# CSE 573: Artificial Intelligence Autumn 2010 

## Lecture 10: Hidden Markov Models 11/2/2010

## Luke Zettlemoyer

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

## Outline

- Probability review
- Random Variables and Events
- Joint / Marginal / Conditional Distributions
- Product Rule, Chain Rule, Bayes’ Rule
- Probabilistic Inference
- Probabilistic sequence models (and inference)
- Markov Chains
- Hidden Markov Models
- Particle Filters


## Announcements

- PS2 grades are out
- PS3 is due Friday -MDPs and RL


## Probability Review

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


## Inference in Ghostbusters

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

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

## Joint Distributions

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

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

$P(T, W)$

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

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

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

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

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


## Events

- An outcome is a joint assignment for all the variables

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

- An event is a set E of outcomes

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

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

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

- Probability that it's hot AND sunny?
- Probability that it's hot?
- Probability that it's hot OR sunny?


## Marginal Distributions

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

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

$P(T)$

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

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

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

$P(W)$

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

## Conditional Distributions

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

Conditional Distributions Joint Distribution

$P(T, W)$

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

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

## Probabilistic Inference

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


## Inference by Enumeration

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

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

## Inference by Enumeration

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

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

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


## The Product Rule

- Sometimes have conditional distributions but want the joint

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

- Example:

$$
P(D \mid W)
$$

$P(D, W)$

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

## The Chain Rule

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

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

- Why is this always true?


## Bayes' Rule

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

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

- Dividing, we get:

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

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


## Ghostbusters, Revisited

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

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

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


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

## Markov Models (Markov Chains)

- A Markov model is:
- a MDP with no actions (and no rewards)
- a chain-structured Bayesian Network (BN)

- A Markov model includes:
- Random variables $X_{t}$ for all time steps $t$ (the state)
- Parameters: called transition probabilities or dynamics, specify how the state evolves over time (also, initial probs)

$$
P\left(X_{1}\right) \quad \text { and } \quad P\left(X_{t} \mid X_{t-1}\right)
$$

## Markov Models (Markov Chains)



- A Markov model defines:
- a joint probability distribution

$$
P\left(X_{1}, \ldots, X_{n}\right)=P\left(X_{1}\right) \prod_{t=2}^{N} P\left(X_{t} \mid X_{t-1}\right)
$$

- One common inference problem:
- Compute marginals $P\left(X_{t}\right)$ for all time steps $t$


## Example: Markov Chain

- Weather:
- States: X = \{rain, sun\}
- Transitions:

- Initial distribution: 1.0 sun
- What's the probability distribution after one step?

$$
\begin{aligned}
P\left(X_{2}=\text { sun }\right)=\quad & P\left(X_{2}=\operatorname{sun} \mid X_{1}=\text { sun }\right) P\left(X_{1}=\text { sun }\right)+ \\
& P\left(X_{2}=\operatorname{sun} \mid X_{1}=\text { rain }\right) P\left(X_{1}=\text { rain }\right) \\
& 0.9 \cdot 1.0+0.1 \cdot 0.0=0.9
\end{aligned}
$$

## Markov Chain Inference

- Question: probability of being in state $x$ at time t?
- Slow answer:
- Enumerate all sequences of length $t$ which end in $s$
- Add up their probabilities

$$
\begin{gathered}
P\left(X_{t}=\operatorname{sun}\right)=\sum_{x_{1} \ldots x_{t-1}} P\left(x_{1}, \ldots x_{t-1}, \text { sun }\right) \\
P\left(X_{1}=\operatorname{sun}\right) P\left(X_{2}=\operatorname{sun} \mid X_{1}=\operatorname{sun}\right) P\left(X_{3}=\operatorname{sun} \mid X_{2}=\operatorname{sun}\right) P\left(X_{4}=\operatorname{sun} \mid X_{3}=\operatorname{sun}\right) \\
P\left(X_{1}=\operatorname{sun}\right) P\left(X_{2}=\operatorname{sain} \mid X_{1}=\operatorname{sun}\right) P\left(X_{3}=\operatorname{sun} \mid X_{2}=\operatorname{rain}\right) P\left(X_{4}=\operatorname{sun} \mid X_{3}=\operatorname{sun}\right)
\end{gathered}
$$

## Mini-Forward Algorithm

- Question: What's $P(X)$ on some day t?
- We don't need to enumerate every sequence!


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

$$
P\left(x_{1}\right)=\text { known }
$$



Forward simulation

## Example

- From initial observation of sun

$$
\begin{gathered}
\left\langle\begin{array}{l}
1.0 \\
0.0
\end{array}\right\rangle
\end{gathered}\left\langle\begin{array}{l}
0.9 \\
0.1
\end{array}\right\rangle\left\langle\begin{array}{l}
0.82 \\
0.18
\end{array}\right\rangle \rightleftarrows \begin{array}{|c}
\left\langle\begin{array}{l}
0.5 \\
0.5
\end{array}\right\rangle \\
\mathrm{P}\left(X_{1}\right)
\end{array} \begin{aligned}
& \mathrm{P}\left(X_{2}\right)
\end{aligned} \mathrm{P}\left(X_{3}\right) \quad \begin{aligned}
& \mathrm{P}\left(X_{\infty}\right)
\end{aligned}
$$

- From initial observation of rain

$$
\left. \begin{array}{l}
0.5 \\
0.5
\end{array}\right\rangle
$$

## Stationary Distributions

- If we simulate the chain long enough:
- What happens?
- Uncertainty accumulates
- Eventually, we have no idea what the state is!
- Stationary distributions:
- For most chains, the distribution we end up in is independent of the initial distribution
- Called the stationary distribution of the chain
- Usually, can only predict a short time out


## Pac-man Markov Chain

Pac-man knows the ghost's initial position, but gets no observations!


## 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)


## Hidden Markov Models

- Markov chains not so useful for most agents
- Eventually you don't know anything anymore
- Need observations to update your beliefs
- Hidden Markov models (HMMs)
- Underlying Markov chain over states S
- You observe outputs (effects) at each time step
- POMDPs without actions (or rewards).
- As a Bayes' net:



## Example



- An HMM is defined by:
- Initial distribution: $P\left(X_{1}\right)$
- Transitions:
$P\left(X_{t} \mid X_{t-1}\right)$
- Emissions:
$P(E \mid X)$


## Hidden Markov Models



- Defines a joint probability distribution:

$$
P\left(X_{1}, \ldots, X_{n}, E_{1}, \ldots, E_{n}\right)=
$$

$$
P\left(X_{1: n}, E_{1: n}\right)=
$$

$$
P\left(X_{1}\right) P\left(E_{1} \mid X_{1}\right) \prod_{t=2}^{N} P\left(X_{t} \mid X_{t-1}\right) P\left(E_{t} \mid X_{t}\right)
$$

## Ghostbusters HMM

- $P\left(X_{1}\right)=$ uniform
- $P\left(X^{\prime} \mid X\right)=$ usually move clockwise, but sometimes move in a random direction or stay in place
- $P(E \mid X)=$ same sensor model as before: red means close, green means far away.

| $1 / 9$ | $1 / 9$ | $1 / 9$ |
| :--- | :--- | :--- |
| $1 / 9$ | $1 / 9$ | $1 / 9$ |
| $1 / 9$ | $1 / 9$ | $1 / 9$ |
| $P\left(X_{1}\right)$ |  |  |



| $1 / 6$ | $1 / 6$ | $1 / 2$ |
| :---: | :---: | :---: |
| 0 | $1 / 6$ | 0 |
| 0 | 0 | 0 |

$P\left(X^{\prime} \mid X=<1,2>\right)$

## HMM Computations

- Given
- joint $P\left(X_{1: n}, E_{1: n}\right)$
- evidence $E_{1: n}=e_{1: n}$
- Inference problems include:
- Filtering, find $P\left(X_{t} \mid e_{1: t}\right)$ for all $t$
- Smoothing, find $P\left(X_{t} \mid e_{1: n}\right)$ for all $t$
- Most probable explanation, find

$$
x_{1: n}^{*}=\operatorname{argmax}_{x_{1: n}} P\left(x_{1: n} \mid e_{1: n}\right)
$$

## 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)


## Filtering / Monitoring

- Filtering, or monitoring, is the task of tracking the distribution $B(X)$ (the belief state) over time
- We start with $B(X)$ 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



Prob


$$
t=0
$$

Sensor model: never more than 1 mistake Motion model: may not execute action with small prob.

## Example: Robot Localization



## Example: Robot Localization



## Example: Robot Localization



## Example: Robot Localization



## Example: Robot Localization



## Inference Recap: Simple Cases



$$
P\left(X_{1} \mid e_{1}\right)
$$

$$
\begin{aligned}
P\left(x_{1} \mid e_{1}\right) & =P\left(x_{1}, e_{1}\right) / P\left(e_{1}\right) \\
& \propto_{X_{1}} P\left(x_{1}, e_{1}\right) \\
& =P\left(x_{1}\right) P\left(e_{1} \mid x_{1}\right)
\end{aligned}
$$


$P\left(X_{2}\right)$

$$
\begin{aligned}
P\left(x_{2}\right) & =\sum_{x_{1}} P\left(x_{1}, x_{2}\right) \\
& =\sum_{x_{1}} P\left(x_{1}\right) P\left(x_{2} \mid x_{1}\right)
\end{aligned}
$$

## Online Belief Updates

- Every time step, we start with current $\mathrm{P}(\mathrm{X} \mid$ evidence $)$
- We update for time:

$$
P\left(x_{t} \mid e_{1: t-1}\right)=\sum_{x_{t-1}} P\left(x_{t-1} \mid e_{1: t-1}\right) \cdot P\left(x_{t} \mid x_{t-1}\right)
$$

- We update for evidence:

$$
P\left(x_{t} \mid e_{1: t}\right) \propto_{X} P\left(x_{t} \mid e_{1: t-1}\right) \cdot P\left(e_{t} \mid x_{t}\right)
$$



## Passage of Time

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

$$
B\left(X_{t}\right)=P\left(X_{t} \mid e_{1: t}\right)
$$

- Then, after one time step passes:

$$
P\left(X_{t+1} \mid e_{1: t}\right)=\sum_{x_{t}} P\left(X_{t+1} \mid x_{t}\right) P\left(x_{t} \mid e_{1: t}\right)
$$

- Or, compactly:

$$
B^{\prime}\left(X^{\prime}\right)=\sum_{x} P\left(X^{\prime} \mid x\right) B(x)
$$

- 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


## Example: Passage of Time

- As time passes, uncertainty "accumulates"

| $<0.01$ | $<0.01$ | $<0.01$ | $<0.01$ | $<0.01$ | $<0.01$ |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| $<0.01$ | $<0.01$ | $<0.01$ | $<0.01$ | $<0.01$ | $<0.01$ |
| $<0.01$ | $<0.01$ | 1.00 | $<0.01$ | $<0.01$ | $<0.01$ |
| $<0.01$ | $<0.01$ | $<0.01$ | $<0.01$ | $<0.01$ | $<0.01$ |
| $<$ |  |  |  |  |  |


| $<0.01$ | $<0.01$ | $<0.01$ | $<0.01$ | $<0.01$ | $<0.01$ |
| :--- | :--- | :--- | :--- | :--- | :--- |
| $<0.01$ | $<0.01$ | 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 | $<0.01$ | $<0.01$ | $<0.01$ |
| $\ll$ |  |  |  |  |  |


| 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 | $<0.01$ | $<0.01$ | $<0.01$ | $<0.01$ |
| 0. |  |  |  |  |  |

$$
\begin{array}{cc}
\mathrm{T}=1 & \mathrm{~T}=2 \\
B^{\prime}\left(X^{\prime}\right)= & \sum_{x} P\left(X^{\prime} \mid x\right) B(x)
\end{array}
$$

$$
\mathrm{T}=5
$$

Transition model: ghosts usually go clockwise

## Observation

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

$$
B^{\prime}\left(X_{t+1}\right)=P\left(X_{t+1} \mid e_{1: t}\right)
$$

- Then:

$$
P\left(X_{t+1} \mid e_{1: t+1}\right) \propto P\left(e_{t+1} \mid X_{t+1}\right) P\left(X_{t+1} \mid e_{1: t}\right)
$$

- Or:

$$
B\left(X_{t+1}\right) \propto P(e \mid X) B^{\prime}\left(X_{t+1}\right)
$$

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


## Example: 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 | 0.03 | 0.05 | $<0.01$ | 0.03 | $<0.01$ |
| 0.03 | 0.03 | $<0.01$ | $<0.01$ | $<0.01$ | $<0.01$ |
| 0 |  |  |  |  |  |

Before observation


After observation

$$
B(X) \propto P(e \mid X) B^{\prime}(X)
$$

## The Forward Algorithm

- We to know: $B_{t}(X)=P\left(X_{t} \mid e_{1: t}\right)$
- We can derive the following updates

$$
\begin{aligned}
P\left(x_{t} \mid e_{1: t}\right) & \propto{ }_{X} P\left(x_{t}, e_{1: t}\right) \\
& =\sum_{x_{t-1}} P\left(x_{t-1}, x_{t}, e_{1: t}\right) \\
& =\sum_{x_{t-1}} P\left(x_{t-1}, e_{1: t-1}\right) P\left(x_{t} \mid x_{t-1}\right) P\left(e_{t} \mid x_{t}\right) \\
& =P\left(e_{t} \mid x_{t}\right) \sum_{x_{t-1}} P\left(x_{t} \mid x_{t-1}\right) P\left(x_{t-1}, e_{1: t-1}\right)
\end{aligned}
$$

- To get $B_{t}(X)$, compute each entry and normalize


## Example HMM



## Summary: Filtering

- Filtering is the inference process of finding a distribution over $X_{T}$ given $e_{1}$ through $e_{T}: P\left(X_{T} \mid e_{1: t}\right)$
- We first compute $\mathrm{P}\left(\mathrm{X}_{1} \mid \mathrm{e}_{1}\right): \quad P\left(x_{1} \mid e_{1}\right) \propto P\left(x_{1}\right) \cdot P\left(e_{1} \mid x_{1}\right)$
- For each trom 2 to $T$, we have $P\left(X_{t-1} \mid e_{1: t-1}\right)$
- Elapse time: compute $P\left(X_{\mathrm{t}} \mid \mathrm{e}_{1: \mathrm{t}-1}\right)$

$$
P\left(x_{t} \mid e_{1: t-1}\right)=\sum_{x_{t-1}} P\left(x_{t-1} \mid e_{1: t-1}\right) \cdot P\left(x_{t} \mid x_{t-1}\right)
$$

- Observe: compute $P\left(X_{t} \mid e_{1:-1-1}, e_{t}\right)=P\left(X_{t} \mid e_{1: t}\right)$

$$
P\left(x_{t} \mid e_{1: t}\right) \propto P\left(x_{t} \mid e_{1: t-1}\right) \cdot P\left(e_{t} \mid x_{t}\right)
$$

## Recap: Reasoning Over Time

- Stationary Markov models


$$
P\left(X_{1}\right) \quad P\left(X \mid X_{-1}\right)
$$

$$
P(E \mid X)
$$

- 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: Filtering

Elapse time: compute $P\left(X_{t} \mid e_{1: t-1}\right)$

$$
P\left(x_{t} \mid e_{1: t-1}\right)=\sum_{x_{t-1}} P\left(x_{t-1} \mid e_{1: t-1}\right) \cdot P\left(x_{t} \mid x_{t-1}\right)
$$

Observe: compute $P\left(X_{t} \mid e_{1: t}\right)$

$$
P\left(x_{t} \mid e_{1: t}\right) \propto P\left(x_{t} \mid e_{1: t-1}\right) \cdot P\left(e_{t} \mid x_{t}\right)
$$


$P\left(X_{1} \mid E_{1}=\right.$ umbrella $) \quad<0.82,0.18>\quad$ Observe
$P\left(X_{2} \mid E_{1}=\right.$ umbrella $) \quad<0.63,0.37>\quad$ Elapse time
$P\left(X_{2} \mid E_{1}=u m b, E_{2}=u m b\right) \quad<0.88,0.12>\quad$ Observe

## Particle Filtering

- Sometimes $|\mathrm{X}|$ is too big to use exact inference
- $|X|$ may be too big to even store $B(X)$
- E.g. X is continuous
- $|\mathrm{X}|^{2}$ may be too big to do updates
- 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


## 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
- $P(x)$ approximated by number of particles with value $x$
- So, many $x$ will have $P(x)=0$ !
- So, many $x$ will have $P(x)=0$ !
- For now, all particles have a weight of 1


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

## Particle Filtering: Elapse Time

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

$$
x^{\prime}=\operatorname{sample}\left(P\left(X^{\prime} \mid x\right)\right)
$$

- This is like prior sampling - samples' frequencies reflect the transition probs
- Here, most samples move clockwise, but some move in another direction or stay in place
- This captures the passage of time
- If we have enough samples, close to the exact values before and after (consistent)



## Particle Filtering: Observe

- Slightly trickier:
- Don't do rejection sampling (why not?)
- We don't sample the observation, we fix it
- This is similar to likelihood weighting, so we downweight our samples based on the evidence

$$
\begin{aligned}
w(x) & =P(e \mid x) \\
B(X) & \propto P(e \mid X) B^{\prime}(X)
\end{aligned}
$$

- Note that, as before, the probabilities don't sum to one, since most have been downweighted (in fact they sum to an approximation of $\mathrm{P}(\mathrm{e})$ )



## 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

Old Particles:
$(3,3) w=0.1$
$(2,1) w=0.9$
$(2,1) w=0.9$
$(3,1) w=0.4$
$(3,2) w=0.3$
$(2,2) w=0.4$
$(1,1) w=0.4$
$(3,1) w=0.4$
$(2,1) w=0.9$
$(3,2) w=0.3$

New Particles:

$$
(2,1) w=1
$$

$(2,1) w=1$
$(2,1) w=1$
$(3,2) w=1$
$(2,2) w=1$
$(2,1) w=1$
$(1,1) w=1$
$(3,1) w=1$
$(2,1) w=1$
$(1,1) w=1$


## 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


