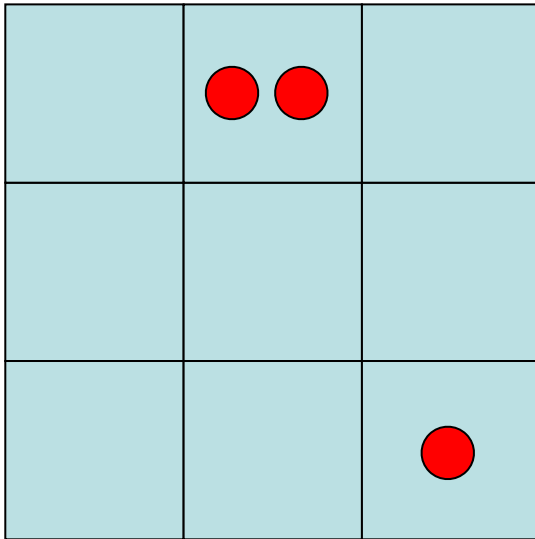


Particles:

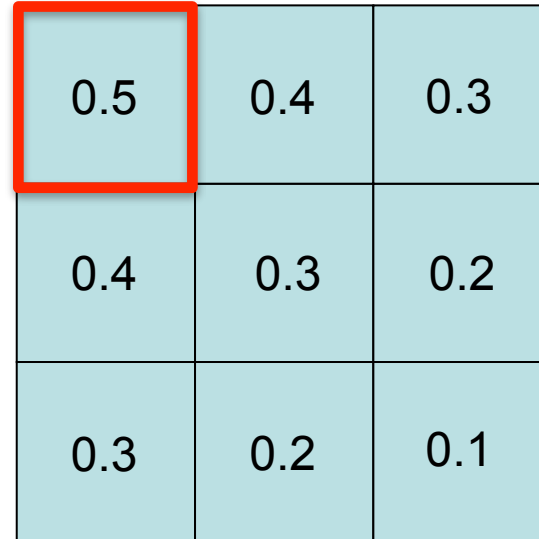
(3,2)  
(2,3)  
(3,2)  
(3,1)  
(3,3)  
(3,2)  
(1,3)  
(2,3)  
(3,2)  
(2,2)



# Example: Observe



+

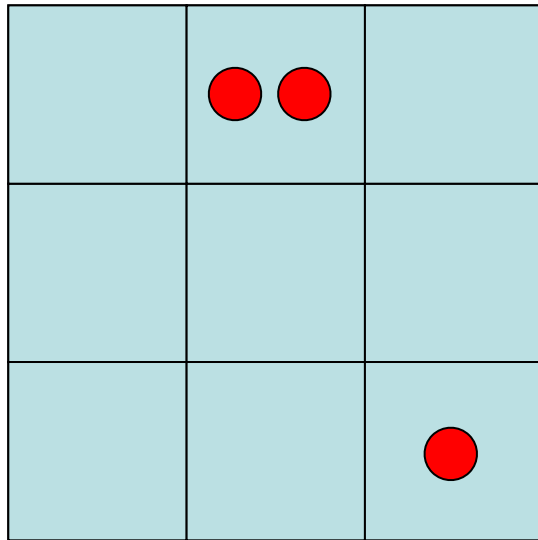


Belief over possible ghost positions before observation

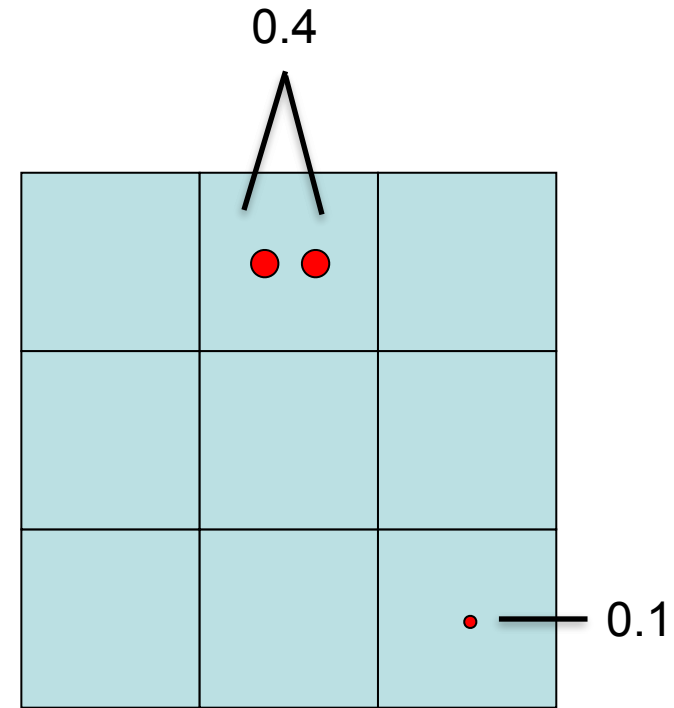
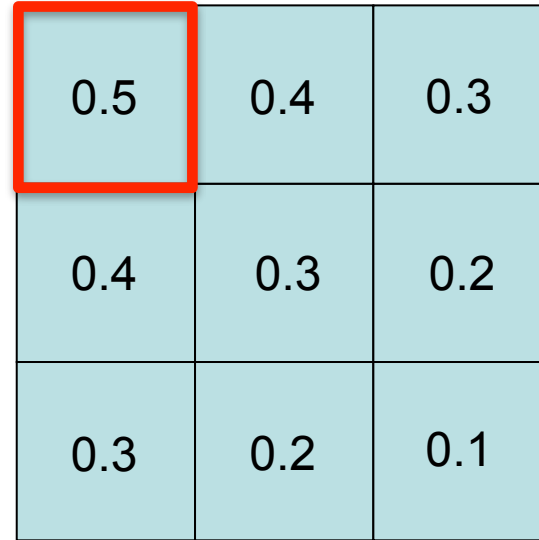
Observation and evidence likelihoods  $p(e | X)$

New belief after observation

# Example: Observe



+



Belief over possible ghost positions before observation

Observation and evidence likelihoods  $p(e | X)$

New belief after observation



# Particle Filtering: Observe

- Slightly trickier:

- Don't sample observation, fix it
- Similar to likelihood weighting, downweight samples based on the evidence

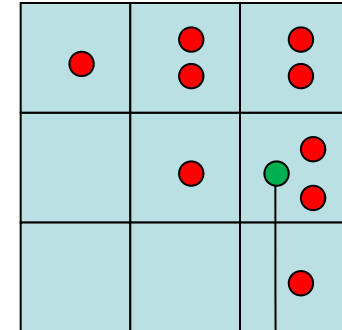
$$w(x) = P(e|x)$$

$$B(X) \propto P(e|X)B'(X)$$

- As before, the probabilities don't sum to one, since all have been downweighted

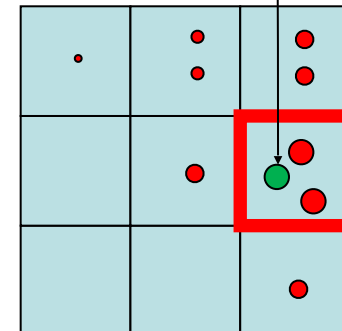
Particles:

(3,2)  
(2,3)  
(3,2)  
(3,1)  
(3,3)  
(3,2)  
(1,3)  
(2,3)  
(3,2)  
(2,2)



Particles:

(3,2) w=.9  
(2,3) w=.2  
(3,2) w=.9  
(3,1) w=.4  
(3,3) w=.4  
(3,2) w=.9  
(1,3) w=.1  
(2,3) w=.2  
(3,2) w=.9  
(2,2) w=.4

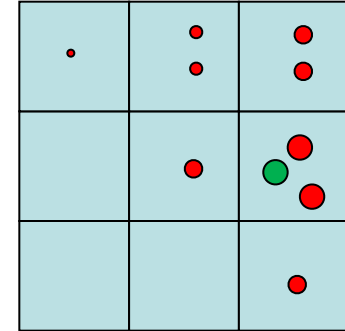


# Particle Filtering: Resample

- Rather than tracking weighted samples, we resample
- N times, we choose from our weighted sample distribution (i.e. draw with replacement)
- This essentially renormalizes the distribution
- Now the update is complete for this time step, continue with the next one

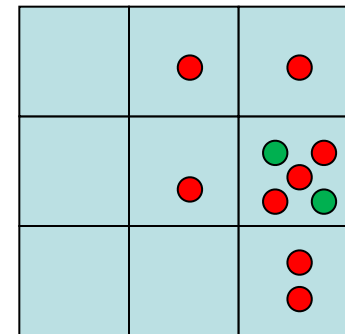
Particles:

(3,2)  $w=.9$   
(2,3)  $w=.2$   
(3,2)  $w=.9$   
(3,1)  $w=.4$   
(3,3)  $w=.4$   
(3,2)  $w=.9$   
(1,3)  $w=.1$   
(2,3)  $w=.2$   
(3,2)  $w=.9$   
(2,2)  $w=.4$



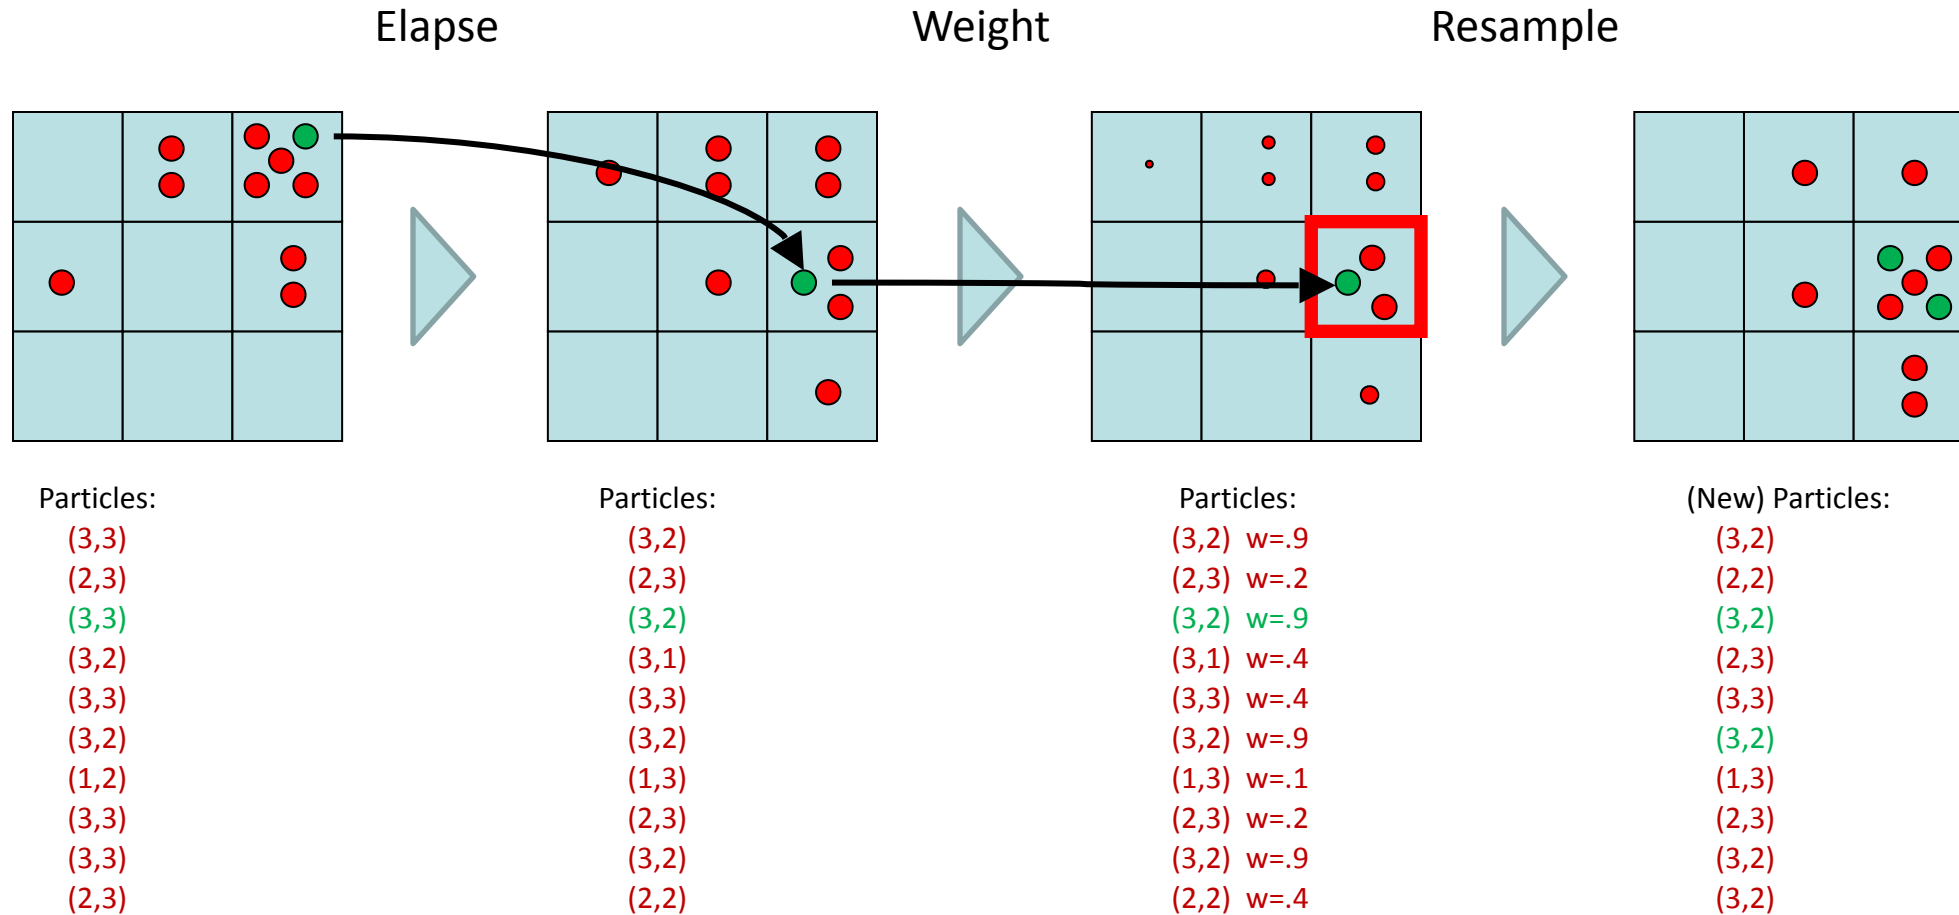
(New) Particles:

(3,2)  
(2,2)  
(3,2)  
(2,3)  
(3,3)  
(3,2)  
(1,3)  
(2,3)  
(3,2)  
(3,2)



# Recap: Particle Filtering

- Particles: track samples of states rather than an explicit distribution



# Moderate Number of Particles

The screenshot shows a PyDev IDE window with a game titled "ghostbusters". The game interface consists of a 6x10 grid of numbers on a black background. The numbers are as follows:

<0.01	0.04	0.01	0.01	<0.01	0.04	0.03	0.02	0.04	0.02
<0.01	0.02	<0.01	0.01	0.02	<0.01	0.02	0.01	0.01	0.02
0.03	0.06	0.01	0.02	<0.01	0.04	0.01	<0.01	0.01	0.01
<0.01	0.01	0.01	0.01	0.04	0.03	0.03	<0.01	0.01	0.04
<0.01	<0.01	0.03	0.03	<0.01	0.02	0.02	0.01	0.05	0.01
0.03	0.02	0.02	0.01	0.01	0.01	<0.01	0.02	0.02	<0.01

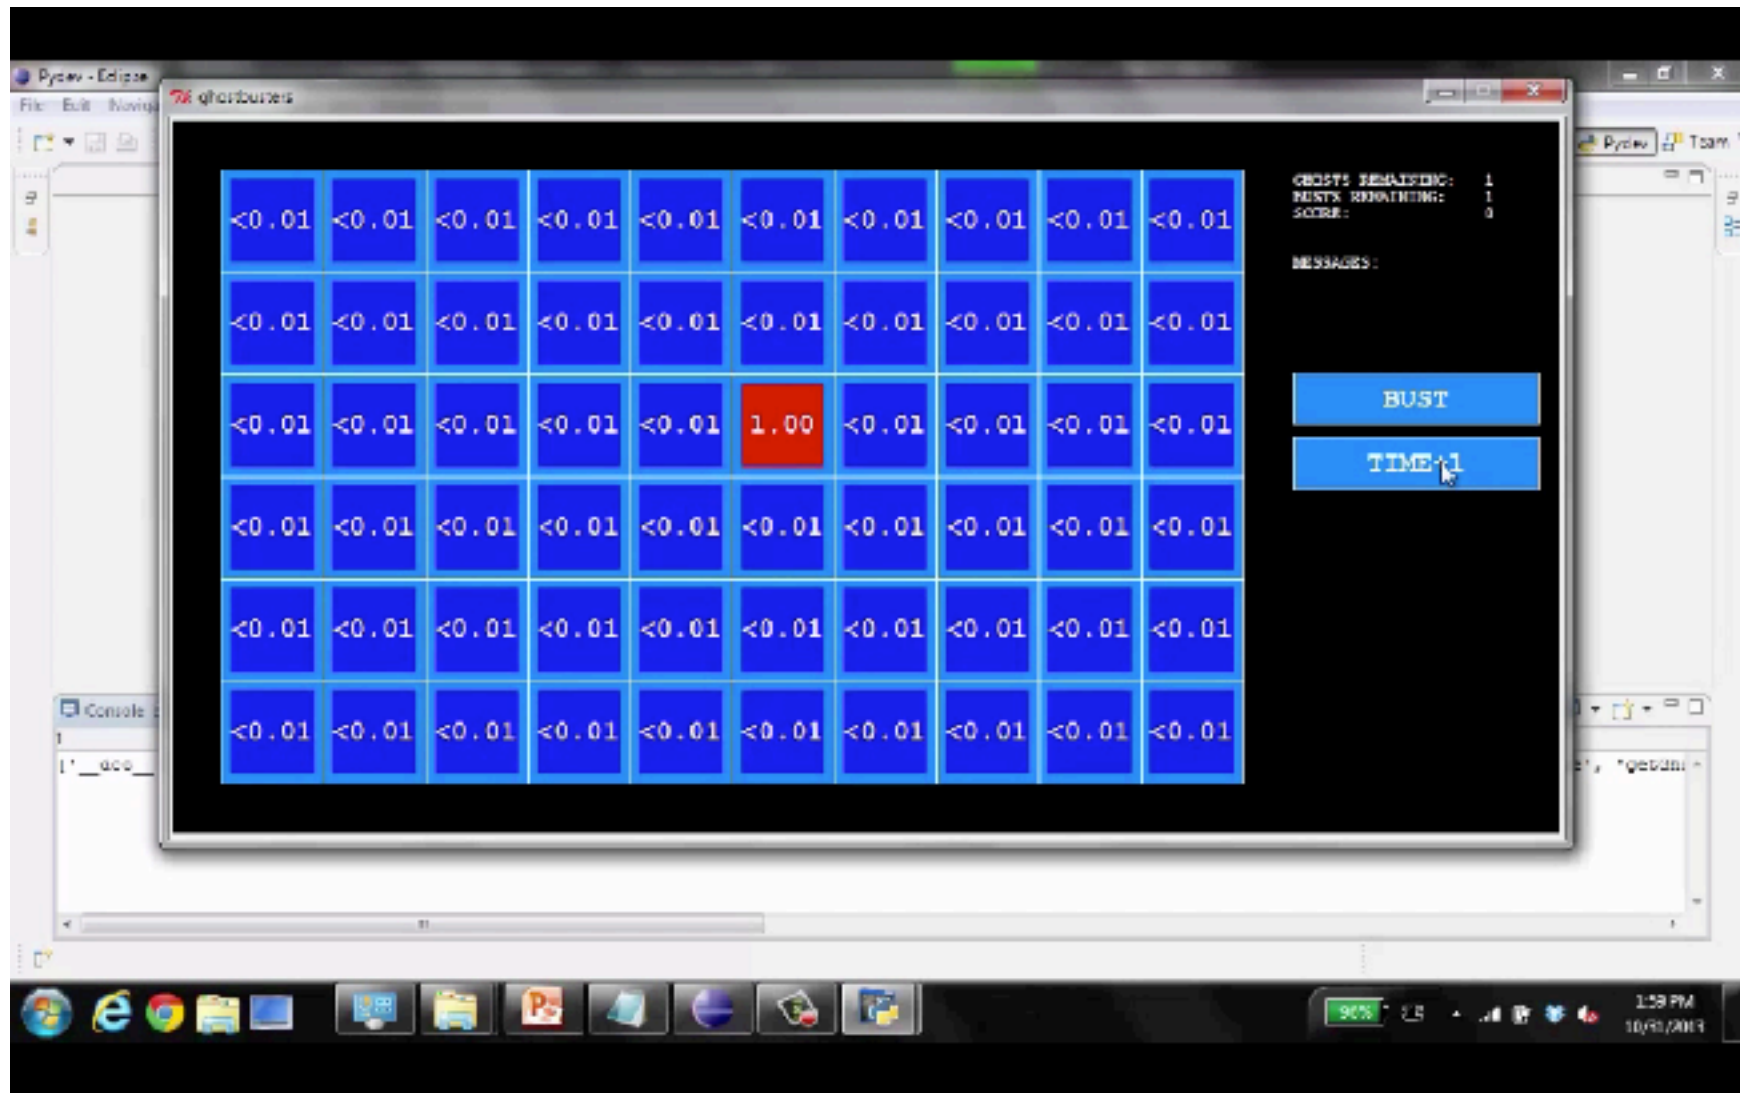
To the right of the grid, the following statistics are displayed:

```
GHOSTS REMAINING: 1  
BUSTS REMAINING: 1  
STRIKE: 0
```

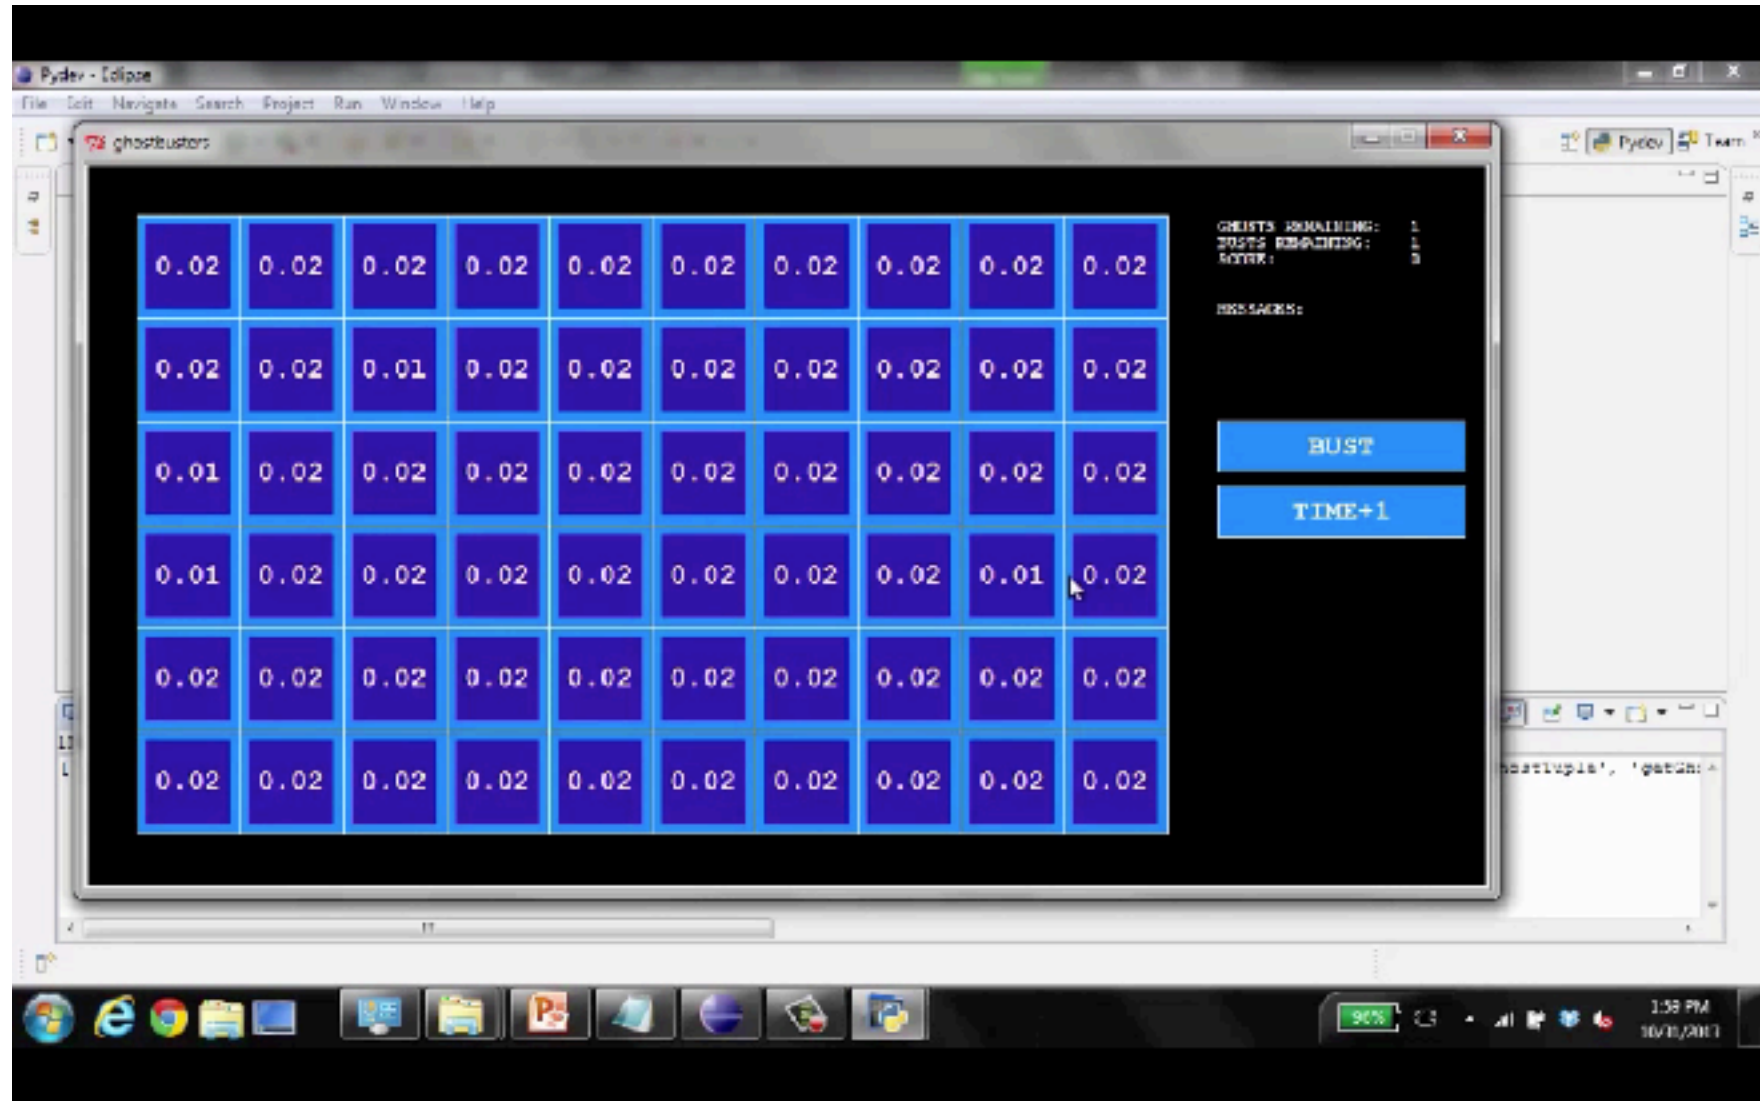
Below the statistics are two buttons: "BUST" and "TIME+1".

The IDE window also shows a console on the right with some output, and a taskbar at the bottom with various application icons and system tray information (90% battery, 1:58 PM, 10/31/2013).

# One Particle

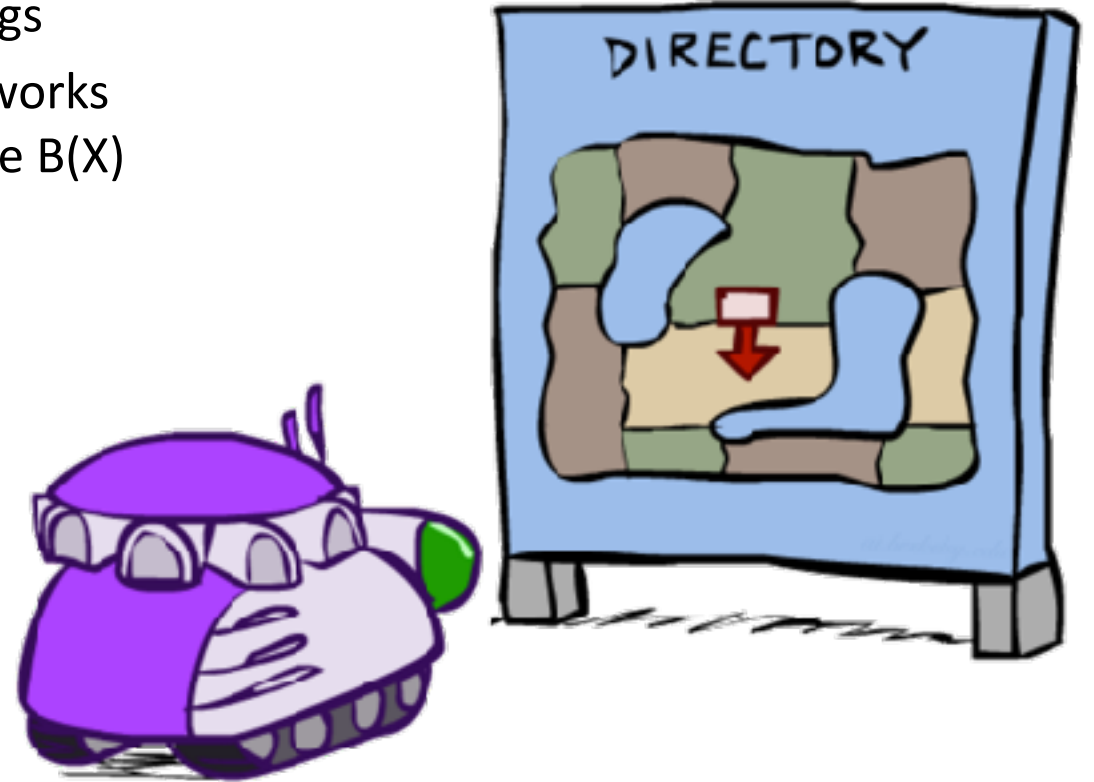
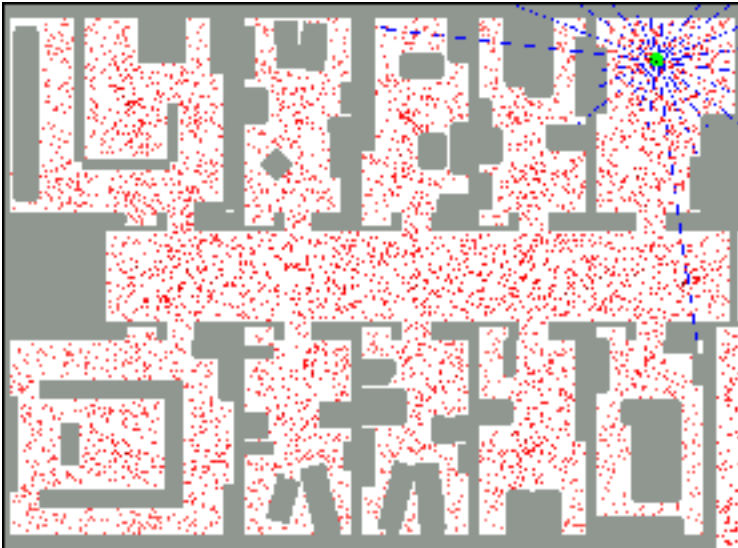


# Huge Number of Particles



# Robot Localization

- In robot localization:
  - We know the map, but not the robot's position
  - Observations may be vectors of range finder readings
  - State space and readings are typically continuous (works basically like a very fine grid) and so we cannot store  $B(X)$
  - Particle filtering is a main technique



# Particle Filter Localization (Sonar)



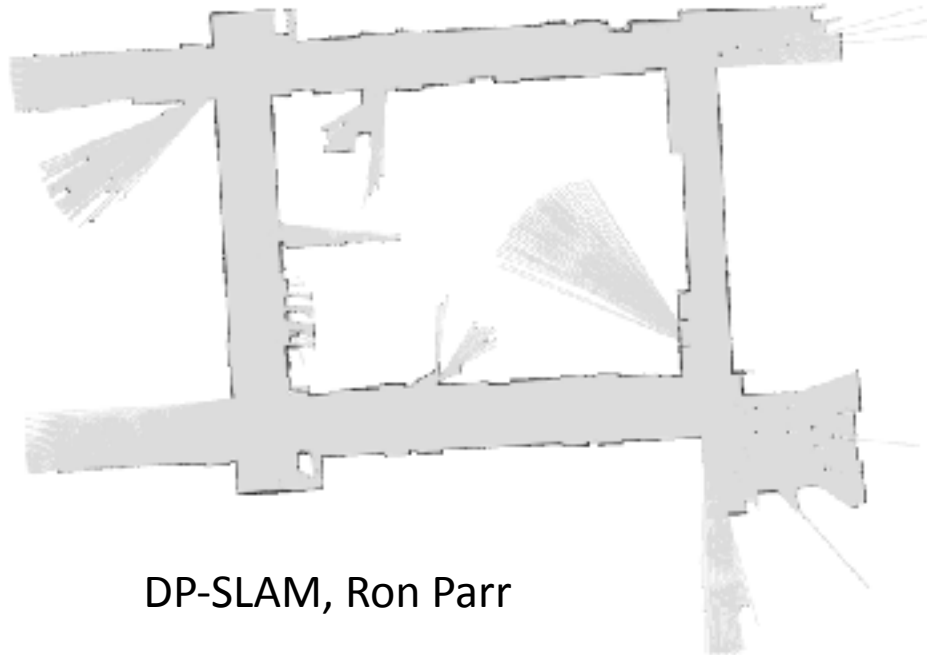
**Global localization with  
sonar sensors**

40000

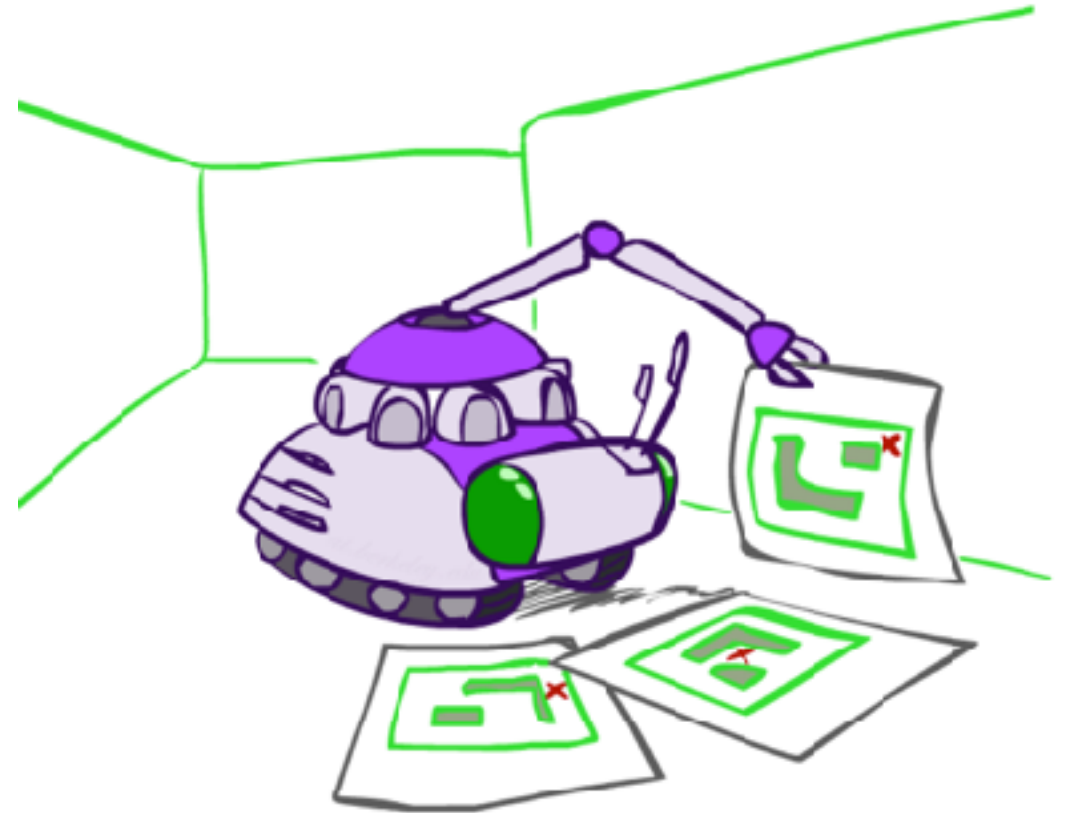


# Robot Mapping

- SLAM: Simultaneous Localization And Mapping
  - We do not know the map or our location
  - State consists of position AND map!
  - Main techniques: Kalman filtering (Gaussian HMMs) and particle methods

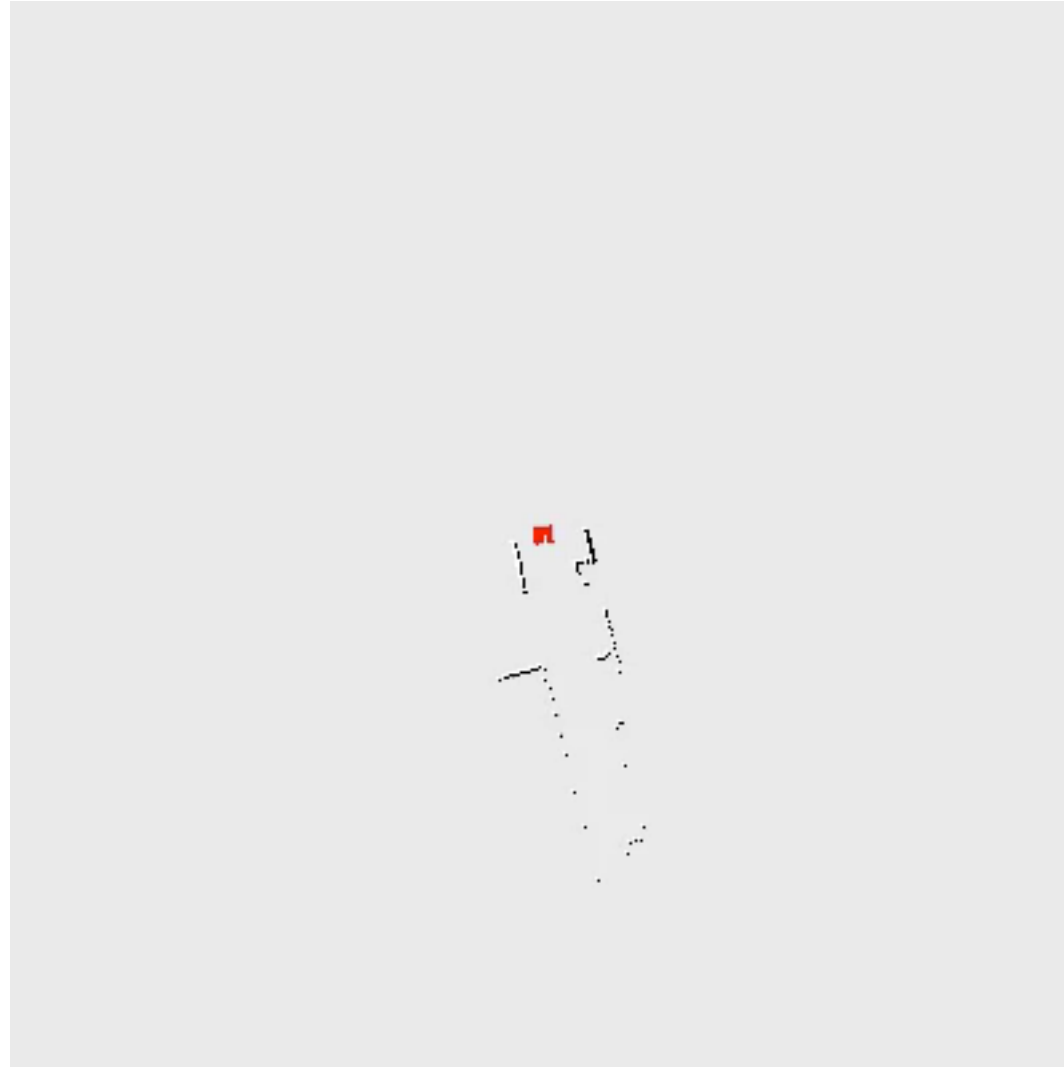


DP-SLAM, Ron Parr

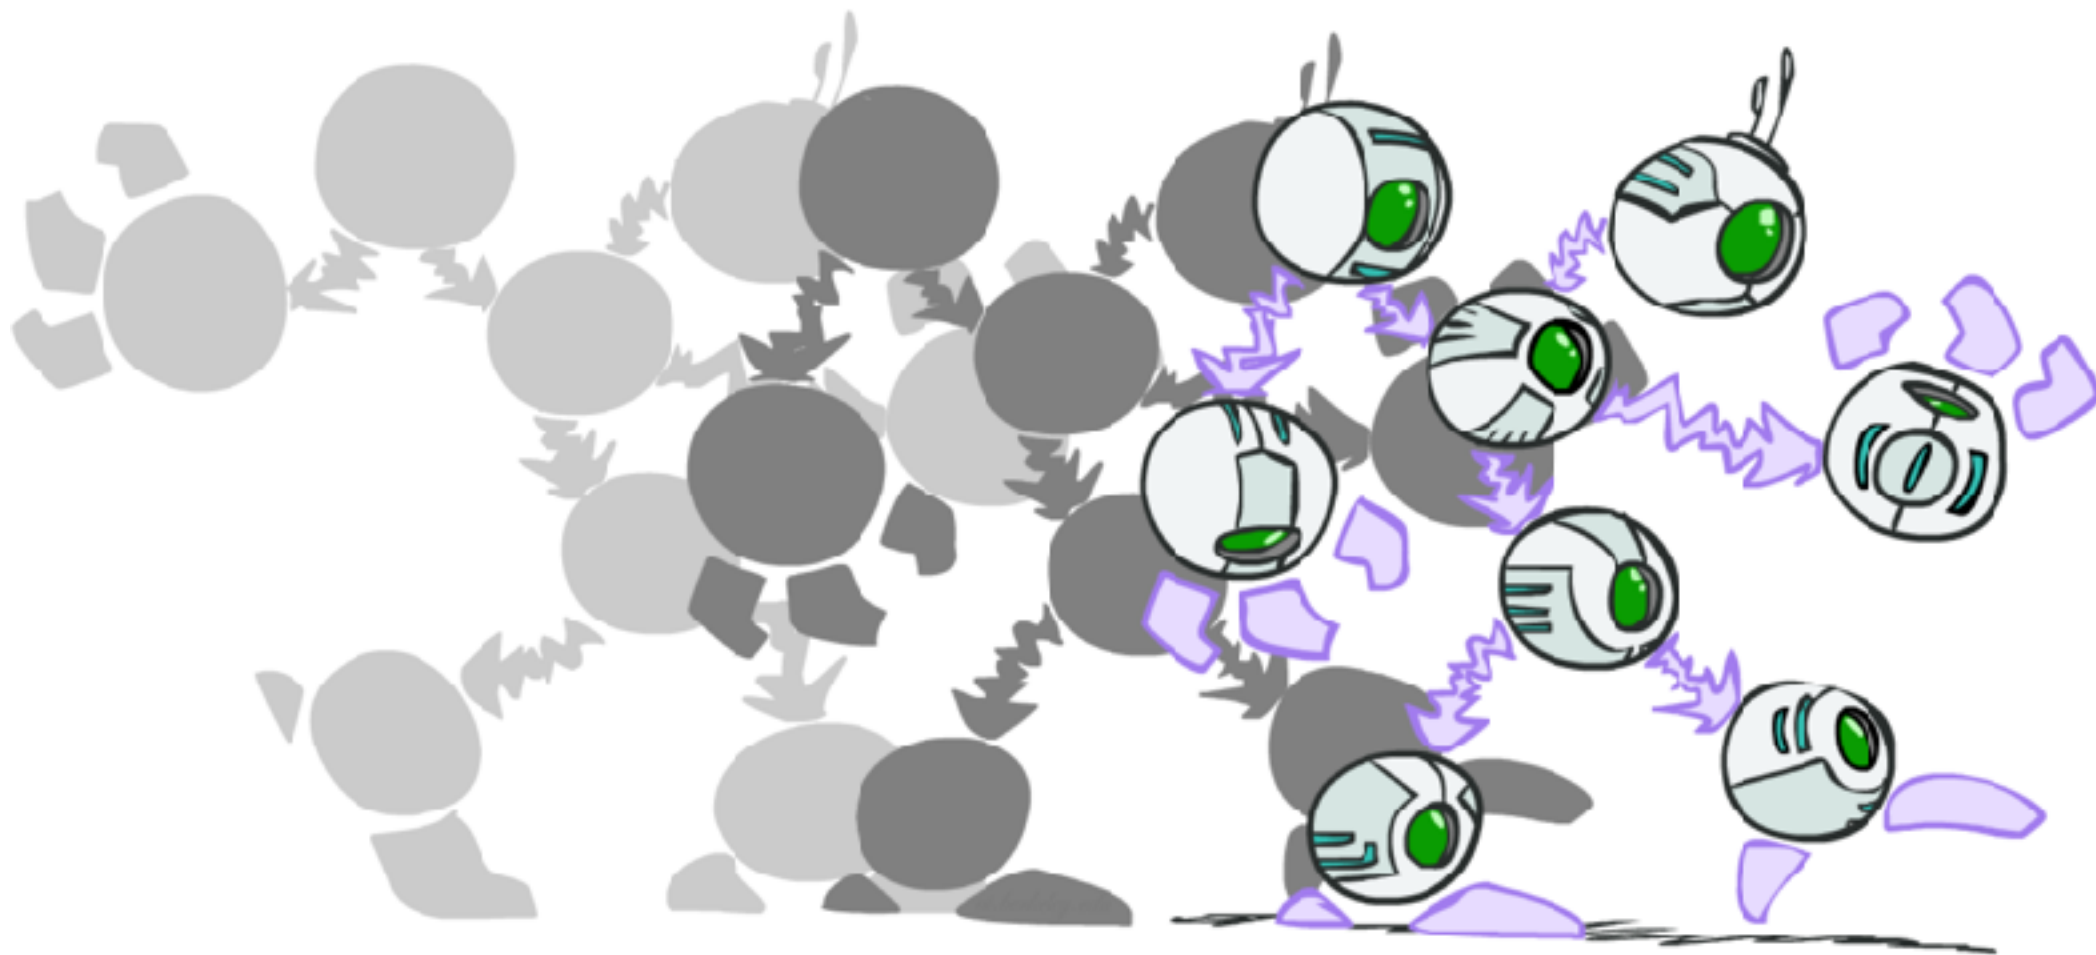


# Particle Filter SLAM – Video 1

---

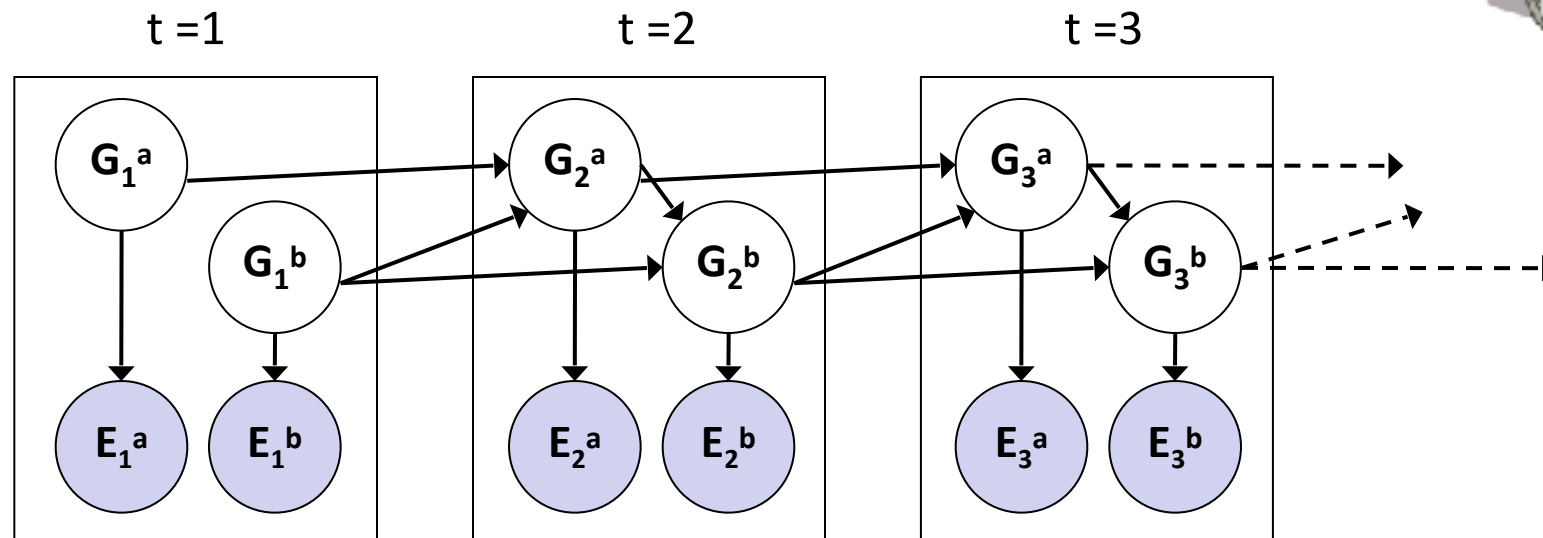


# Dynamic Bayes Nets



# Dynamic Bayes Nets (DBNs)

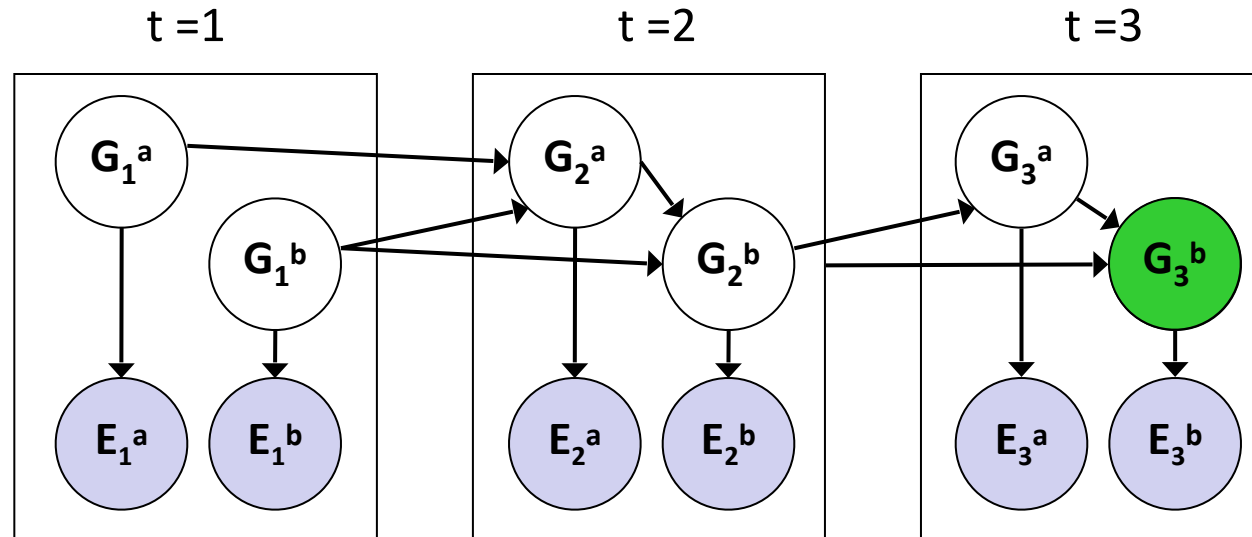
- We want to track multiple variables over time, using multiple sources of evidence
- Idea: Repeat a fixed Bayes net structure at each time
- Variables from time  $t$  can condition on those from  $t-1$



- Dynamic Bayes nets are a generalization of HMMs

# Exact Inference in DBNs

- Variable elimination applies to dynamic Bayes nets
- Procedure: “unroll” the network for  $T$  time steps, then eliminate variables until  $P(X_T | e_{1:T})$  is computed



- Online belief updates: Eliminate all variables from the previous time step; store factors for current time only

# DBN Particle Filters

- A particle is a complete sample for a time step
- **Initialize:** Generate prior samples for the  $t=1$  Bayes net
  - Example particle:  $\mathbf{G}_1^a = (3,3)$   $\mathbf{G}_1^b = (5,3)$
- **Elapse time:** Sample a successor for each particle
  - Example successor:  $\mathbf{G}_2^a = (2,3)$   $\mathbf{G}_2^b = (6,3)$
- **Observe:** Weight each entire sample by the likelihood of the evidence conditioned on the sample
  - Likelihood:  $P(\mathbf{E}_1^a | \mathbf{G}_1^a) * P(\mathbf{E}_1^b | \mathbf{G}_1^b)$
- **Resample:** Select samples (tuples of values) in proportion to their likelihood (weight)

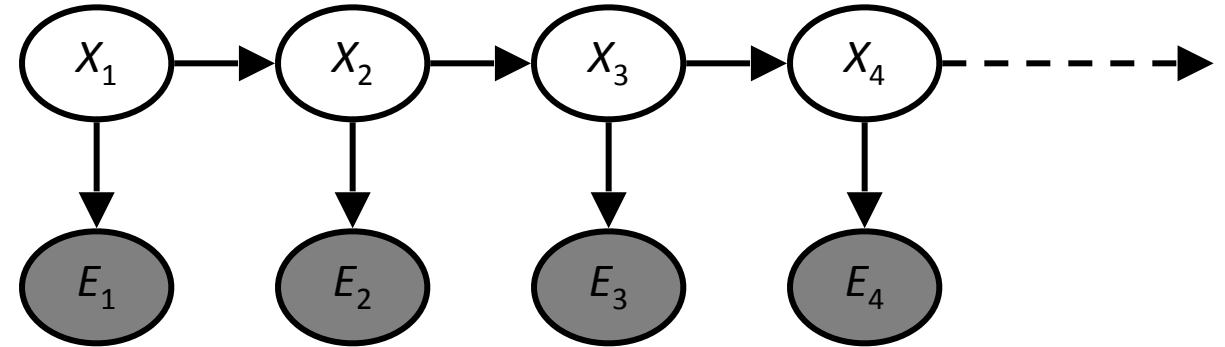
# Most Likely Explanation



# HMMs: MLE Queries

- HMMs defined by

- States  $X$
- Observations  $E$
- Initial distribution:  $P(X_1)$
- Transitions:  $P(X|X_{-1})$
- Emissions:  $P(E|X)$

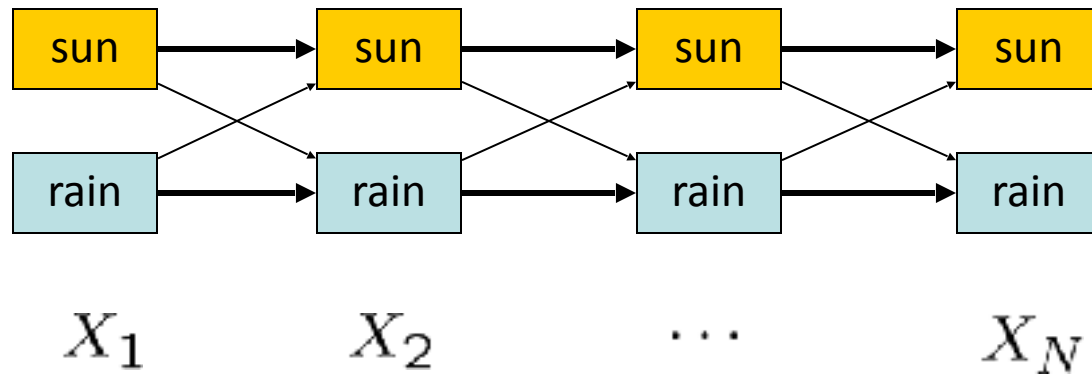


- New query: most likely explanation:  $\arg \max_{x_{1:t}} P(x_{1:t}|e_{1:t})$
- New method: the Viterbi algorithm
- Question: Why not just apply filtering and predict most likely value of each variable separately?



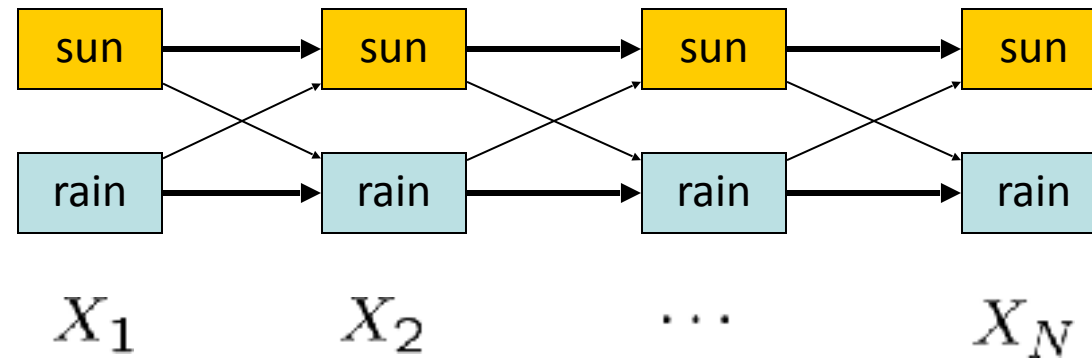
# State Trellis

- State trellis: graph of states and transitions over time



- Each arc represents some transition  $x_{t-1} \rightarrow x_t$
- Each arc has weight  $P(x_t|x_{t-1})P(e_t|x_t)$
- Each path is a sequence of states
- The product of weights on a path is that sequence's probability along with the evidence
- Forward algorithm computes sums of all paths to each node, Viterbi computes best paths
- Exponentially many paths, but dynamic programming can find best path in linear time!

# Forward / Viterbi Algorithms



Forward Algorithm (Sum)

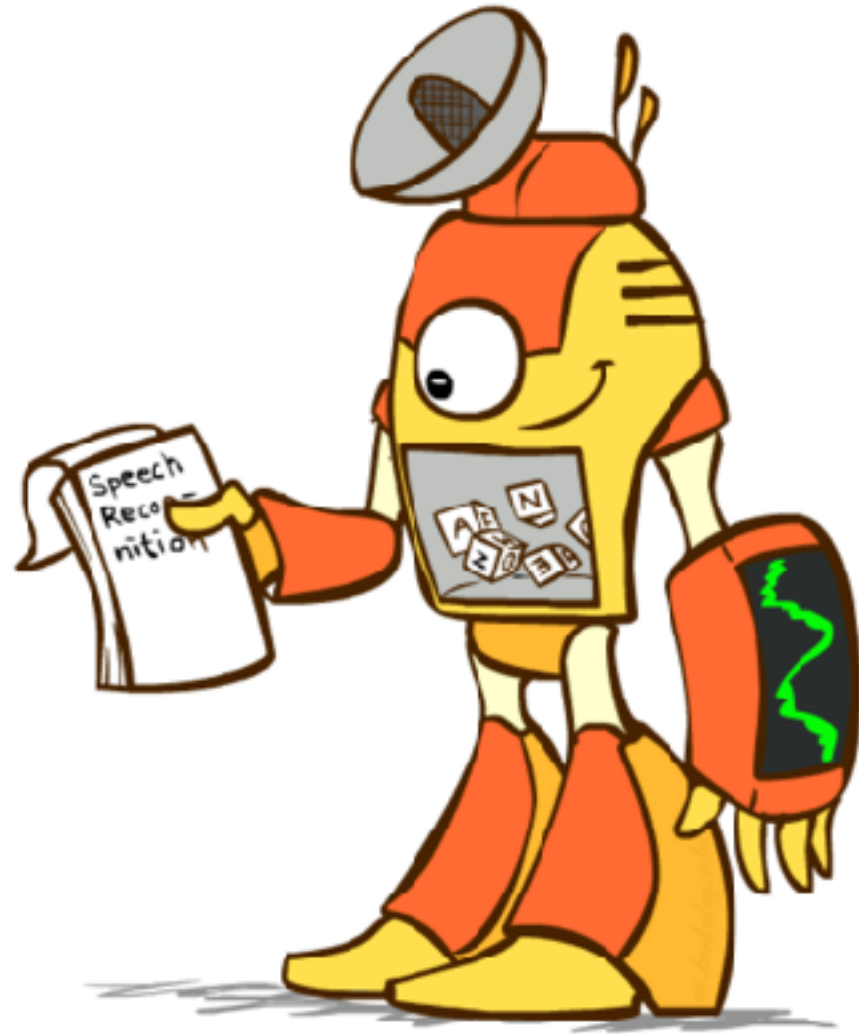
$$\begin{aligned} f_t[x_t] &= P(x_t, e_{1:t}) \\ &= P(e_t|x_t) \sum_{x_{t-1}} P(x_t|x_{t-1}) f_{t-1}[x_{t-1}] \end{aligned}$$

Viterbi Algorithm (Max)

$$\begin{aligned} m_t[x_t] &= \max_{x_{1:t-1}} P(x_{1:t-1}, x_t, e_{1:t}) \\ &= P(e_t|x_t) \max_{x_{t-1}} P(x_t|x_{t-1}) m_{t-1}[x_{t-1}] \end{aligned}$$

# Speech Recognition

---



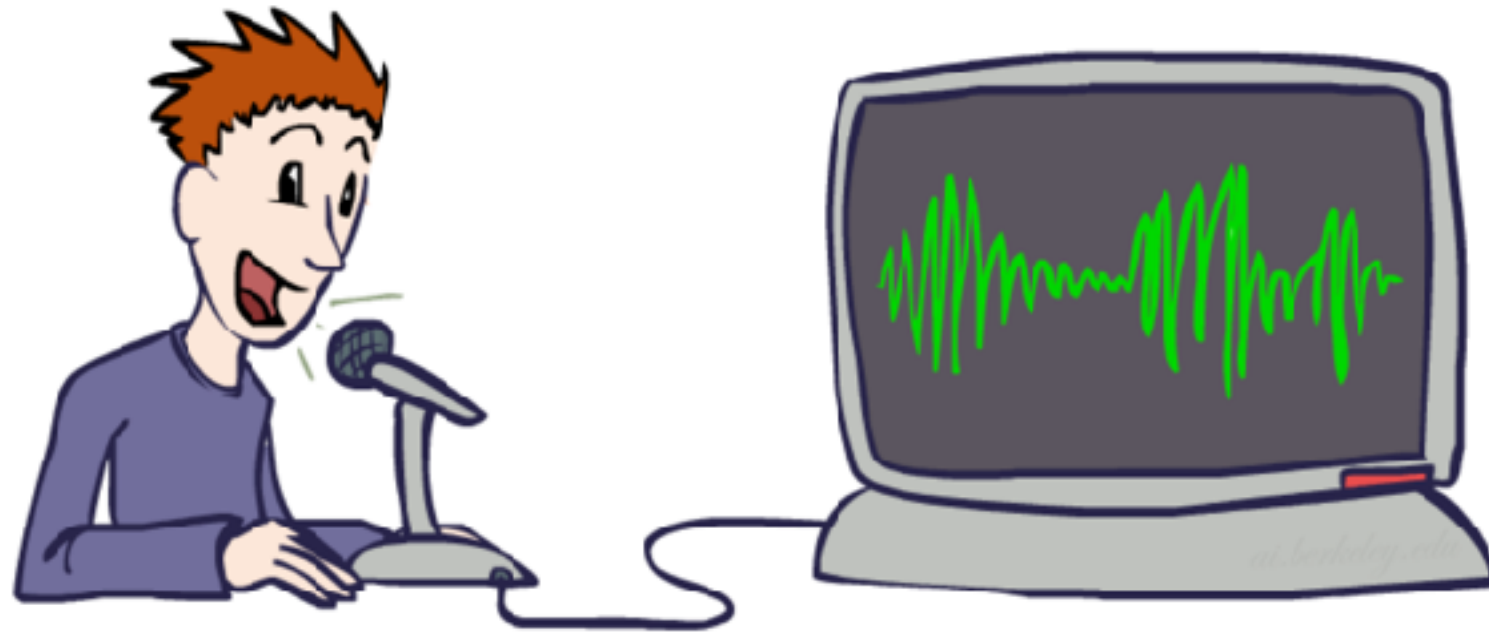
# Speech Recognition in Action

---



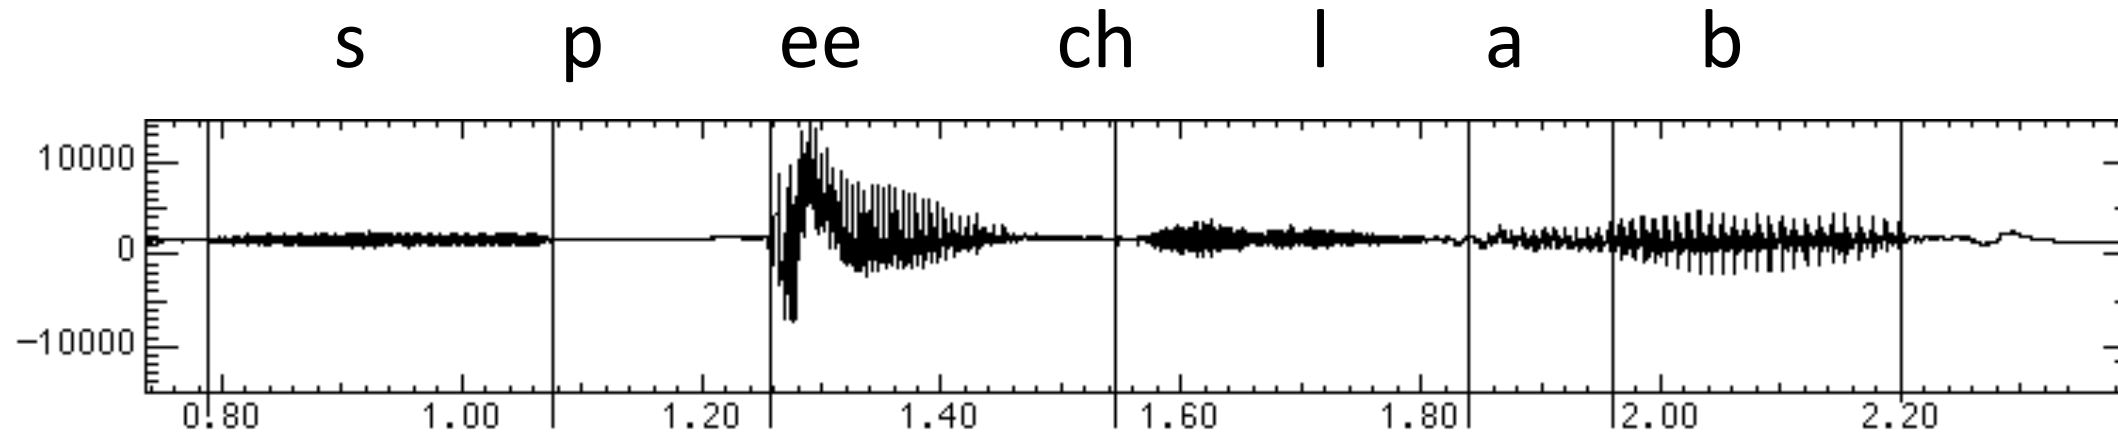
# Digitizing Speech

---



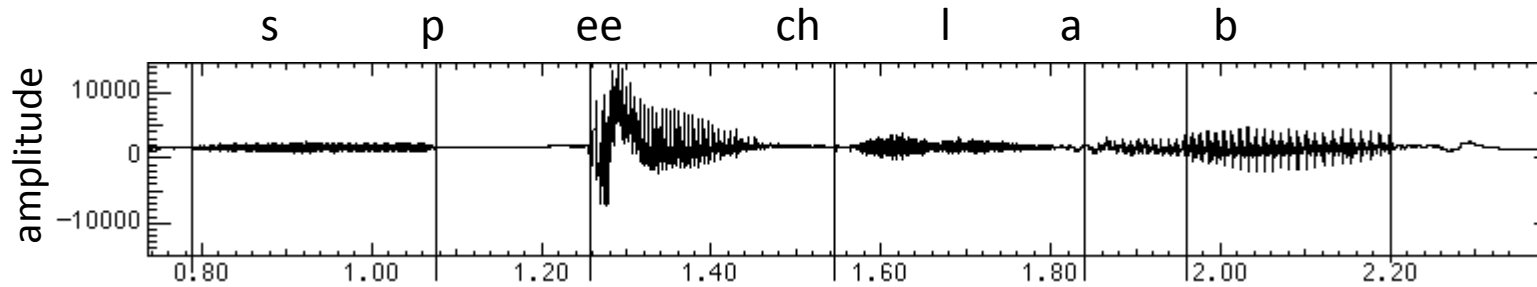
# Speech waveforms

- Speech input is an acoustic waveform

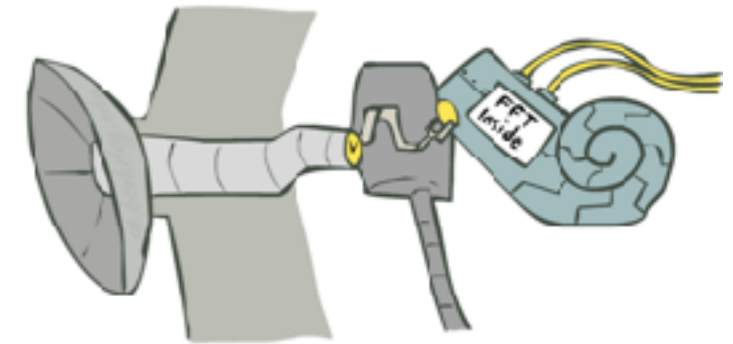
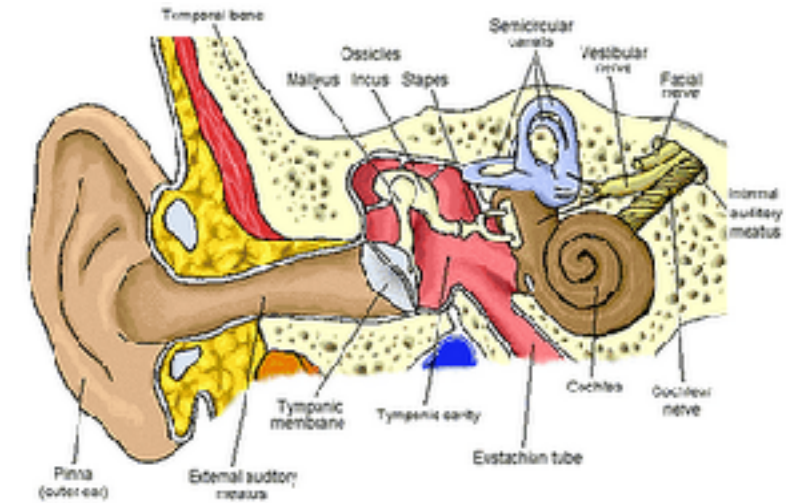
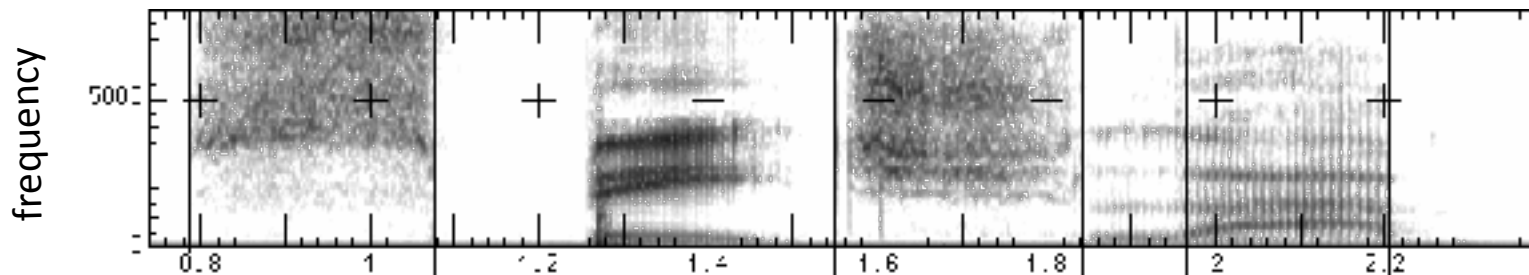


# Spectral Analysis

- Frequency gives pitch; amplitude gives volume
  - Sampling at ~8 kHz (phone), ~16 kHz (mic) (kHz=1000 cycles/sec)

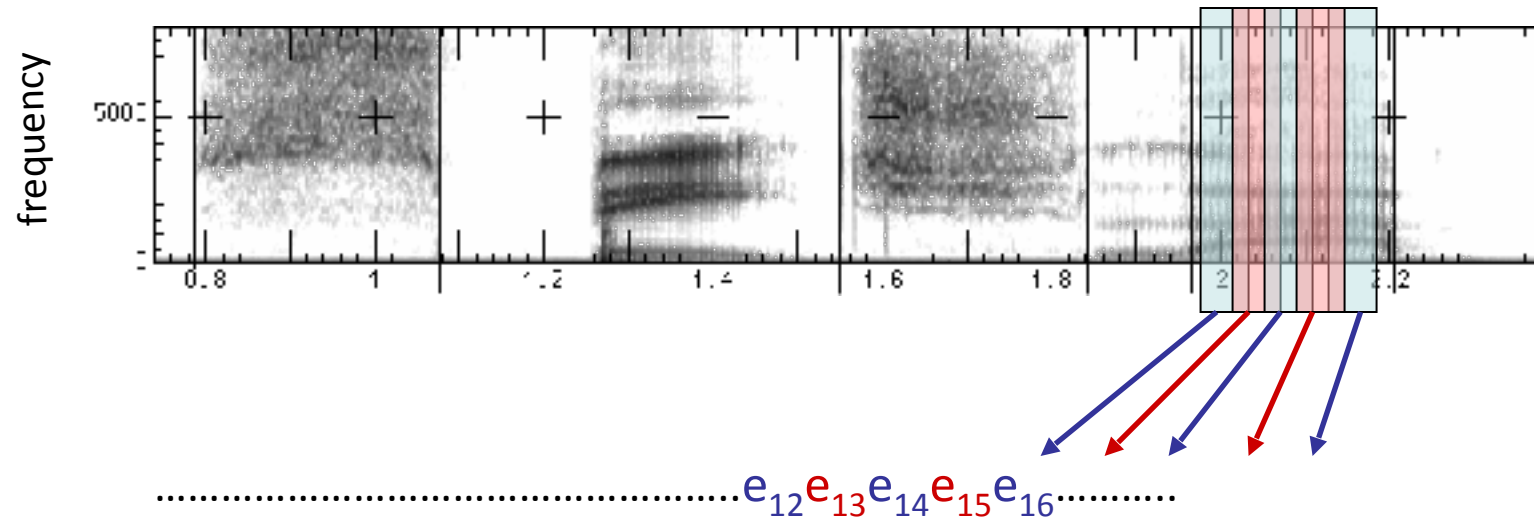


- Fourier transform of wave displayed as a spectrogram
  - Darkness indicates energy at each frequency



# Acoustic Feature Sequence

- Time slices are translated into acoustic feature vectors (~39 real numbers per slice)



- These are the observations  $E$ , now we need the hidden states  $X$



# Speech State Space

---

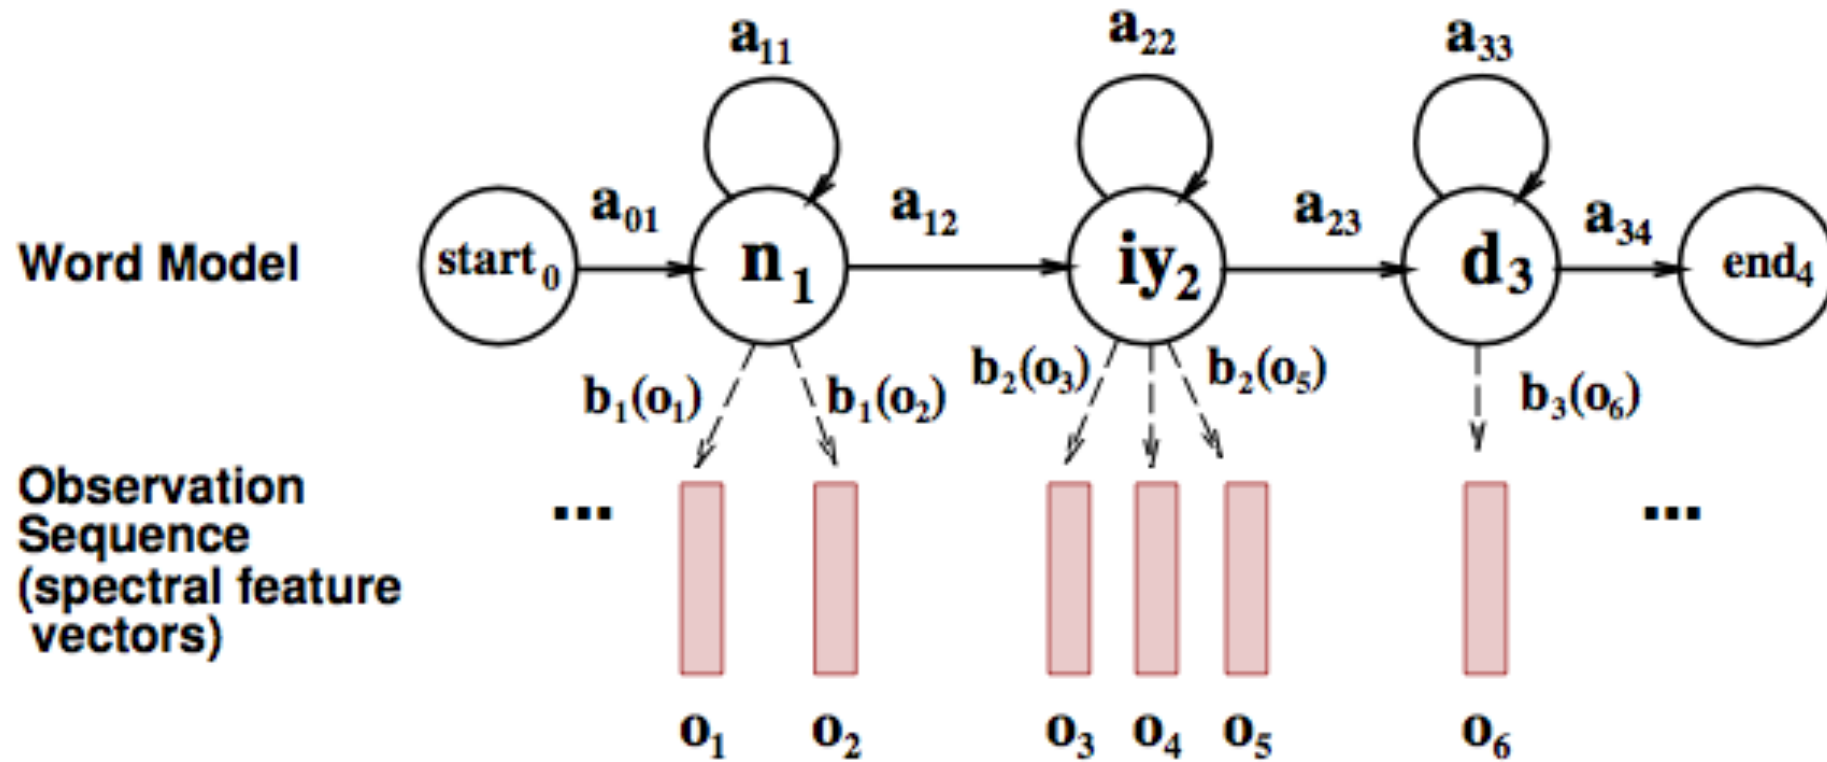
## ■ HMM Specification

- $P(E|X)$  encodes which acoustic vectors are appropriate for each phoneme (each kind of sound)
- $P(X|X')$  encodes how sounds can be strung together

## ■ State Space

- We will have one state for each sound in each word
- Mostly, states advance sound by sound
- Build a little state graph for each word and chain them together to form the state space  $X$

# States in a Word



# Transitions with a Bigram Model

Training Counts

198015222	the first
194623024	the same
168504105	the following
158562063	the world
...	
14112454	the door
-----	
23135851162	the *

$$\hat{P}(\text{door}|\text{the}) = \frac{14112454}{23135851162}$$
$$= 0.0006$$

# Decoding (Viterbi)

- Finding the words given the acoustics is an HMM inference problem
- Which state sequence  $x_{1:T}$  is most likely given the evidence  $e_{1:T}$ ?

$$x_{1:T}^* = \arg \max_{x_{1:T}} P(x_{1:T}|e_{1:T}) = \arg \max_{x_{1:T}} P(x_{1:T}, e_{1:T})$$

- From the sequence  $x$ , we can simply read off the words

