# Markov Models and Hidden Markov Models

Robert Platt
Northeastern University

Some images and slides are used from:
1. CS188 UC Berkeley
2. RN, AIMA

# Markov Models

We have already seen that an MDP provides a useful framework for modeling stochastic control problems.

Markov Models: model any kind of temporally dynamic system.

# Probability again: Independence

Two random variables, *x* and *y*, are independent when:

$$\forall(x, y), P(x, y) = P(x)P(y) \quad \Longleftrightarrow \quad x \perp\!\!\!\perp y$$

$$x \not\perp\!\!\!\perp y$$

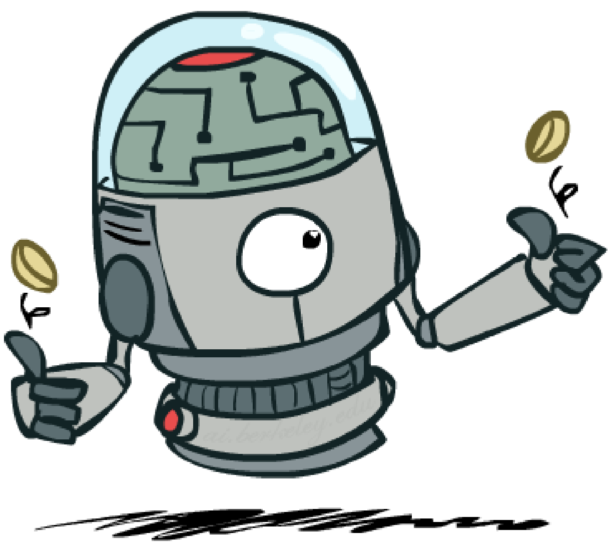

The outcomes of two different coin flips are usually independent of each other

Image: Berkeley CS188 course notes (downloaded Summer 2015)

# Probability again: Independence

If: $$P(x, y) = P(x)P(y)$$

Then: $$P(x) = P(x|y)$$

$$P(y) = P(y|x)$$

Why?

# Are T and W independent?

$$P_1(T, W)$$

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

# Are T and W independent?

$P(T)$

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

$P_1(T, W)$

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

$P(W)$

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

# Are T and W independent?

$P(T)$

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

$P_1(T, W)$

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

$P_2(T, W) = P(T)P(W)$

| T | W | P |
|---|---|---|
| hot | sun | 0.3 |
| hot | rain | 0.2 |
| cold | sun | 0.3 |
| cold | rain | 0.2 |

$P(W)$

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

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Conditional independence

Independence: $$\forall(x, y), P(x, y) = P(x)P(y)$$

$$x \perp\!\!\!\perp y$$

Conditional independence: $$\forall(x, y, z), P(x, y|z) = P(x|z)P(y|z)$$

$$x \perp\!\!\!\perp y|z$$

Equivalent statements of conditional independence: $$P(x|z) = P(x|z, y)$$

$$P(y|z) = P(y|z, x)$$

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Conditional independence: example

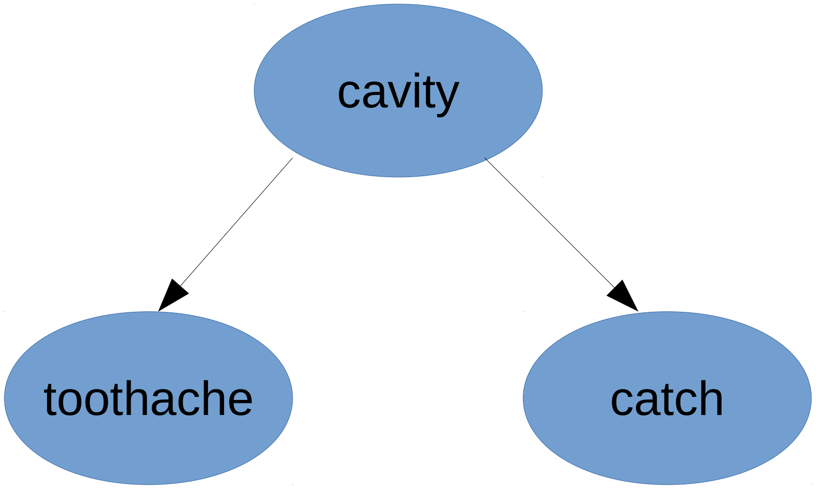


P(toothache, catch | cavity) = P(toothache | cavity) = P(catch | cavity)

or...

P(toothache | cavity) = P(toothache | cavity, catch)

P(catch | cavity) = P(catch | cavity, toothache)

# Conditional independence: example

- What about this domain:
  - Traffic
  - Umbrella
  - Raining

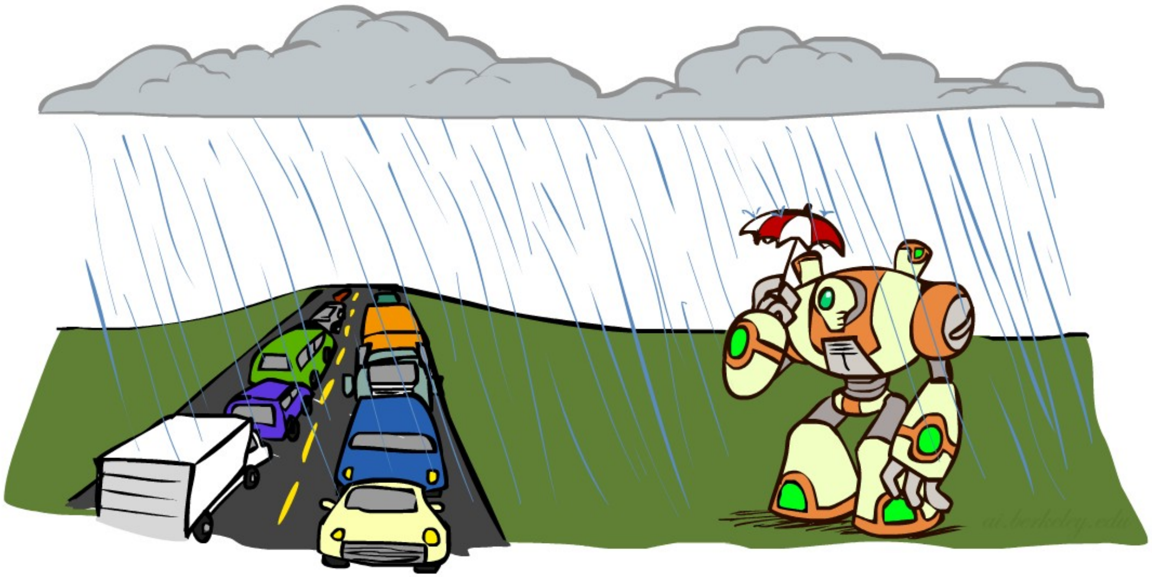

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Conditional independence: example

- What about this domain:
  - Fire
  - Smoke
  - Alarm

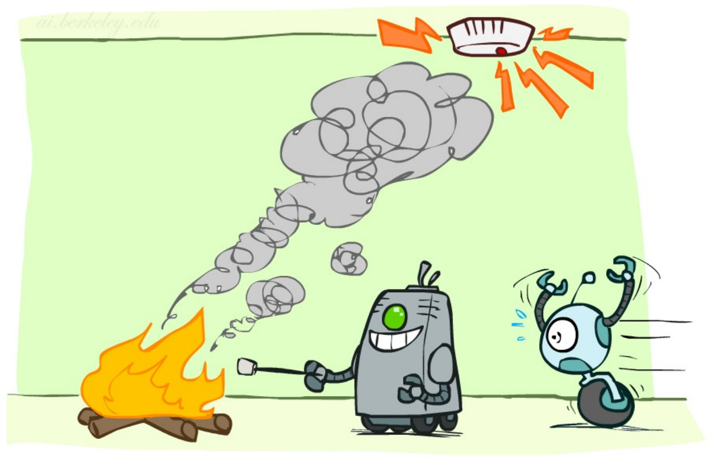

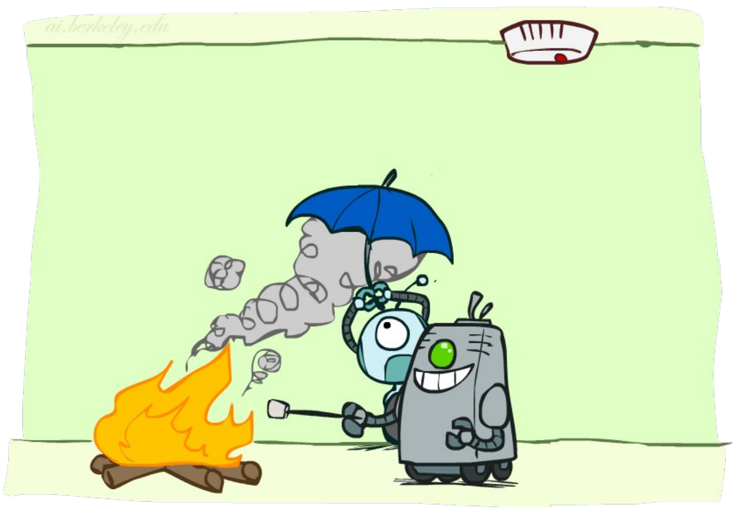

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Markov Processes

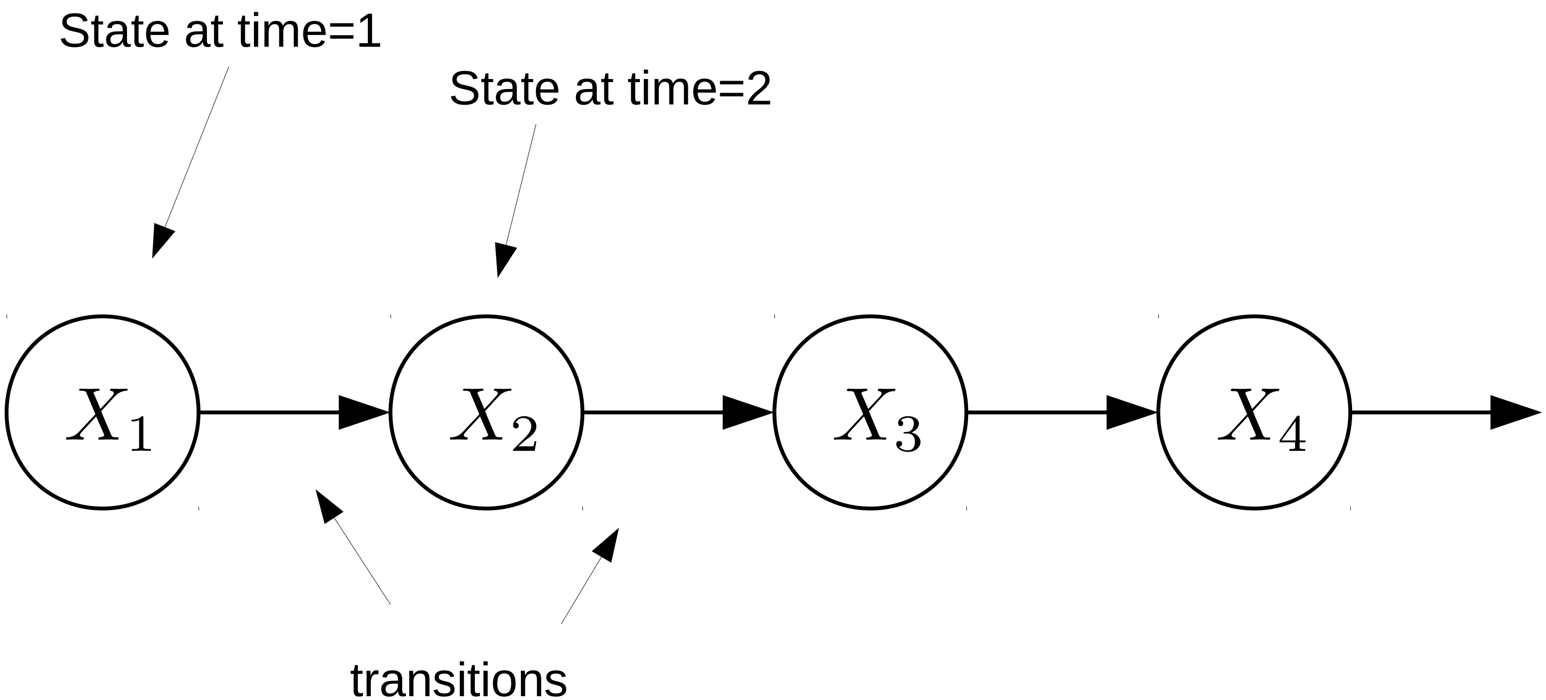

# Markov Processes

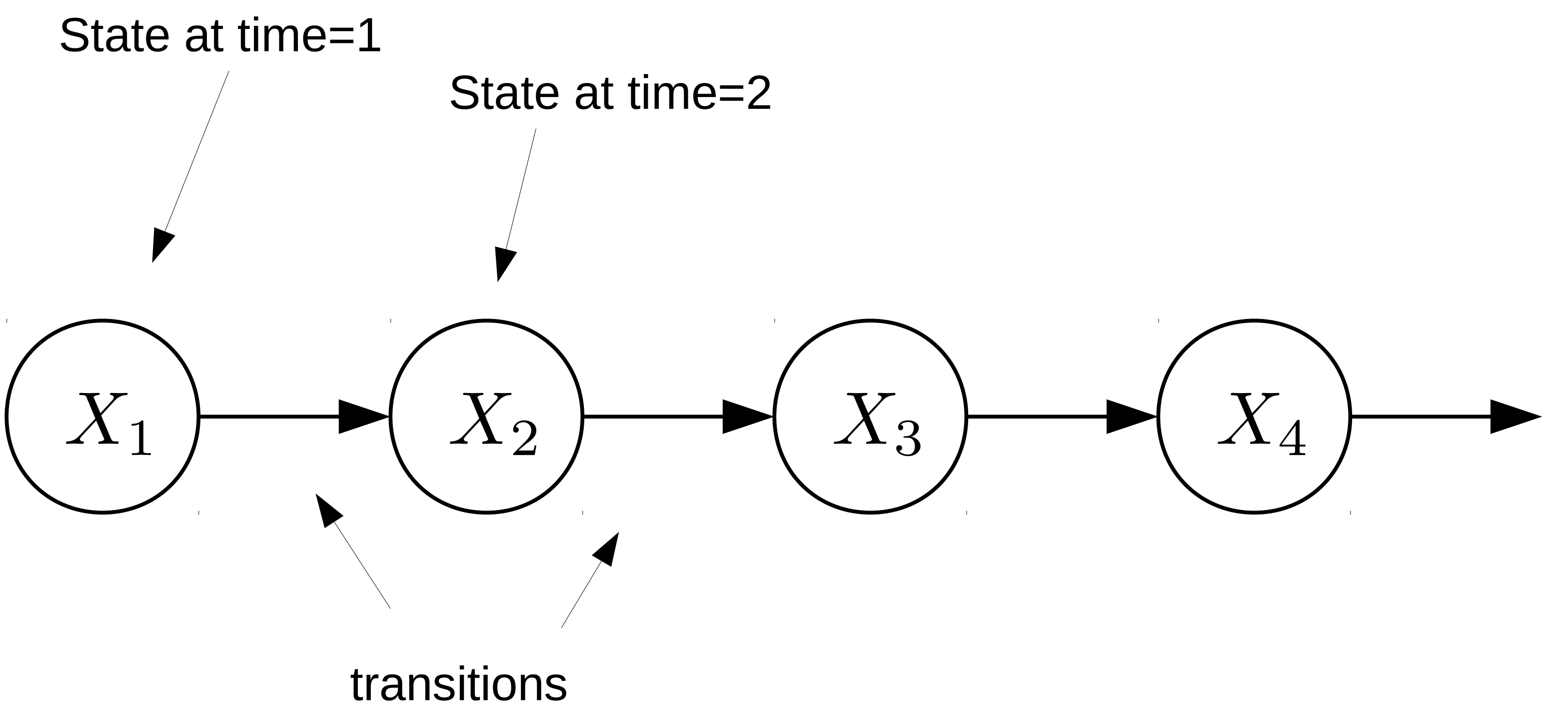


Since this is a Markov process, we assume transitions are Markov:

Process model: $$P(X_t|X_{t-1}) = P(X_t|X_{t-1}, \dots, X_1)$$

Markov assumption: $$X_t \perp\!\!\!\perp X_{t-2}|X_{t-1}$$

# Markov Processes

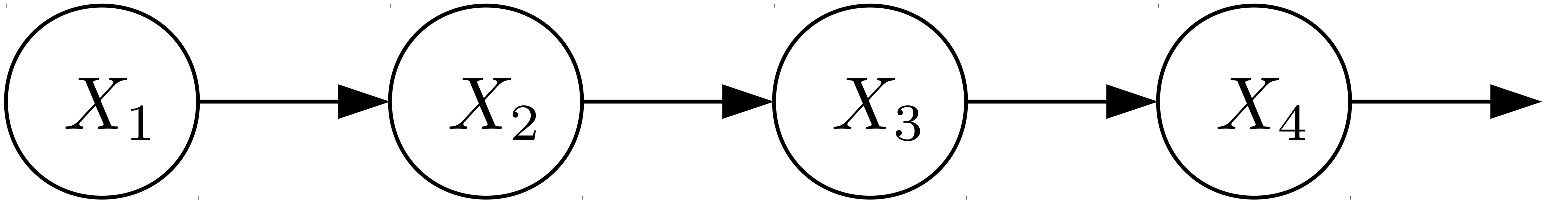


How do we calculate: $P(X_1, X_2, X_3, X_4) = ?$

$$P(X_1, X_2, X_3, X_4) = P(X_1)P(X_2|X_1)P(X_3|X_2, X_1)P(X_4|X_3, X_2, X_1)$$

# Markov Processes

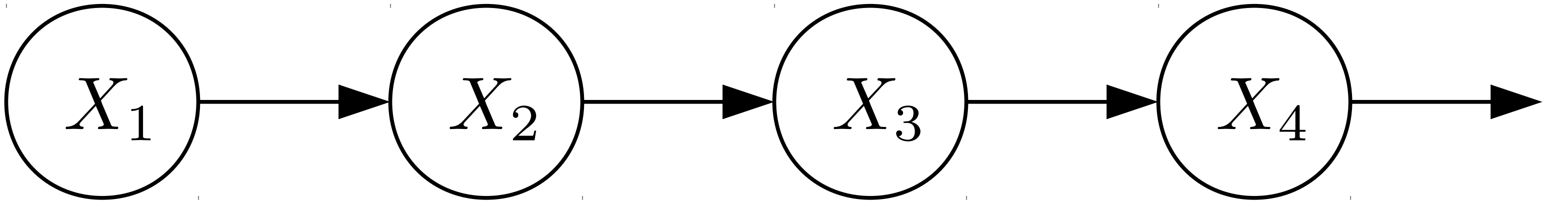


How do we calculate: $P(X_1, X_2, X_3, X_4) = ?$

$$P(X_1, X_2, X_3, X_4) = \underbrace{P(X_1)P(X_2|X_1)}_{P(X_2, X_1)}P(X_3|X_2, X_1)P(X_4|X_3, X_2, X_1)$$

# Markov Processes

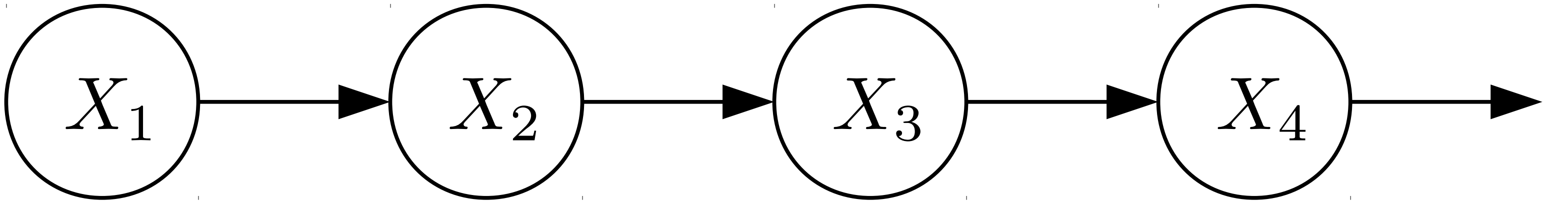

How do we calculate: $P(X_1, X_2, X_3, X_4) =?$

$$P(X_1, X_2, X_3, X_4) = \underbrace{P(X_1)P(X_2|X_1)P(X_3|X_2, X_1)}_{P(X_3, X_2, X_1)}P(X_4|X_3, X_2, X_1)$$

# Markov Processes

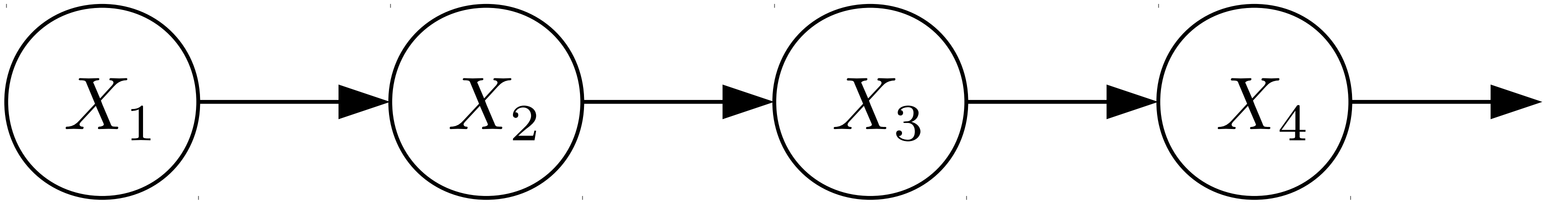


How do we calculate: $P(X_1, X_2, X_3, X_4) = ?$

$$P(X_1, X_2, X_3, X_4) = P(X_1)P(X_2|X_1)P(X_3|X_2, X_1)P(X_4|X_3, X_2, X_1)$$

Can we simplify this expression?

# Markov Processes

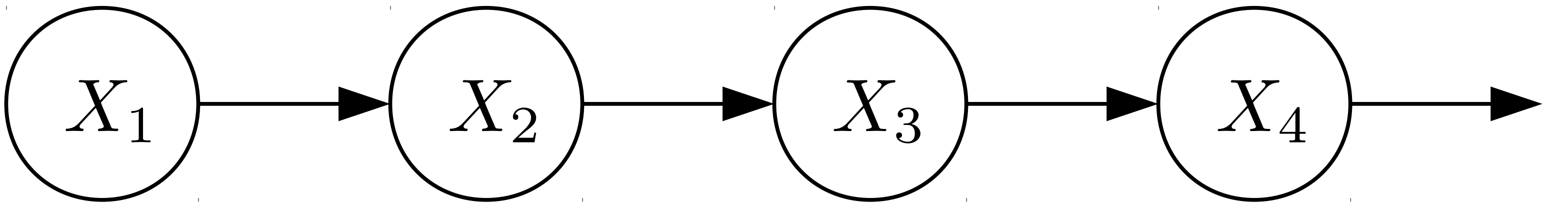

How do we calculate: $P(X_1, X_2, X_3, X_4) = ?$

$$P(X_1, X_2, X_3, X_4) = P(X_1)P(X_2|X_1)\underbrace{P(X_3|X_2, X_1)}_{P(X_3|X_2)}\underbrace{P(X_4|X_3, X_2, X_1)}_{P(X_4|X_3)}$$

# Markov Processes

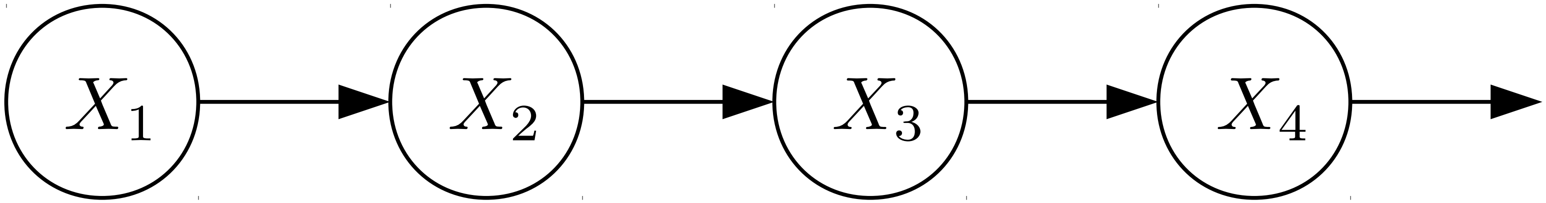


How do we calculate: $P(X_1, X_2, X_3, X_4) = ?$

$$P(X_1, X_2, X_3, X_4) = P(X_1)P(X_2|X_1)P(X_3|X_2, X_1)P(X_4|X_3, X_2, X_1)$$

$$P(X_1, X_2, X_3, X_4) = P(X_1)P(X_2|X_1)P(X_3|X_2)P(X_4|X_3)$$

# Markov Processes

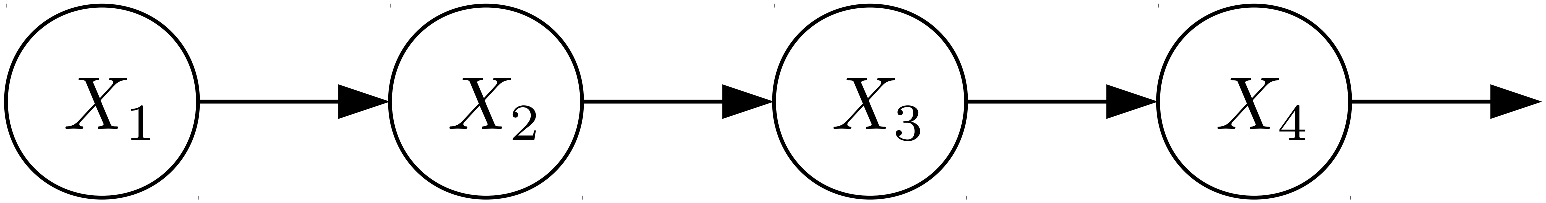


How do we calculate: $P(X_1, X_2, X_3, X_4) = ?$

$$P(X_1, X_2, X_3, X_4) = P(X_1)P(X_2|X_1)P(X_3|X_2, X_1)P(X_4|X_3, X_2, X_1)$$

$$P(X_1, X_2, X_3, X_4) = P(X_1)P(X_2|X_1)P(X_3|X_2)P(X_4|X_3)$$

In general: $P(X_1, X_2, \ldots, X_T) = P(X_1) \prod_{t=1}^{T-1} P(X_{t+1}|X_t)$

# Markov Processes

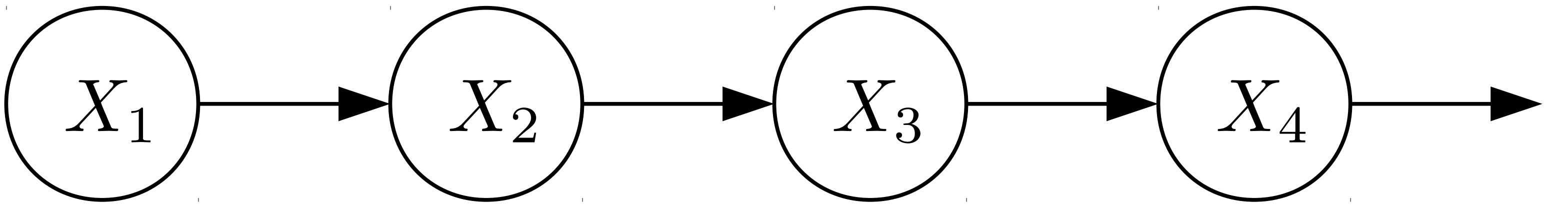


How do we calculate: $P(X_1, X_2, X_3, X_4) = ?$

$$P(X_1, X_2, X_3, X_4) = P(X_1)P(X_2|X_1)P(X \ldots X_2, X_1)$$

Process model

$$P(X_1, X_2, X_3, X_4) = P(X_1)P(X_2|X_1)P(X_3|X_2)P(X_4|X_3)$$

In general: $P(X_1, X_2, \ldots, X_T) = P(X_1) \prod_{t=1}^{T-1} P(X_{t+1}|X_t)$

# Markov Processes: example

- States: X = {rain, sun}
- Initial distribution: 1.0 sun
- Process model: $P(X_t \mid X_{t-1})$:

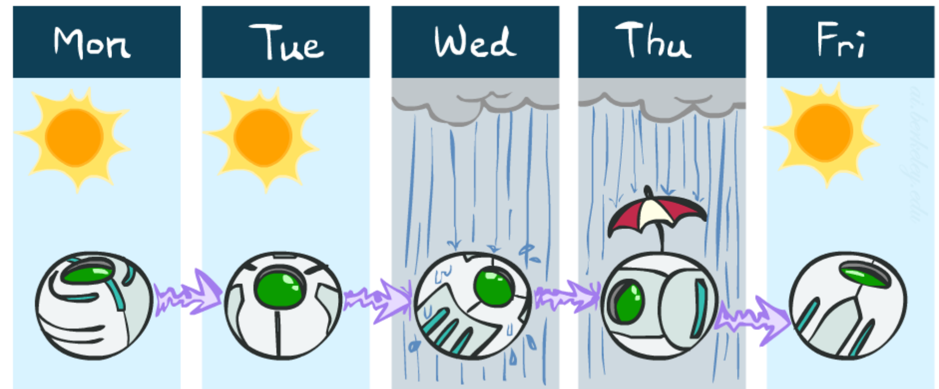


| $X_{t-1}$ | $X_t$ | $P(X_t \mid X_{t-1})$ |
|---|---|---|
| sun | sun | 0.9 |
| sun | rain | 0.1 |
| rain | sun | 0.3 |
| rain | rain | 0.7 |

Two new ways of representing the same CPT

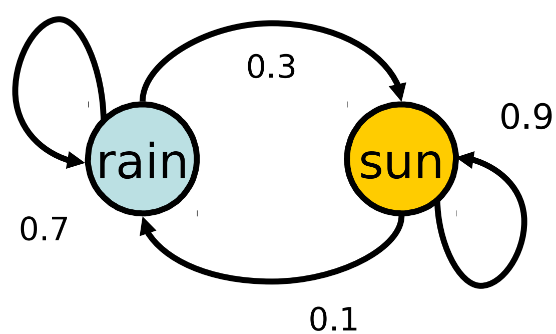


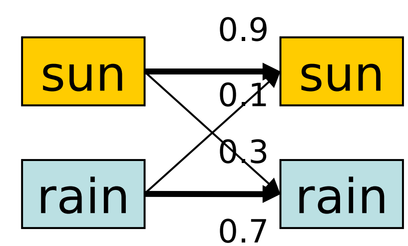


Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Simulating dynamics forward

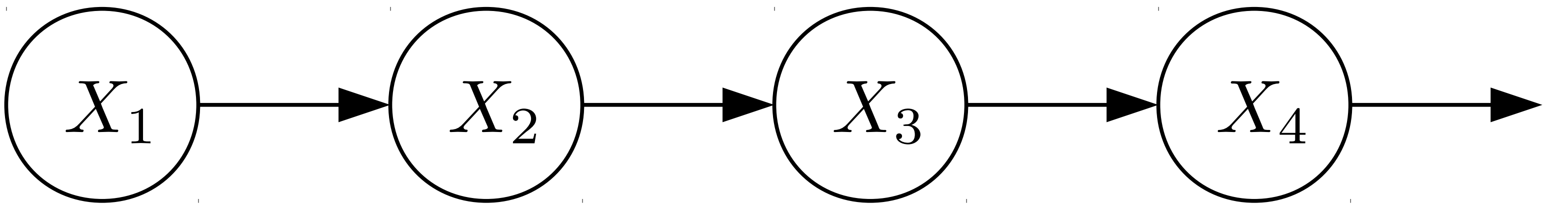


Joint distribution: $$P(X_1, X_2, \ldots, X_T) = P(X_1) \prod_{t=1}^{T-1} P(X_{t+1}|X_t)$$

But, suppose we want to predict the state at time T, given a prior distribution at time 1?

$$P(X_2) = \sum_{X_1} P(X_1)P(X_2|X_1)$$

$$P(X_3) = \sum_{X_2} P(X_2)P(X_3|X_2)$$

$$\vdots$$

$$P(X_T) = \sum_{X_{T-1}} P(X_{T-1})P(X_T|X_{T-1})$$

# Markov Processes: example

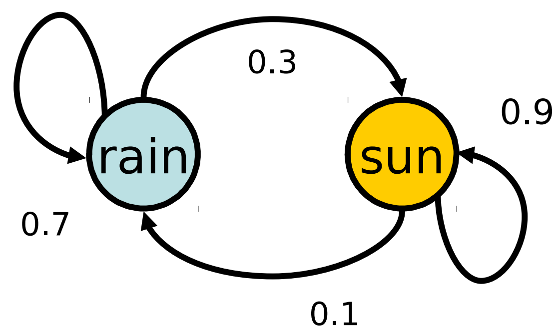


- Initial distribution: 1.0 sun

- What is the probability distribution after one step?

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

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

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Simulating dynamics forward

- From initial observation of sun

$$\left\langle \begin{array}{c} 1.0 \\ 0.0 \end{array} \right\rangle \quad \left\langle \begin{array}{c} 0.9 \\ 0.1 \end{array} \right\rangle \quad \left\langle \begin{array}{c} 0.84 \\ 0.16 \end{array} \right\rangle \quad \left\langle \begin{array}{c} 0.804 \\ 0.196 \end{array} \right\rangle \Rightarrow \left\langle \begin{array}{c} 0.75 \\ 0.25 \end{array} \right\rangle$$

$P(X_1)$ $P(X_2)$ $P(X_3)$ $P(X_4)$ $P(X_\infty)$

- From initial observation of rain

$$\left\langle \begin{array}{c} 0.0 \\ 1.0 \end{array} \right\rangle \quad \left\langle \begin{array}{c} 0.3 \\ 0.7 \end{array} \right\rangle \quad \left\langle \begin{array}{c} 0.48 \\ 0.52 \end{array} \right\rangle \quad \left\langle \begin{array}{c} 0.588 \\ 0.412 \end{array} \right\rangle \Rightarrow \left\langle \begin{array}{c} 0.75 \\ 0.25 \end{array} \right\rangle$$

$P(X_1)$ $P(X_2)$ $P(X_3)$ $P(X_4)$ $P(X_\infty)$

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

$$\left\langle \begin{array}{c} p \\ 1-p \end{array} \right\rangle \quad \cdots \Rightarrow \left\langle \begin{array}{c} 0.75 \\ 0.25 \end{array} \right\rangle$$

$P(X_1)$ $P(X_\infty)$

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Simulating dynamics forward

- From initial observation of sun

$$\begin{pmatrix} 1.0 \\ 0.0 \end{pmatrix} \quad \begin{pmatrix} 0.9 \\ 0.1 \end{pmatrix} \quad \begin{pmatrix} 0.84 \\ 0.16 \end{pmatrix} \quad \begin{pmatrix} 0.804 \\ 0.196 \end{pmatrix} \Rightarrow \begin{pmatrix} 0.75 \\ 0.25 \end{pmatrix}$$

$P(X_1)$ $P(X_2)$ $P(X_3)$ $P(X_4)$ $P(X_\infty)$

- From i … in

$$\begin{pmatrix} 0.0 \\ 1.0 \end{pmatrix} \quad \ldots \quad \begin{pmatrix} \ldots 88 \\ \ldots 12 \end{pmatrix} \Rightarrow \begin{pmatrix} 0.75 \\ 0.25 \end{pmatrix}$$

$P(X_1)$ … $X_4)$ $P(X_\infty)$

- From y … bution $P(X_1)$:

$$\begin{pmatrix} p \\ 1-p \end{pmatrix} \quad \ldots \quad \Rightarrow \begin{pmatrix} 0.75 \\ 0.25 \end{pmatrix}$$

$P(X_1)$ $P(X_\infty)$

This is called the *stationary distribution*

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Hidden Markov Models (HMMs)

<u>Hidden Markov Models</u>: markov models applied to estimation problems

- speech to text
- tracking in computer vision
- robot localization

# Hidden Markov Models (HMMs)

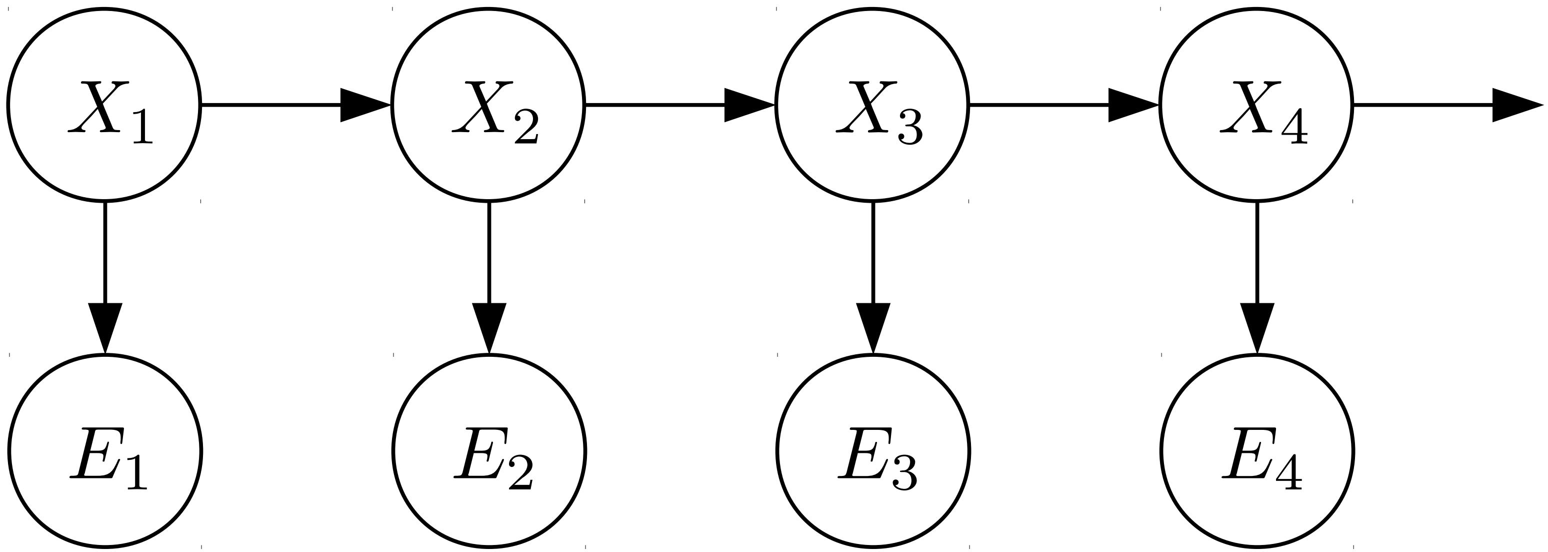


Called an “emission”

State, $X_t$ , is assumed to be unobserved

However, you get to make one observation, $E_t$ , on each timestep.

# Hidden Markov Models (HMMs)

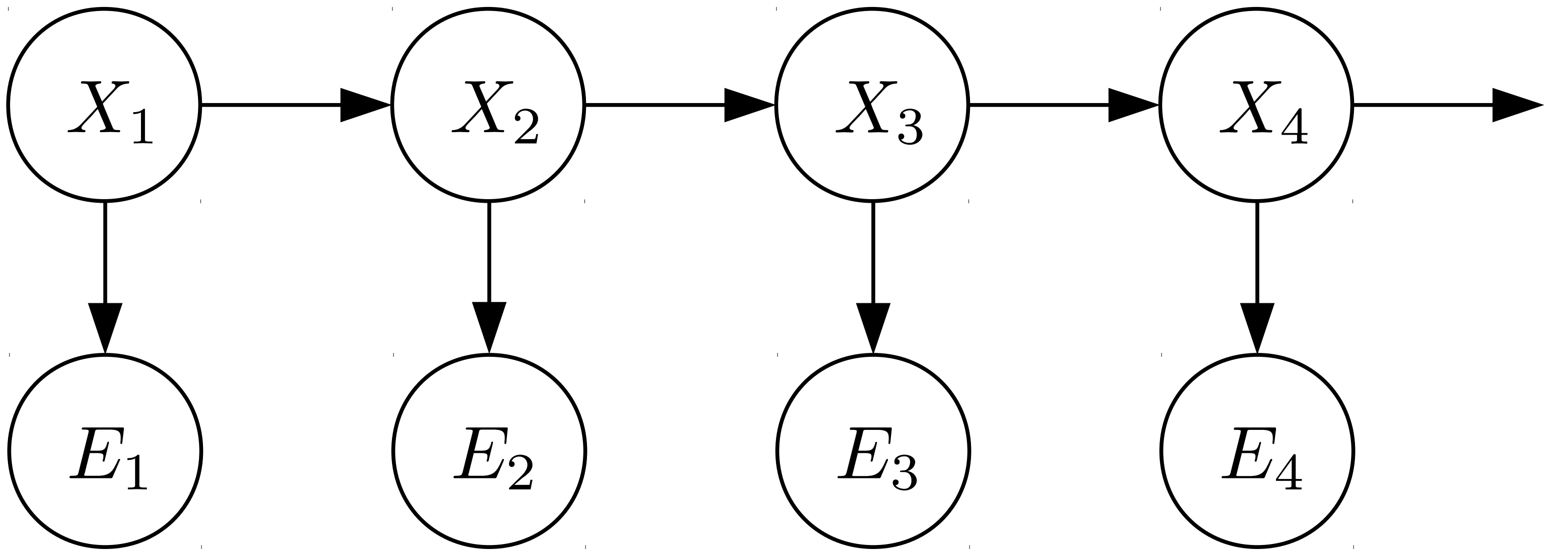

Sensor Markov Assumption: the current observation depends only on current state:

$$P(E_t|X_t, X_{t-1}, \ldots, X_1) = P(E_t|X_t)$$

$$E_t \perp\!\!\!\perp X_{t-1}|X_t$$

# HMM example

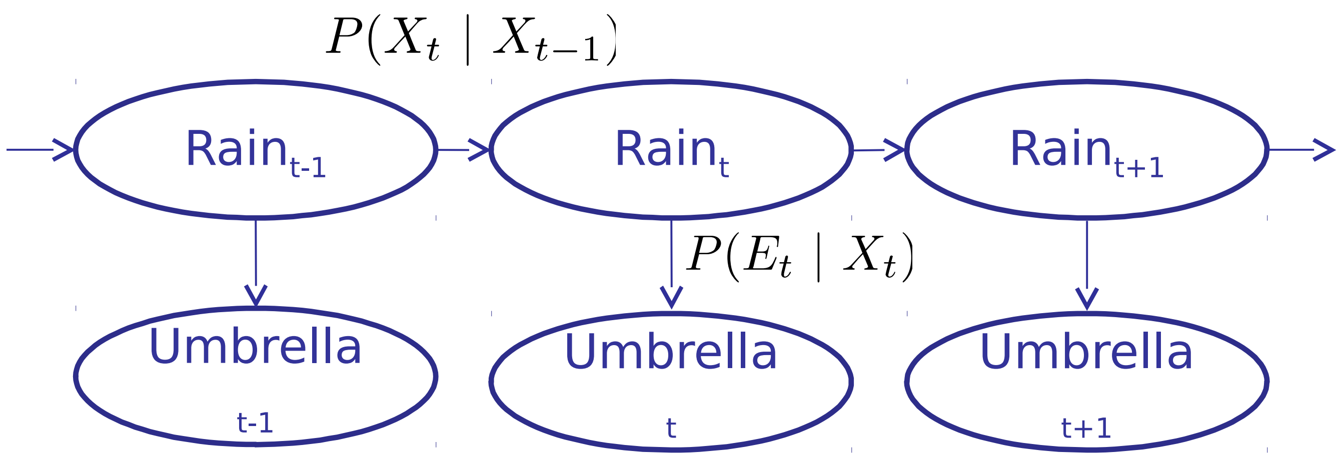


| $R_t$ | $R_{t+1}$ | $P(R_{t+1}\|R_t)$ |
|---|---|---|
| +r | +r | 0.7 |
| +r | -r | 0.3 |
| -r | +r | 0.3 |
| -r | -r | 0.7 |

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

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

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Real world HMM applications

- 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)

# HMM Filtering

- 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

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# HMM Filtering

Given a prior distribution, $P(X_1)$ , and a series of observations, $E_1, \ldots, E_T$ , calculate the posterior distribution: $P(X_t|E_1, \ldots, E_T)$

Two steps:

$P(X_t|E_{1:t})$ → $P(X_{t+1}|E_{1:t})$ → $P(X_{t+1}|E_{1:t+1})$

Process update          Observation update

# HMM Filtering

Given a prior distribution, $P(X_1)$ , and a series of observations, $E_1, \ldots, E_T$ , calculate the posterior distribution: $P(X_t|E_1, \ldots, E_T)$

Two steps:

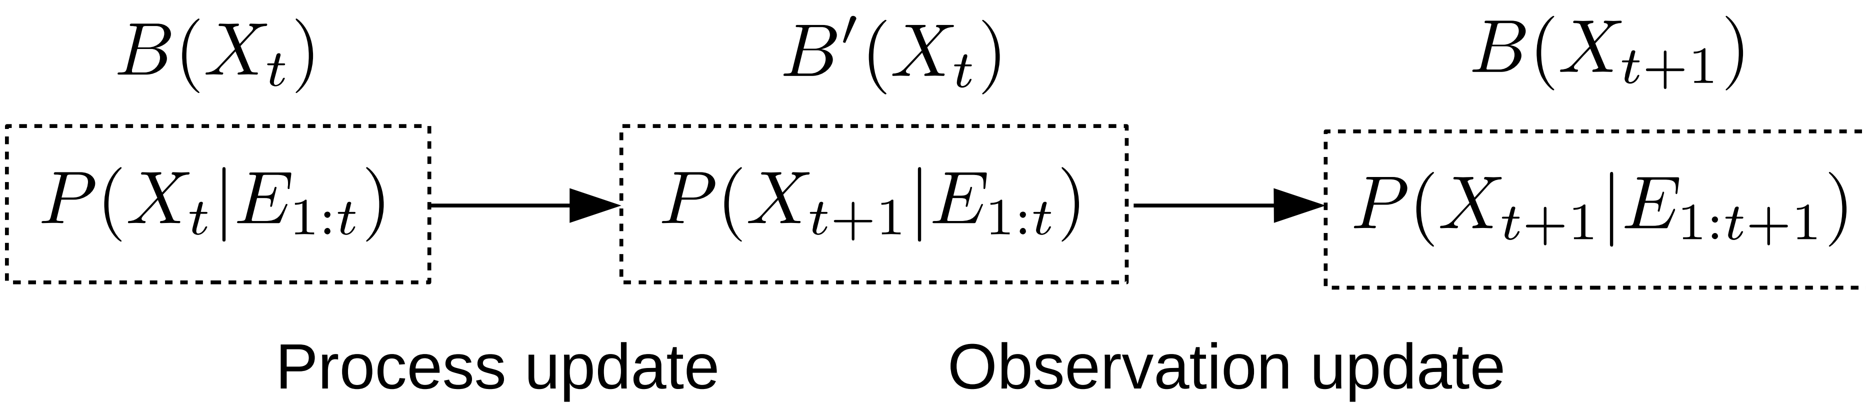

# HMM Filtering

Given a prior distribution, $P(X_1)$ , and a series of observations, $E_1, \ldots, E_T$ , calculate the posterior distribution: $P$

"Beliefs"

Two steps:

$B(X_t)$ $B'(X_t)$ $B(X_{t+1})$

$P(X_t|E_{1:t})$ → $P(X_{t+1}|E_{1:t})$ → $P(X_{t+1}|E_{1:t+1})$

Process update Observation update

# Process update

$B(X_t)$ : $P(X_t|e_{1:t})$ $\longrightarrow$ $B'(X_t)$ : $P(X_{t+1}|e_{1:t})$

$$P(X_{t+1}|e_{1:t}) = \sum_{X_t} P(X_{t+1}|X_t, e_{1:t}) P(X_t|e_{1:t})$$

$$B'(X_{t+1}) = \sum_{X_t} P(X_{t+1}|X_t, e_{1:t}) B(X_t)$$

This is just forward simulation of the Markov Model

# Process update: example

- As time passes, uncertainty "accumulates"

(Transition model: ghosts usually go clockwise)

| | | | | | |
|---|---|---|---|---|---|
| <0.01 | <0.01 | <0.01 | <0.01 | <0.01 | <0.01 |
| <0.01 | <0.01 | <0.01 | <0.01 | <0.01 | <0.01 |
| <0.01 | <0.01 | 1.00 | <0.01 | <0.01 | <0.01 |
| <0.01 | <0.01 | <0.01 | <0.01 | <0.01 | <0.01 |

T = 1

| | | | | | |
|---|---|---|---|---|---|
| <0.01 | <0.01 | <0.01 | <0.01 | <0.01 | <0.01 |
| <0.01 | <0.01 | 0.06 | <0.01 | <0.01 | <0.01 |
| <0.01 | 0.76 | 0.06 | 0.06 | <0.01 | <0.01 |
| <0.01 | <0.01 | 0.06 | <0.01 | <0.01 | <0.01 |

T = 2

| | | | | | |
|---|---|---|---|---|---|
| 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 | 0.03 | 0.05 | <0.01 | 0.03 | <0.01 |
| 0.03 | 0.03 | <0.01 | <0.01 | <0.01 | <0.01 |

T = 5

$$B'(X_{t+1}) = \sum_{X_t} P(X_{t+1}|X_t, e_{1:t})B(X_t)$$

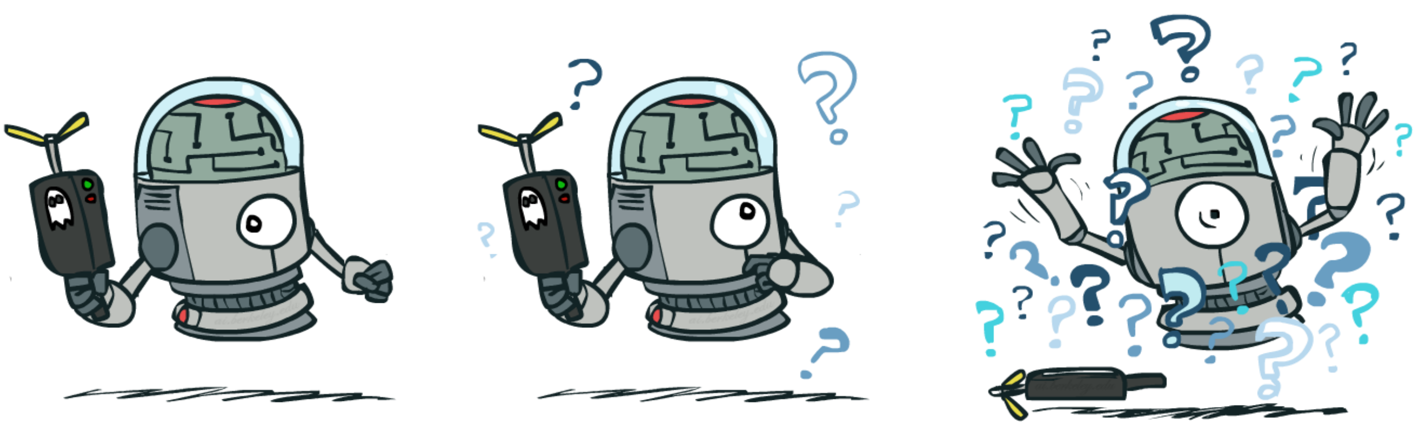

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Observation update

$B'(X_t)$ $\qquad$ $B(X_{t+1})$

$P(X_{t+1}|e_{1:t}) \rightarrow P(X_{t+1}|e_{1:t+1})$

$$P(X_{t+1}|e_{1:t+1}) = \eta P(e_{t+1}|X_{t+1}) P(X_{t+1}|e_{1:t})$$

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

Where $\eta = \dfrac{1}{P(e_{t+1})}$ is a normalization factor

# Observation update

- 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 | 0.03 | 0.05 | <0.01 | 0.03 | <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

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

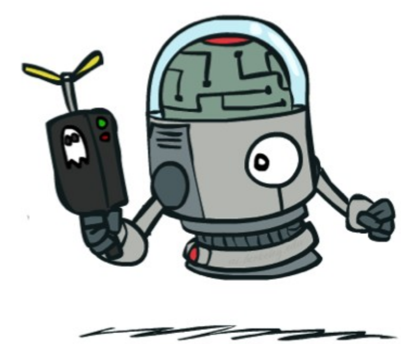

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Robot localization example

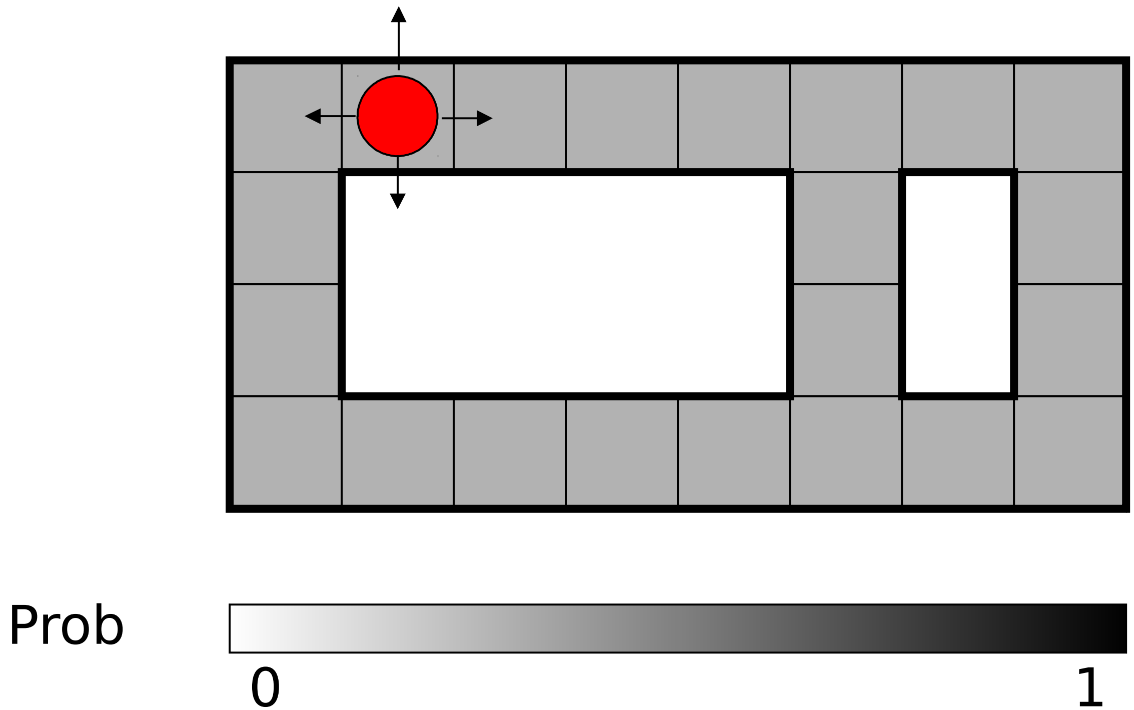


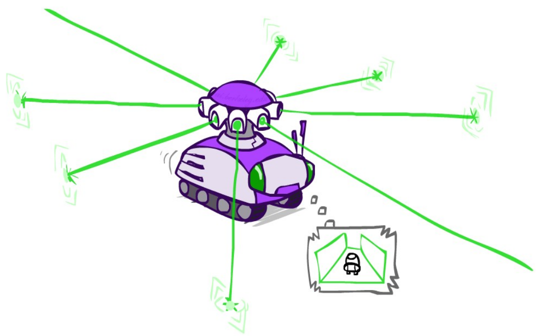

Observation model: can read in which directions there is a wall, never more than 1 mistake

Process model: may not execute action with small prob.

# Robot localization example

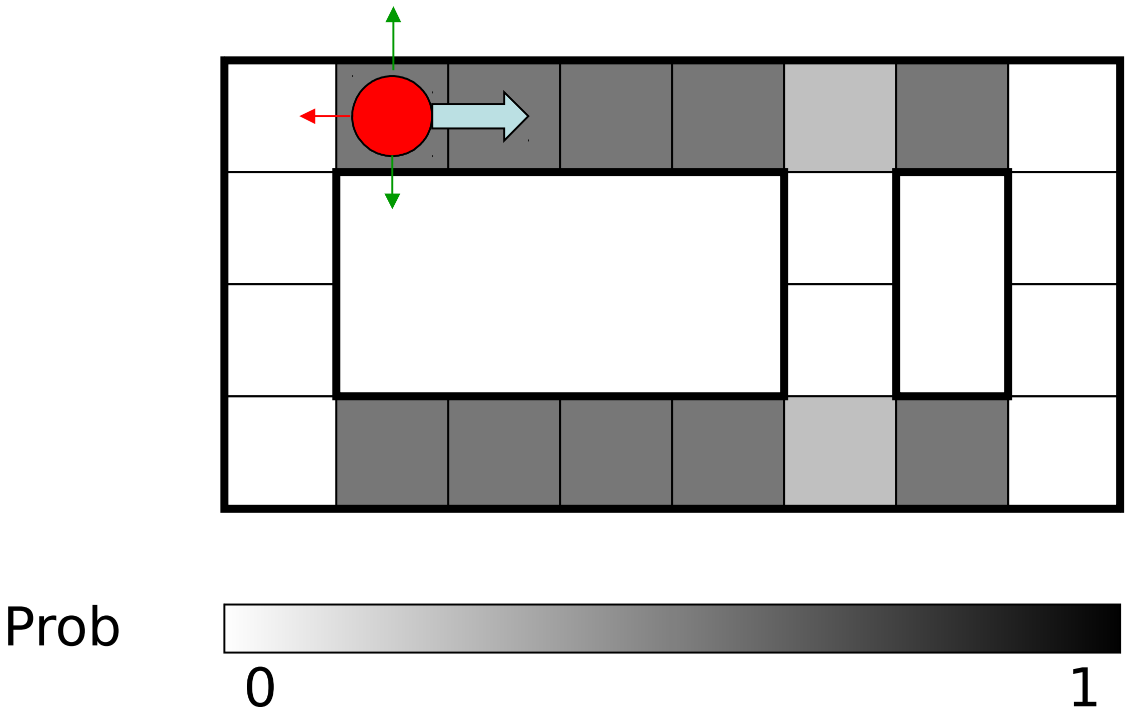


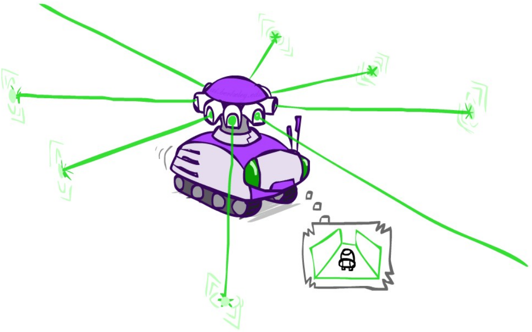

Lighter grey: was possible to get the reading, but less likely b/c required 1 mistake

# Robot localization example

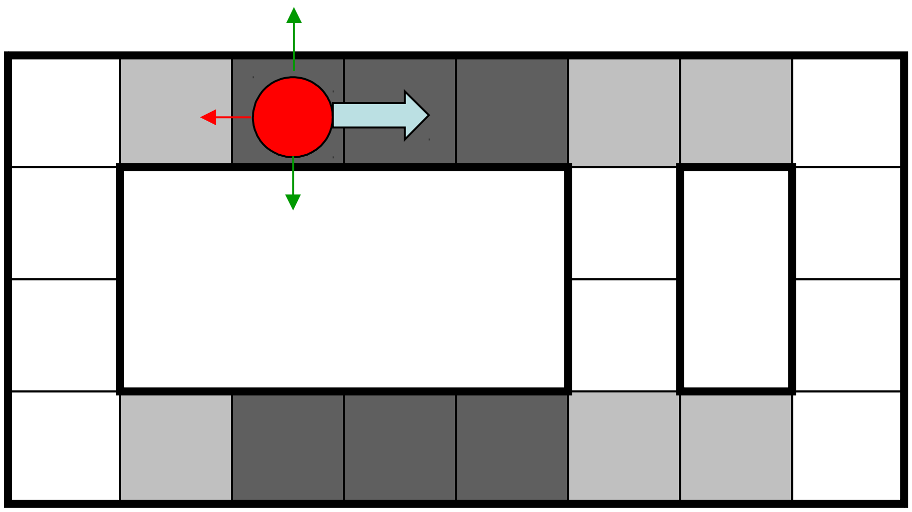

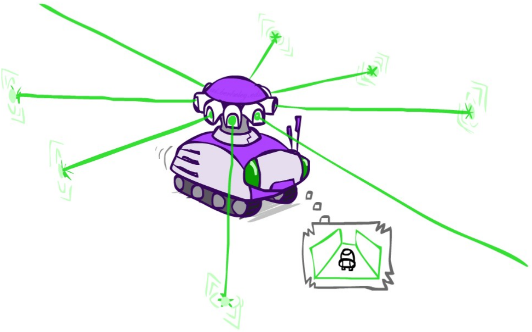


Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Robot localization example

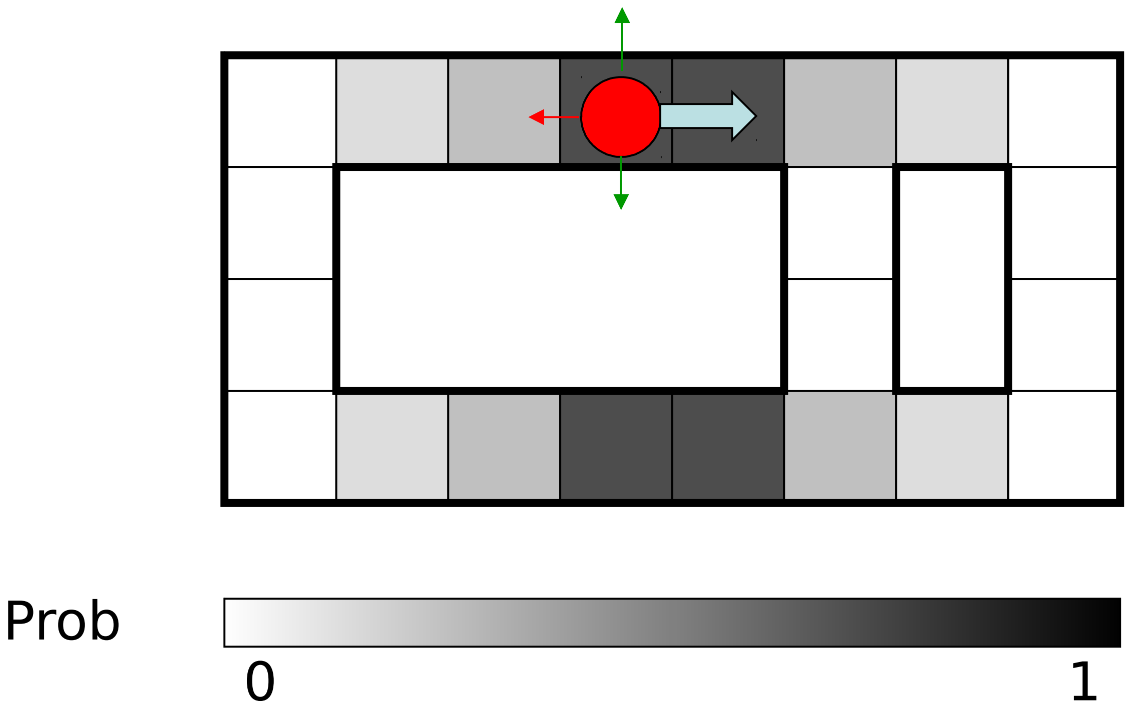


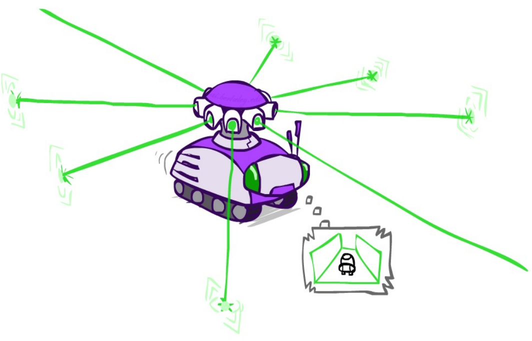

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Robot localization example

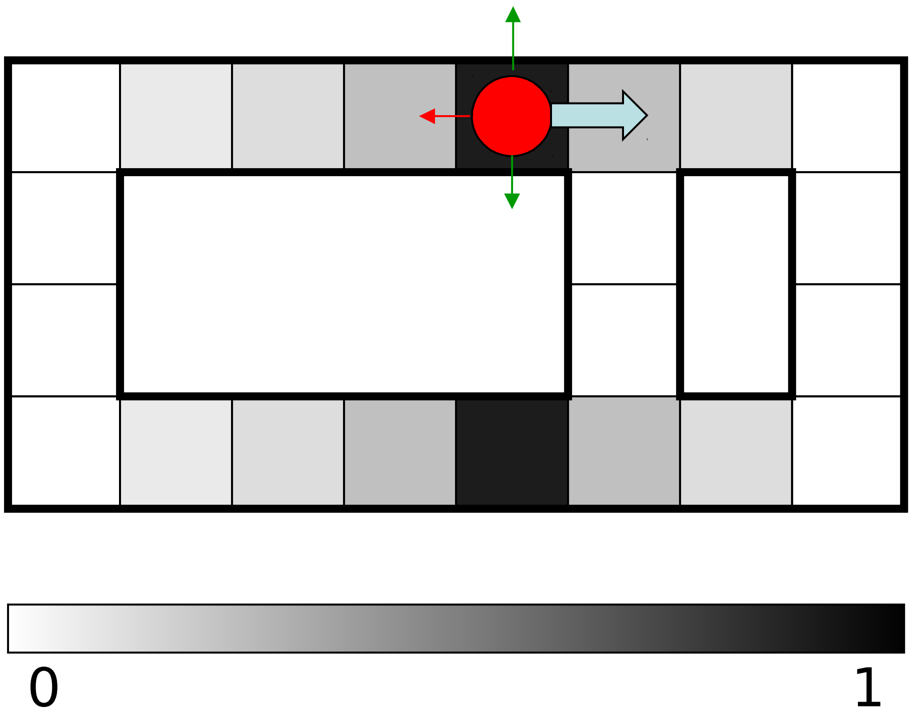


Prob

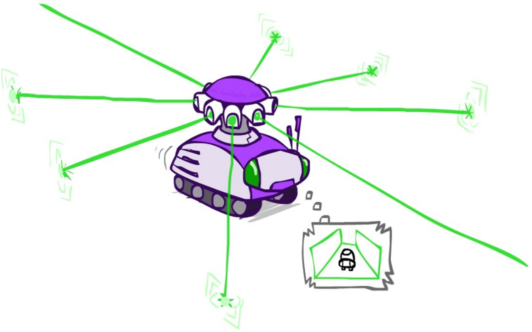

# Robot localization example

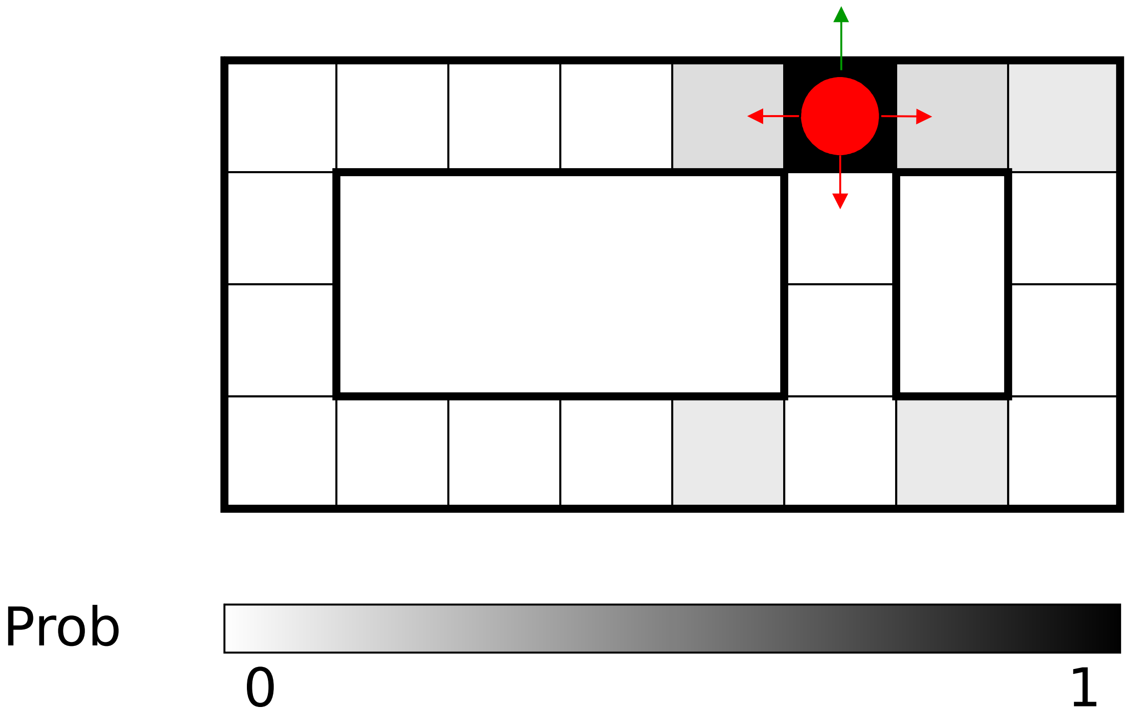


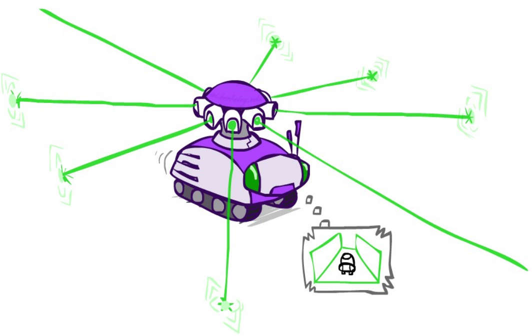

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Weather HMM example

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

B(+r) = 0.5
B(-r) = 0.5

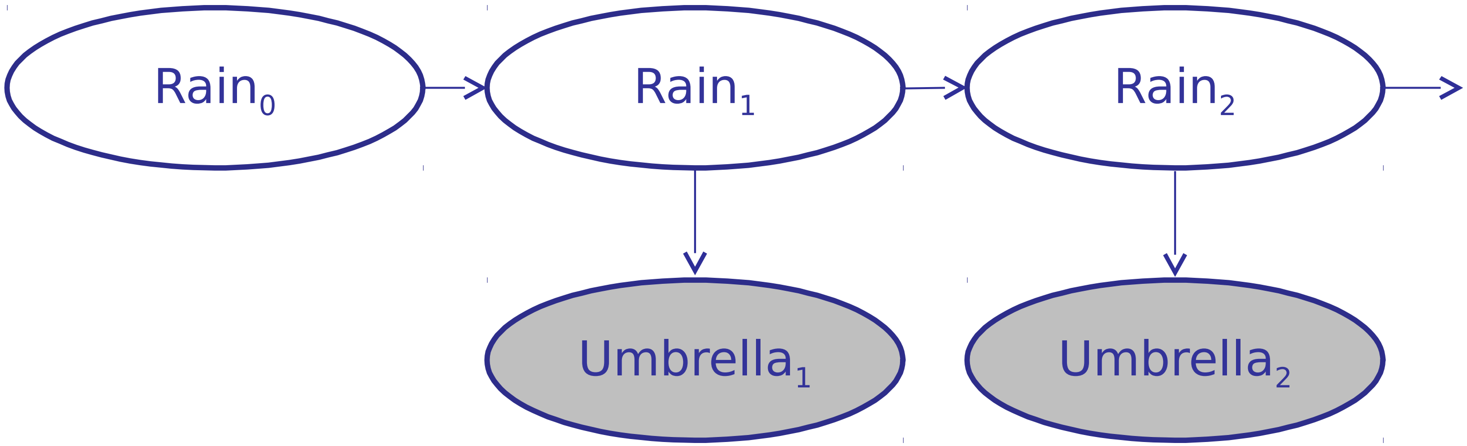


| $R_t$ | $U_t$ | $P(U_t\|R_t)$ |
|---|---|---|
| +r | +u | 0.9 |
| +r | -u | 0.1 |
| -r | +u | 0.2 |
| -r | -u | 0.8 |

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Weather HMM example

B'(+r) = 0.5
B'(-r) = 0.5

B(+r) = 0.5
B(-r) = 0.5

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

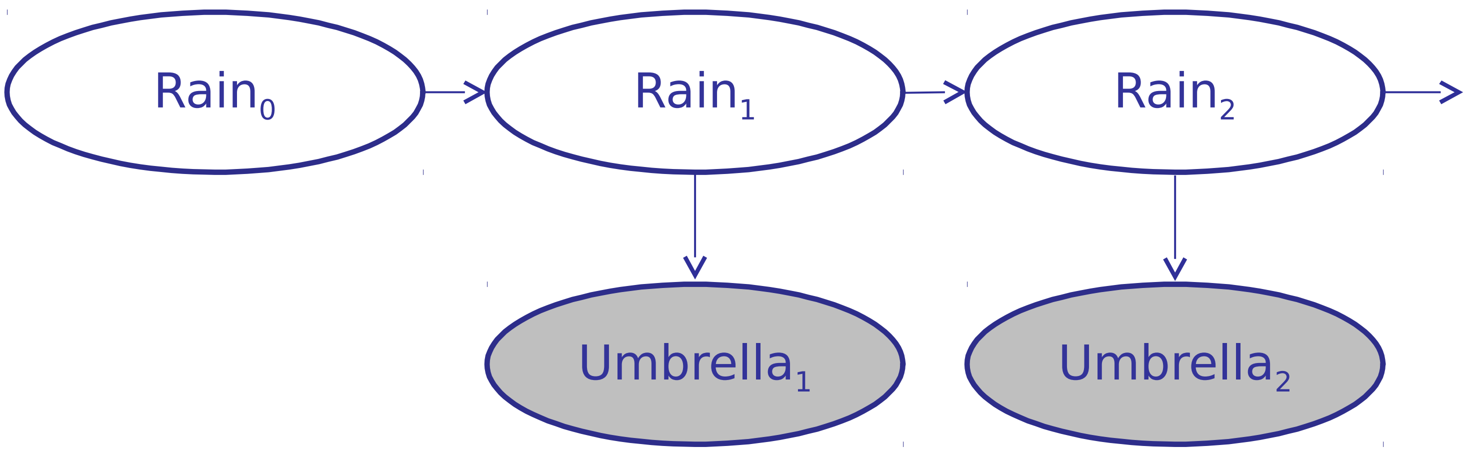


| $R_t$ | $U_t$ | $P(U_t\|R_t)$ |
|---|---|---|
| +r | +u | 0.9 |
| +r | -u | 0.1 |
| -r | +u | 0.2 |
| -r | -u | 0.8 |

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Weather HMM example

B'(+r) = 0.5
B'(-r) = 0.5

B(+r) = 0.5
B(-r) = 0.5

B(+r) = 0.818
B(-r) = 0.182

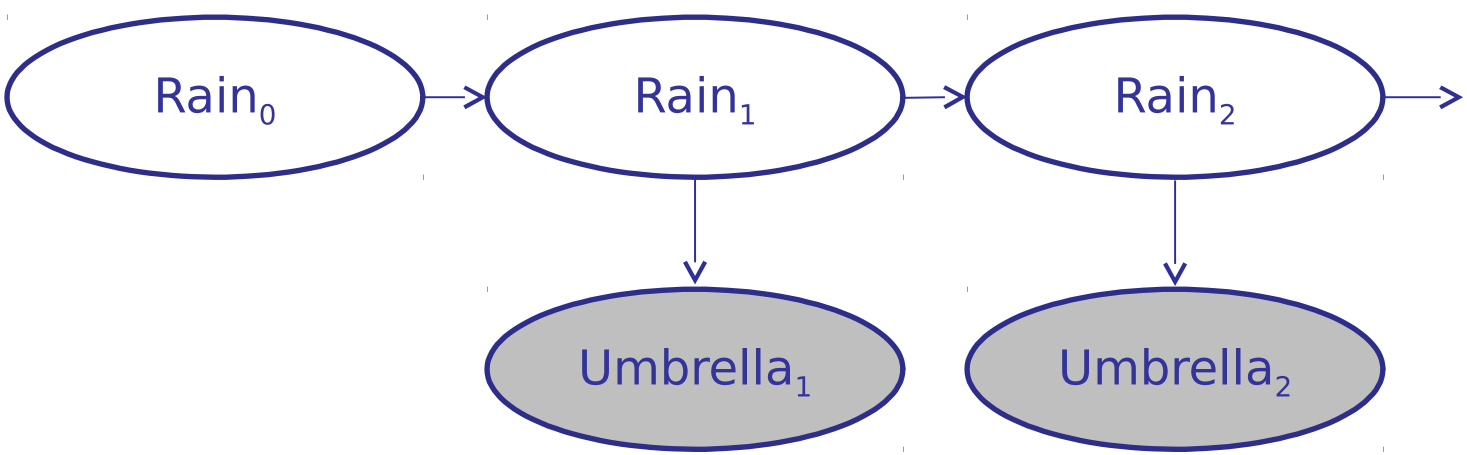


| $R_t$ | $R_{t+1}$ | $P(R_{t+1} \mid R_t)$ |
|---|---|---|
| +r | +r | 0.7 |
| +r | -r | 0.3 |
| -r | +r | 0.3 |
| -r | -r | 0.7 |

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

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Weather HMM example

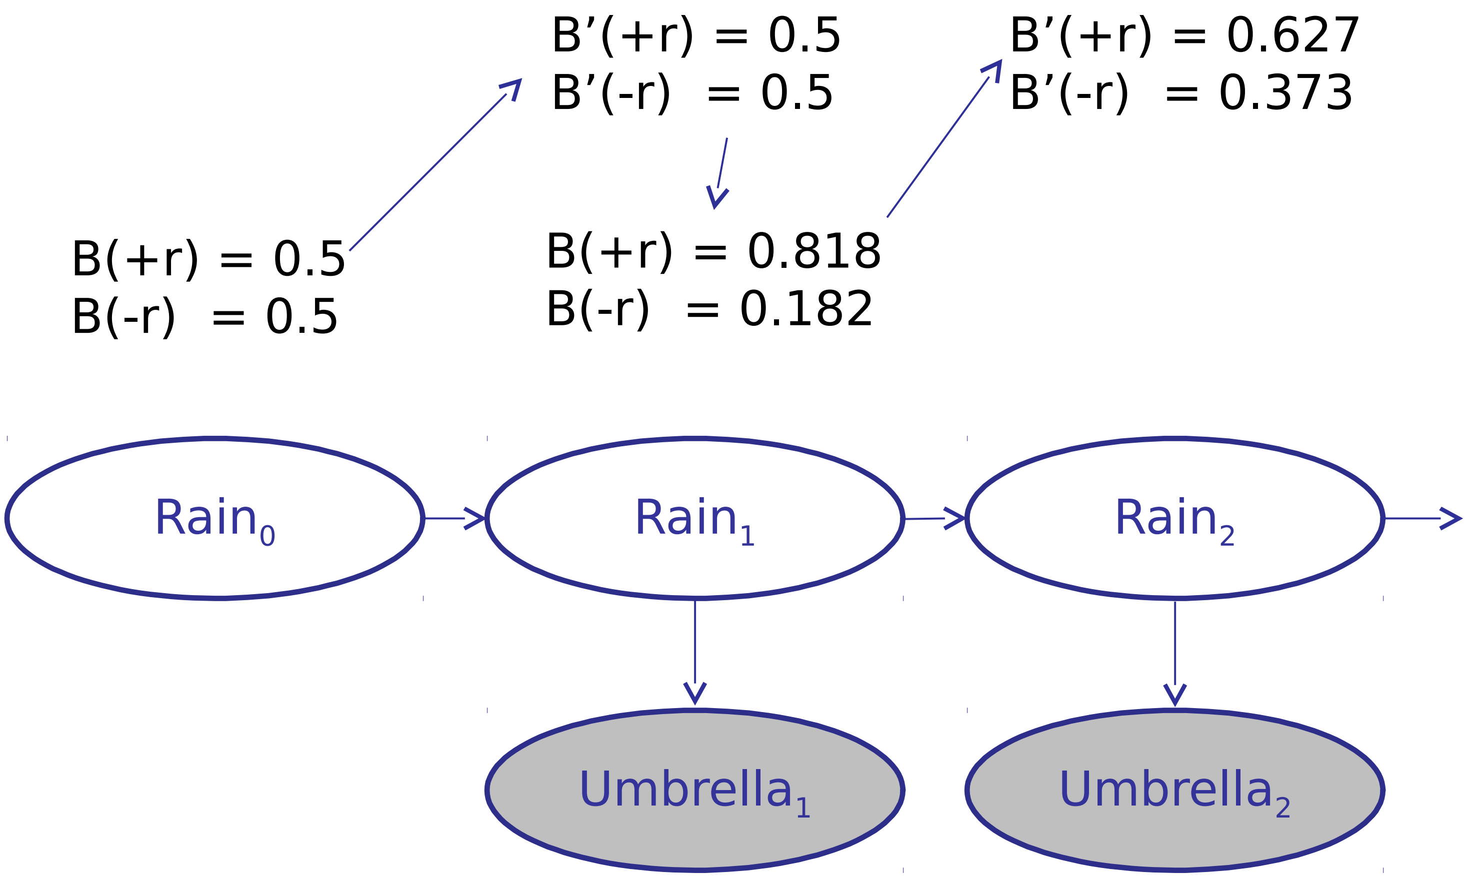


| $R_t$ | $R_{t+1}$ | $P(R_{t+1}\|R_t)$ |
|---|---|---|
| +r | +r | 0.7 |
| +r | -r | 0.3 |
| -r | +r | 0.3 |
| -r | -r | 0.7 |

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

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Weather HMM example

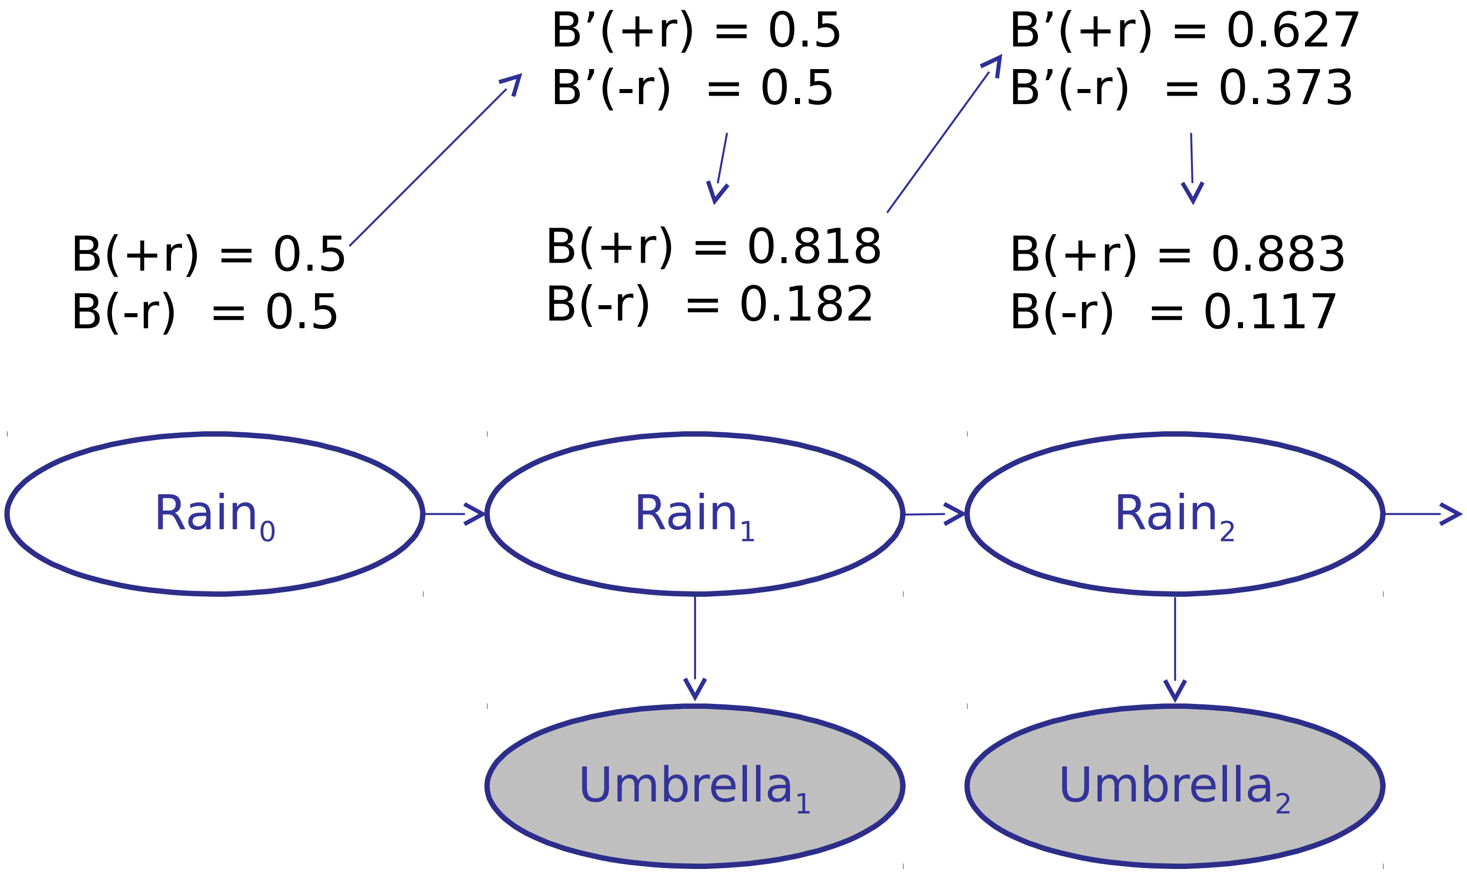


| $R_t$ | $R_{t+1}$ | $P(R_{t+1} \mid R_t)$ |
|---|---|---|
| +r | +r | 0.7 |
| +r | -r | 0.3 |
| -r | +r | 0.3 |
| -r | -r | 0.7 |

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

# Particle Filtering

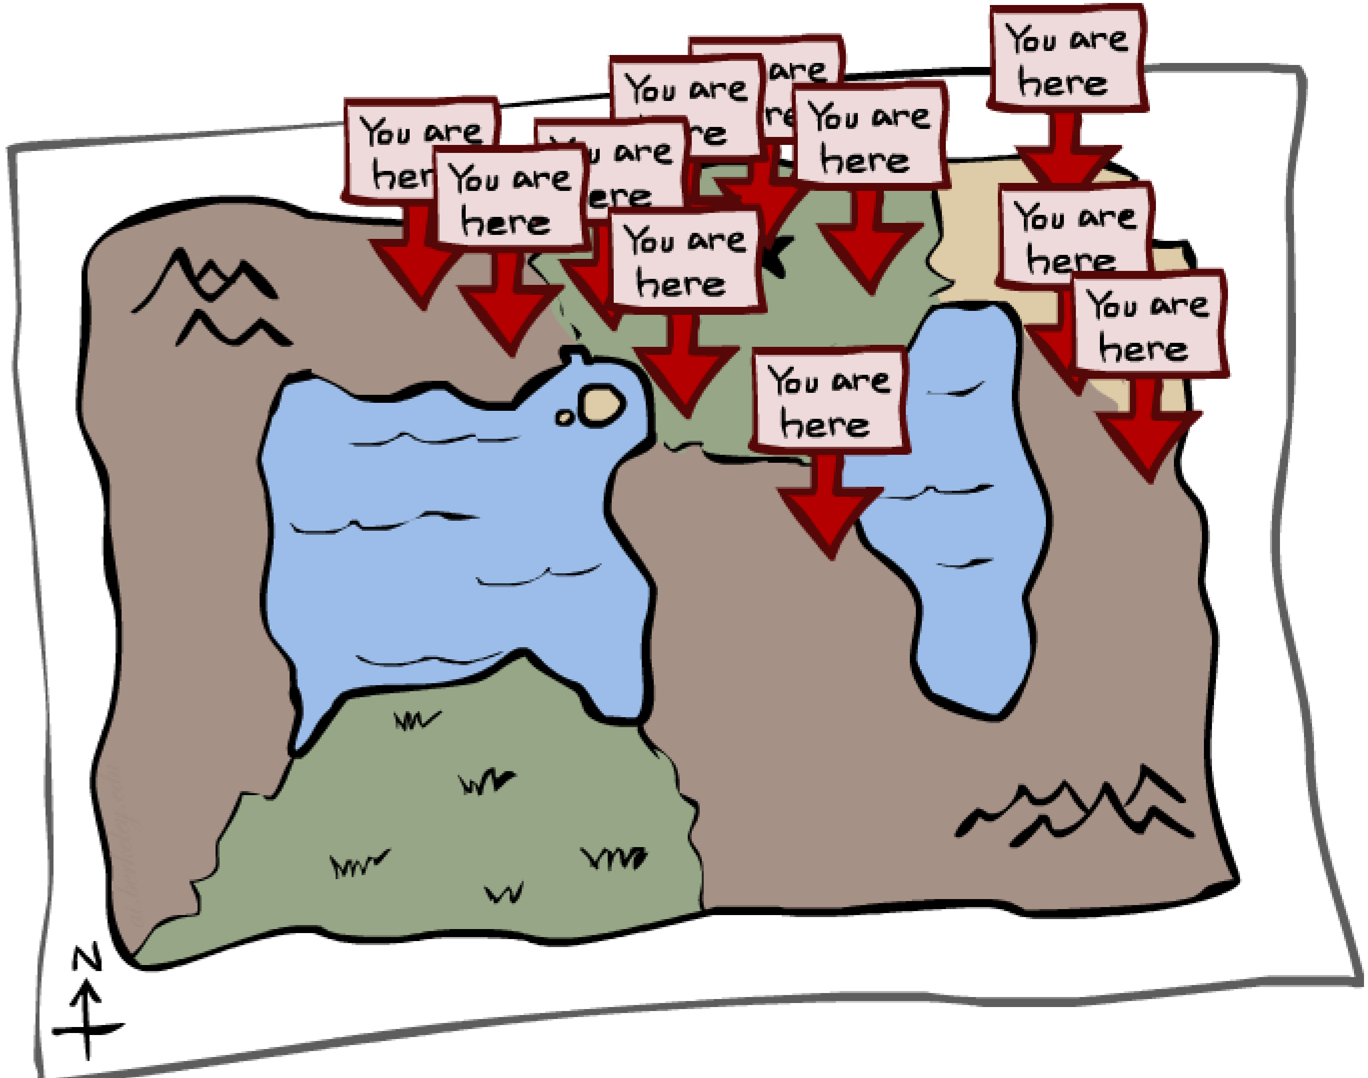


Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Representation: Particles

- Our representation of P(X) is now a list of N particles (samples)
  - Generally, N << |X|
  - Storing map from X to counts would defeat the point

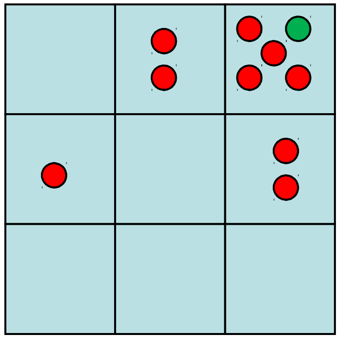

Particles:
(3,3)
(2,3)
(3,3)
(3,2)
(3,3)
(3,2)
(1,2)
(3,3)
(3,3)
(2,3)

- 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: Elapse Time

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

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

  - This is like prior sampling – samples' 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)

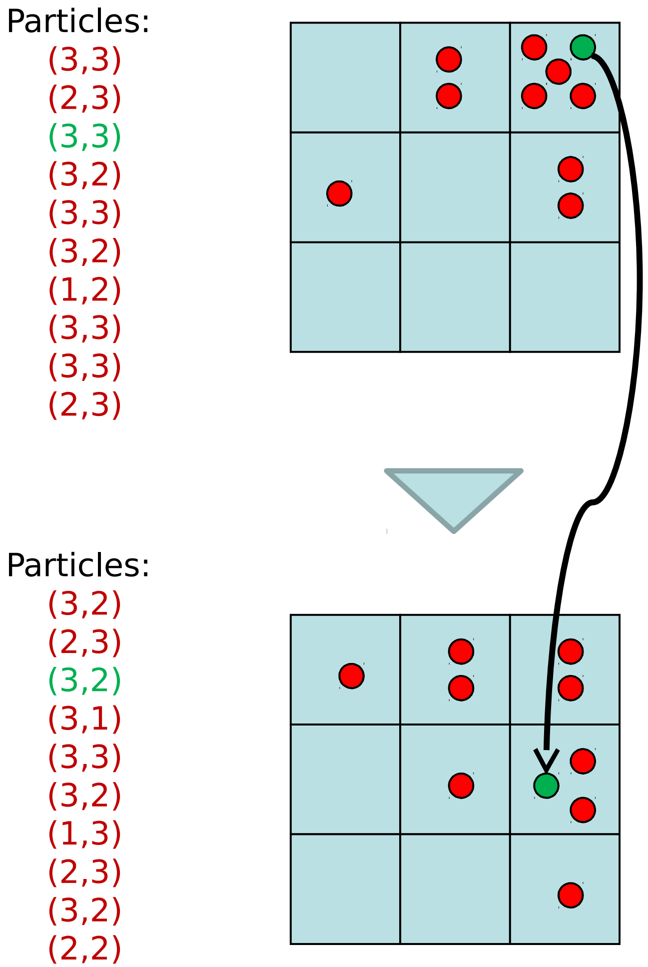


Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# 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 (in fact they now sum to (N times) an approximation of P(e))

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

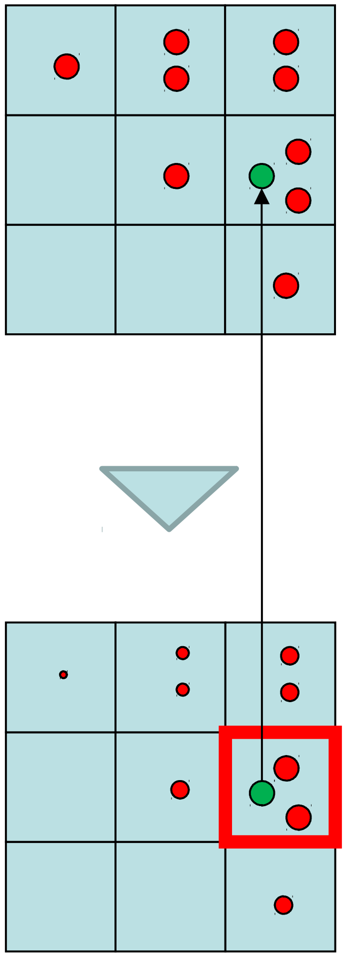

# Particle Filtering: Resample

- Rather than tracking weighted samples, we resample
- N times, we choose from our weighted sample distribution (i.e. draw with replacement)
- This is equivalent to renormalizing 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)

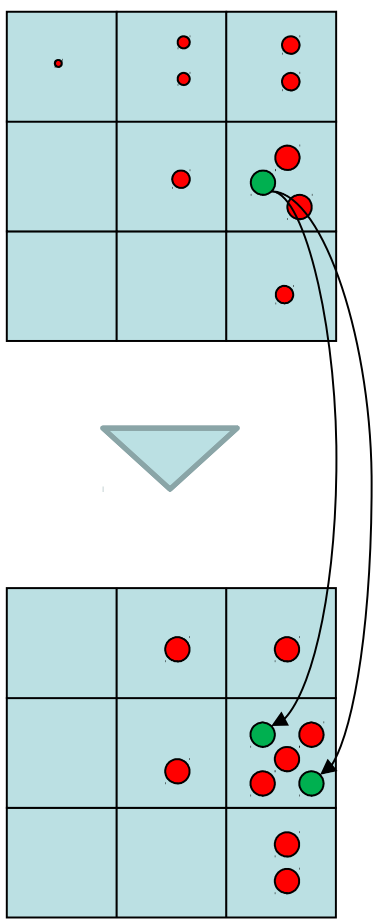

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Recap: Particle Filtering

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

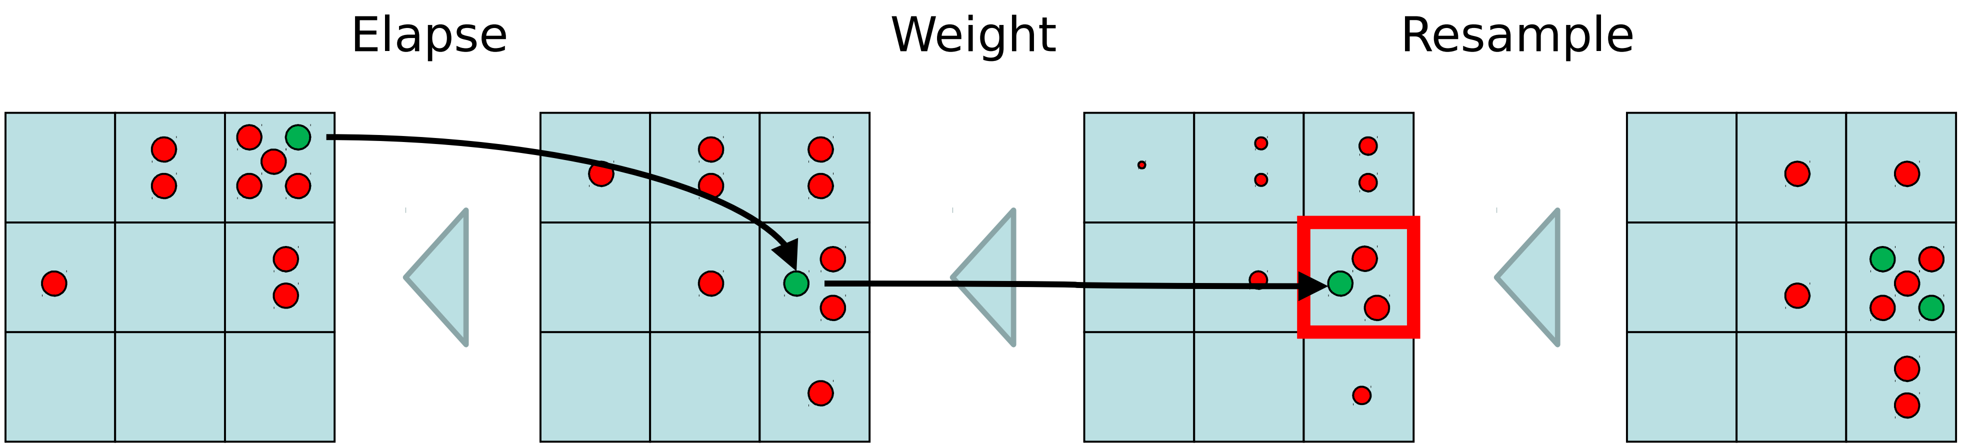


Particles:
(3,3)
(2,3)
(3,3)
(3,2)
(3,3)
(3,2)
(1,2)
(3,3)
(3,3)
(2,3)

Particles:
(3,2)
(2,3)
(3,2)
(3,1)
(3,3)
(3,2)
(1,3)
(2,3)
(3,2)
(2,2)

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)

Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# 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

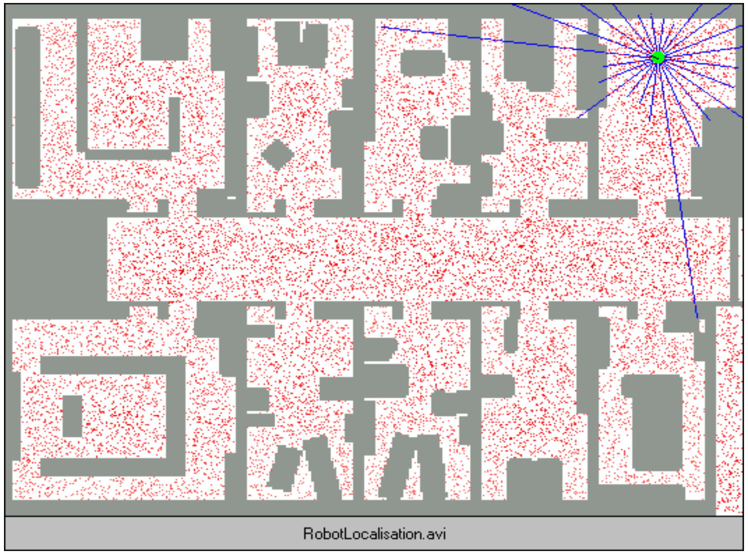


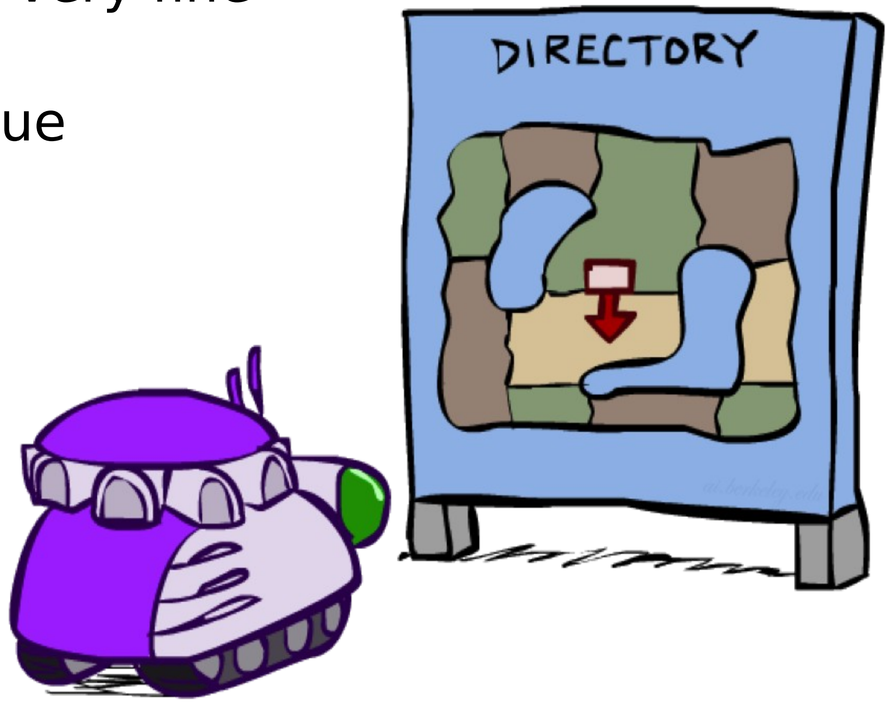


Slide: Berkeley CS188 course notes (downloaded Summer 2015)

# Particle Filter Localization (Sonar)

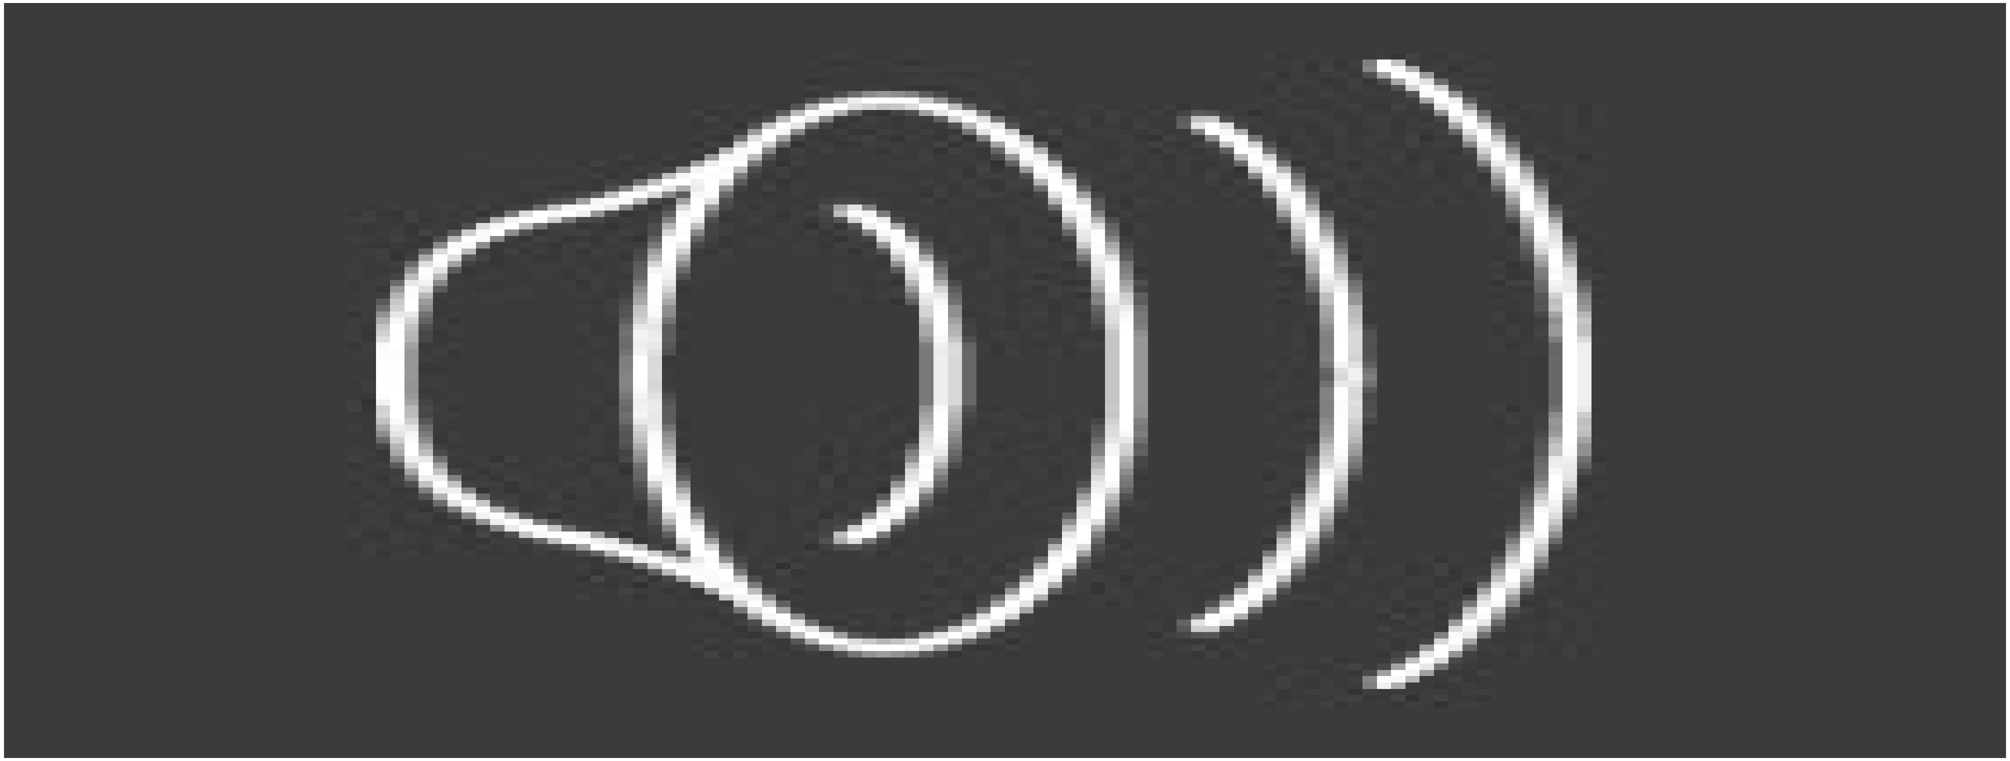

# Particle Filter Localization (Laser)

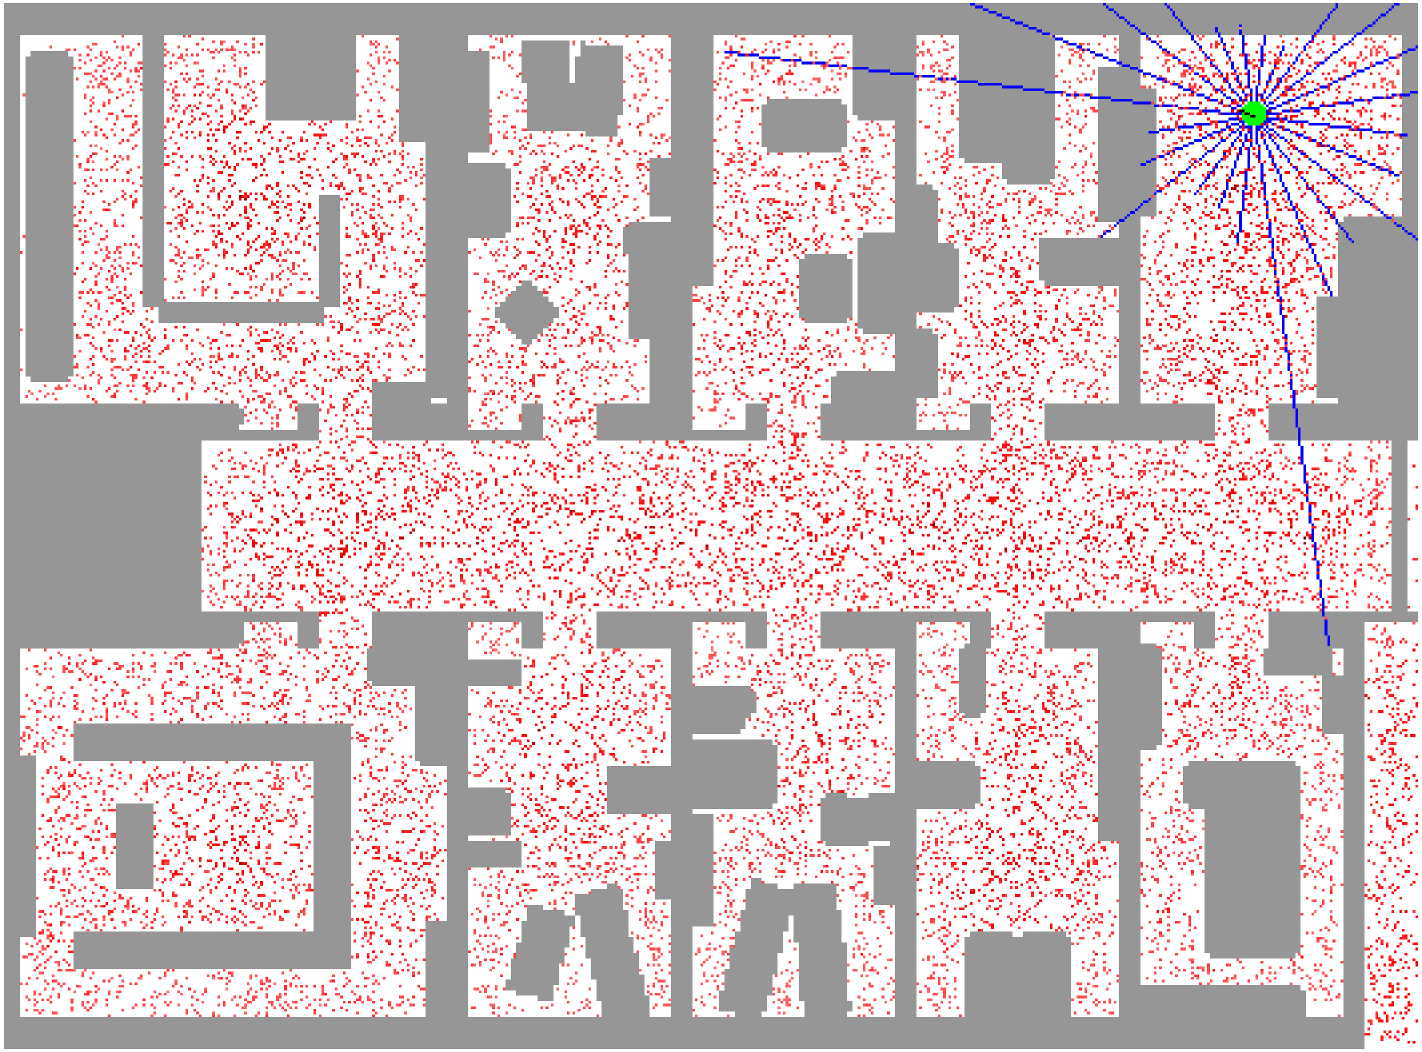

# Dynamic Bayes Nets

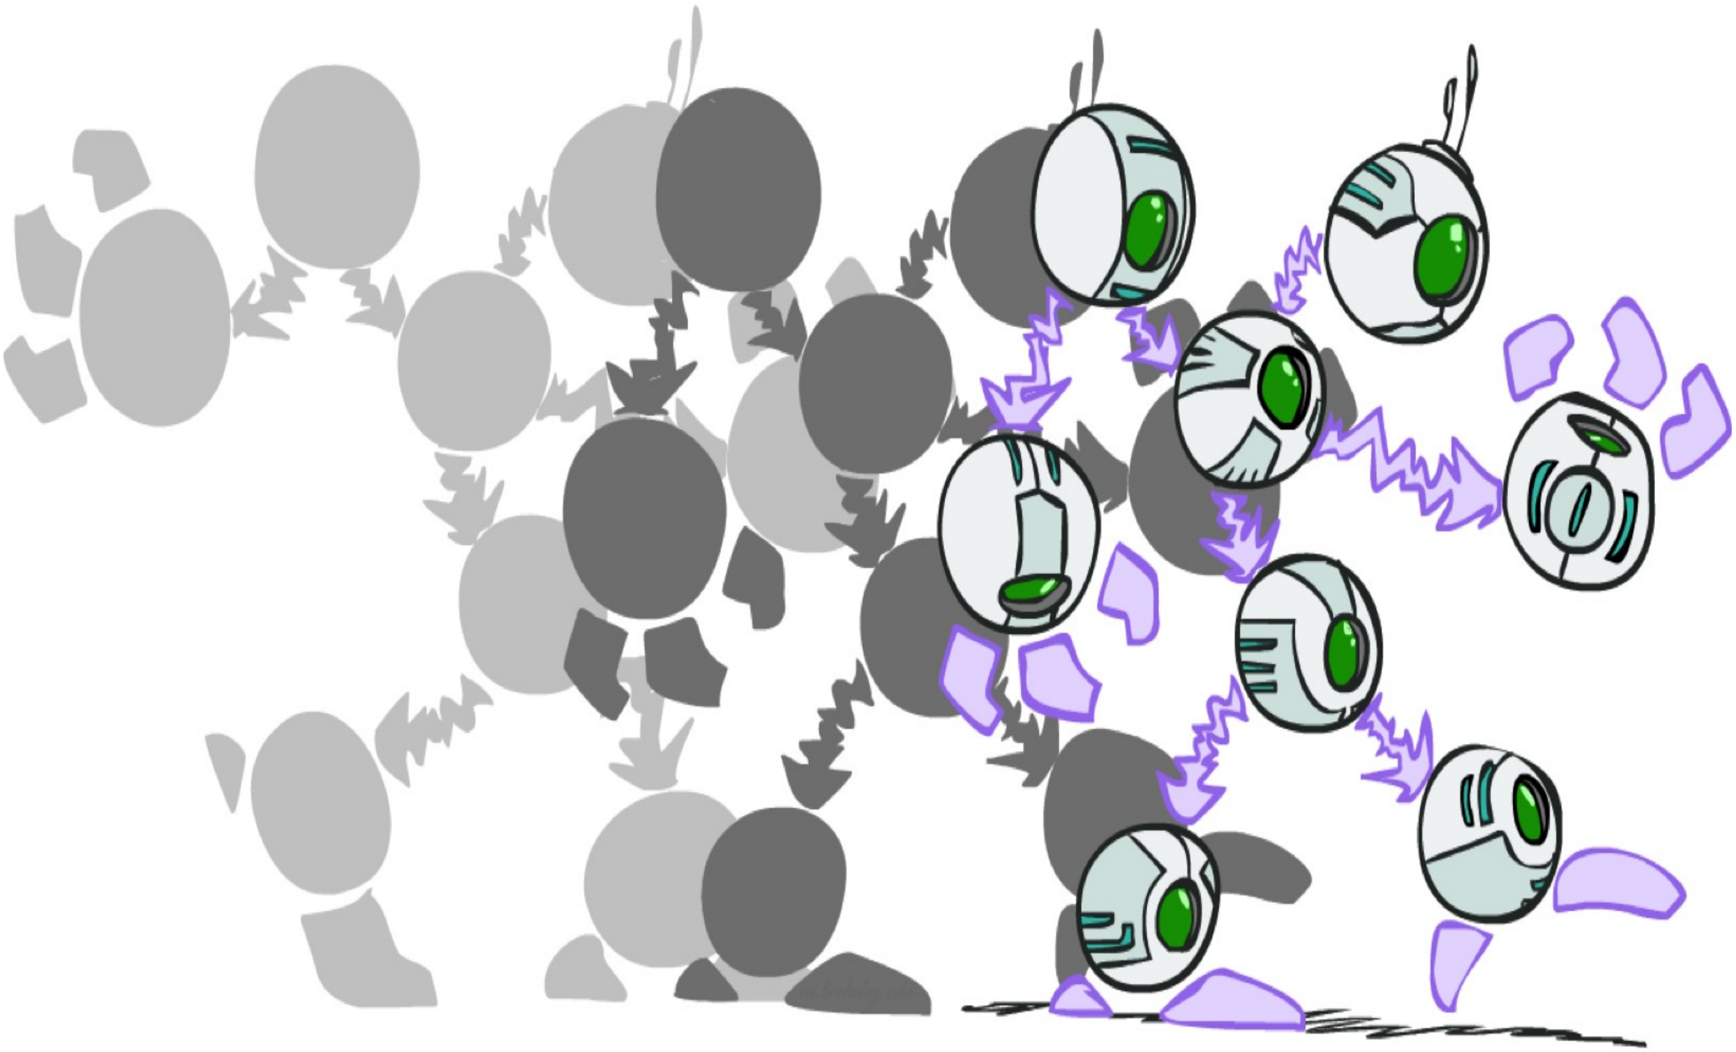

# Dynamic Bayes Nets (DBNs)

- 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*

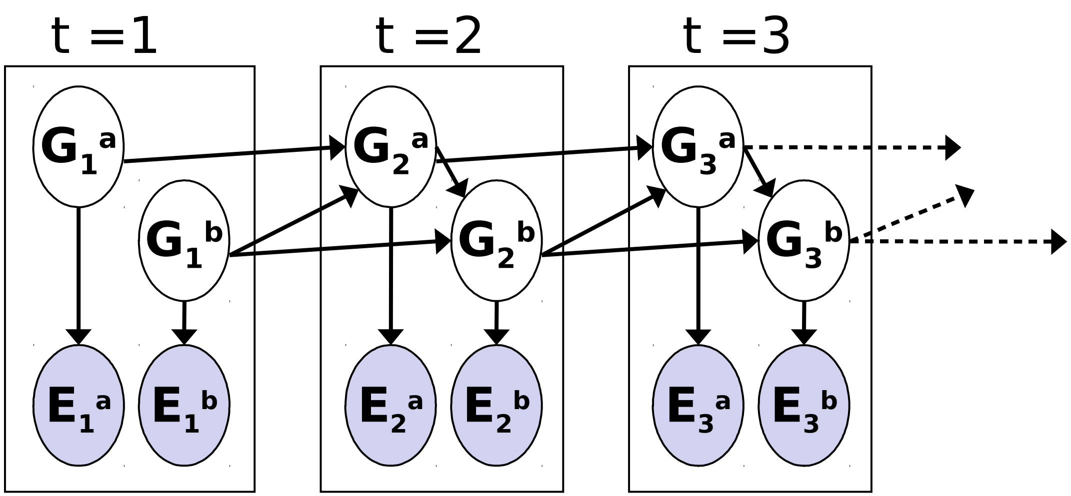


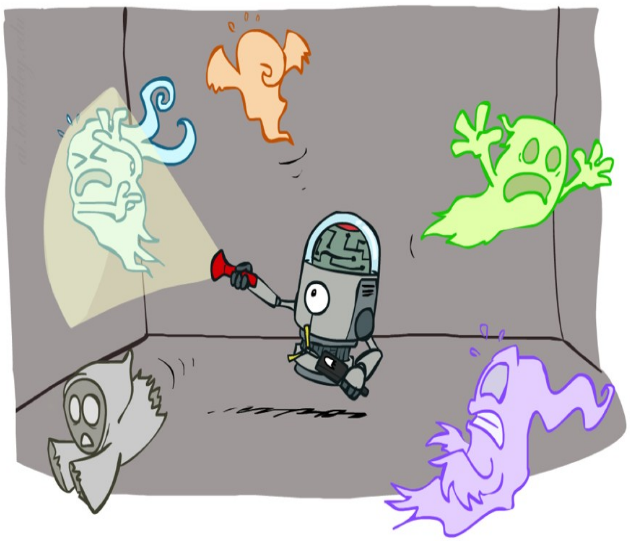

- Dynamic Bayes nets are a generalization of HMMs

# 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: $P(\mathbf{E_1^a} \mid \mathbf{G_1^a}) * P(\mathbf{E_1^b} \mid \mathbf{G_1^b})$
- **Resample:** Select prior samples (tuples of values) in proportion to their likelihood