

## Probabilistic Reasoning over Time

- Goal: Represent and reason about changes in the world over time
- Examples:
  - WUMPUS evidence (stench, breeze, scream) arrives over time
  - Monitoring a diabetic patient
  - Inferring the current location of a robot from its sensor data


## Umbrella World

- Suppose you are a security guard robot at an underground installation. You never go outside, but you would like to know what the weather is.
- Each morning, you see the Director come in. Some mornings he has a wet umbrella; other mornings he has no umbrella.

# Notation

- State variables (is it raining on day i?):  $R_0, R_1, R_2, \dots$
- Evidence variables (is he carrying an umbrella on day i?):  $U_1, U_2, U_3, \dots$
- $X_{a:b}$  denotes  $X_a, X_{a+1}, \dots, X_{b-1}, X_b$

# Hidden Markov Model




- Markov assumption:  
 $P(R_t | R_{1:t-1}) = P(R_t | R_{t-1})$   
Captures the “dynamics” of the world. For example, rainy days and non-rainy days come in “groups”
- Sensor model:  $P(U_t | R_t)$
- Stationarity: True for all times  $t$

# Probability Distributions


$R_t$	$R_{t-1}=\text{no}$	$R_{t-1}=\text{yes}$
no	0.7	0.3
yes	0.3	0.7

$U_t$	$R_t=\text{no}$	$R_t=\text{yes}$
no	0.8	0.1
yes	0.2	0.9



We can view the HMM as a probabilistic finite state machine

# Joint Distribution



$$\underline{P}(R_{0:n}, U_{0:n}) = \underline{P}(R_0) \prod_{t=1} \underline{P}(R_t | R_{t-1}) \cdot \underline{P}(U_t | R_t)$$

Can be generalized to multiple state variables (e.g., position, velocity, and acceleration) and multiple sensors (e.g., motor speed, battery level, wheel shaft encoders)

# Temporal Reasoning Tasks

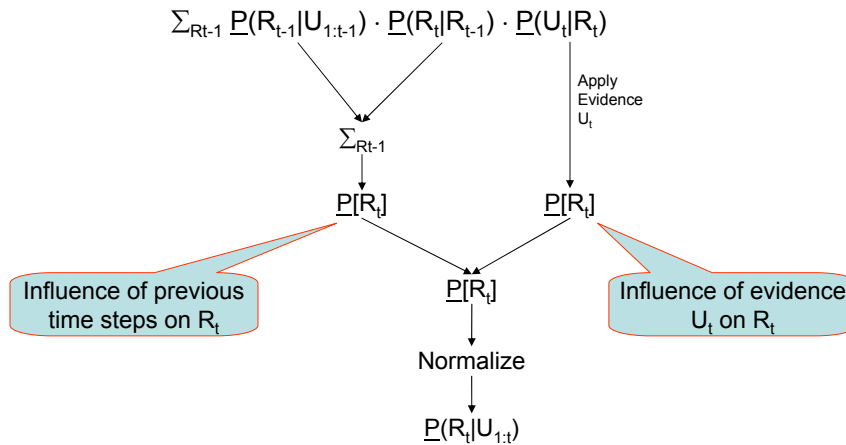
- *Filtering or Monitoring*: Compute the *belief state* given the history of sensor readings.  $\underline{P}(R_t|U_{1:t})$
- *Prediction*: Predict future state for some  $k > 0$ .  $\underline{P}(R_{t+k}|U_{1:t})$
- *Smoothing*: Reconstruct a previous state given subsequent evidence.  $\underline{P}(R_k|U_{1:t})$
- *Most Likely Explanation*: Reconstruct entire sequence of states given entire sequence of sensor readings.  $\operatorname{argmax}_{R_{1:n}} \underline{P}(R_{1:n}|U_{1:n})$

# Filtering by Variable Elimination

$$\begin{aligned}
 \underline{P}(R_1|U_1) &= \text{Normalize}[\text{ApplyEvidence}[U_1, \sum_{R_0} \underline{P}(R_0) \cdot \underline{P}(R_1|R_0) \cdot \underline{P}(U_1|R_1)]] \\
 &= \text{Normalize}[\sum_{R_0} \underline{P}(R_0) \cdot \underline{P}(R_1|R_0) \cdot \underline{P}[R_1]] \\
 &= \text{Normalize}[\underline{P}[R_1] \cdot \sum_{R_0} \underline{P}(R_0) \cdot \underline{P}(R_1|R_0)] \\
 &= \text{Normalize}[\underline{P}[R_1] \cdot \underline{P}[R_1]] \\
 &= \text{Normalize}[\underline{P}[R_1]]
 \end{aligned}$$

$$\begin{aligned}
 \underline{P}(R_2|U_{1:2}) &= \text{Normalize}[\text{ApplyEvidence}[U_{1:2}, \\
 &\quad \sum_{R_{0:1}} \underline{P}(R_0) \cdot \underline{P}(R_1|R_0) \cdot \underline{P}(U_1|R_1) \cdot \underline{P}(R_2|R_1) \cdot \underline{P}(U_2|R_2)]] \\
 &= \text{Normalize}[\sum_{R_{0:1}} \underline{P}(R_0) \cdot \underline{P}(R_1|R_0) \cdot \underline{P}[R_1] \cdot \underline{P}(R_2|R_1) \cdot \underline{P}[R_2]] \\
 &= \text{Normalize}[\sum_{R_1} [\sum_{R_0} \underline{P}(R_0) \cdot \underline{P}(R_1|R_0)] \cdot \underline{P}[R_1] \cdot \underline{P}(R_2|R_1) \cdot \underline{P}[R_2]] \\
 &= \text{Normalize}[\sum_{R_1} \underline{P}[R_1] \cdot \underline{P}[R_1] \cdot \underline{P}(R_2|R_1) \cdot \underline{P}[R_2]] \\
 &= \text{Normalize}[\underline{P}[R_2] \cdot \underline{P}[R_2]] \\
 &= \text{Normalize}[\underline{P}[R_2]]
 \end{aligned}$$

## General Pattern



(c) 2003 Thomas G. Dietterich

9

## The Forward Algorithm

Define:

$$\text{Forward}(\underline{P}(R_{t-1}|U_{1:t-1}), U_t) = \sum_{R_{t-1}} \underline{P}(R_{t-1}|U_{1:t-1}) \cdot \underline{P}(R_t|R_{t-1}) \cdot \text{ApplyEvidence}[U_t, \underline{P}(U_t|R_t)]$$

Then filtering can be written recursively as:

$$\underline{P}(R_t|U_{1:t}) = \text{Normalize}[\text{Forward}(\underline{P}(R_{t-1}|U_{1:t-1}), U_t)]$$

In general, we can iterate over multiple time steps:

$$\text{Forward}(\underline{P}(R_i|U_{1:i-1}), U_{i:t}) = \text{Forward}(\text{Forward}(\underline{P}(R_i|U_{1:i-1}), U_i), U_{i+1:t}) \text{ while } i \leq t$$

(c) 2003 Thomas G. Dietterich

10

# Example: Day 1

- day 1: Umbrella.  $U_1 = \text{yes}$

$$\underline{P}(R_1) = \text{Normalize}[\text{Forward}(\underline{P}(R_0), \text{yes})]$$

$$\text{Normalize} \left[ \sum_{R_0} \begin{array}{|c|c|} \hline R_0 & P(R_0) \\ \hline \text{no} & 0.5 \\ \hline \text{yes} & 0.5 \\ \hline \end{array} \cdot \begin{array}{|c|c|c|} \hline R_1 & R_0=\text{no} & R_0=\text{yes} \\ \hline \text{no} & 0.7 & 0.3 \\ \hline \text{yes} & 0.3 & 0.7 \\ \hline \end{array} \cdot \begin{array}{|c|c|c|} \hline U_1 & R_1=\text{no} & R_1=\text{yes} \\ \hline \text{no} & 0.8 & 0.1 \\ \hline \text{yes} & 0.2 & 0.9 \\ \hline \end{array} \right]$$

$$\text{Normalize} \left[ \sum_{R_0} \begin{array}{|c|c|c|} \hline R_1 & R_0=\text{no} & R_0=\text{yes} \\ \hline \text{no} & 0.7 * 0.5 & 0.3 * 0.5 \\ \hline \text{yes} & 0.3 * 0.5 & 0.7 * 0.5 \\ \hline \end{array} \cdot \begin{array}{|c|c|c|} \hline U_1 & R_1=\text{no} & R_1=\text{yes} \\ \hline \text{no} & 0.8 & 0.1 \\ \hline \text{yes} & 0.2 & 0.9 \\ \hline \end{array} \right]$$

# Example: Day 1 (continued)

$$\text{Normalize} \left[ \sum_{R_0} \begin{array}{|c|c|c|} \hline R_1 & R_0=\text{no} & R_0=\text{yes} \\ \hline \text{no} & 0.35 & 0.15 \\ \hline \text{yes} & 0.15 & 0.35 \\ \hline \end{array} \cdot \begin{array}{|c|c|c|} \hline U_1 & R_1=\text{no} & R_1=\text{yes} \\ \hline \text{no} & 0.8 & 0.1 \\ \hline \text{yes} & 0.2 & 0.9 \\ \hline \end{array} \right]$$

$$\text{Normalize} \left[ \begin{array}{|c|c|} \hline R_1 & \\ \hline \text{no} & 0.50 \\ \hline \text{yes} & 0.50 \\ \hline \end{array} \cdot \begin{array}{|c|c|c|} \hline U_1 & R_1=\text{no} & R_1=\text{yes} \\ \hline \text{no} & 0.8 & 0.1 \\ \hline \text{yes} & 0.2 & 0.9 \\ \hline \end{array} \right]$$

$$\text{Normalize} \left[ \begin{array}{|c|c|} \hline R_1 & P(R_1) \\ \hline \text{no} & 0.10 \\ \hline \text{yes} & 0.45 \\ \hline \end{array} \right] = \begin{array}{|c|c|} \hline R_1 & P(R_1) \\ \hline \text{no} & 0.18 \\ \hline \text{yes} & 0.82 \\ \hline \end{array}$$

# Example: Day 2

- Day 2:  $U_2 = \text{yes}$

$$\text{Normalize} \left[ \sum_{R_1} \begin{array}{|c|c|} \hline R_1 & P(R_1) \\ \hline \text{no} & 0.18 \\ \hline \text{yes} & 0.82 \\ \hline \end{array} \cdot \begin{array}{|c|c|c|} \hline R_2 & R_1=\text{no} & R_1=\text{yes} \\ \hline \text{no} & 0.7 & 0.3 \\ \hline \text{yes} & 0.3 & 0.7 \\ \hline \end{array} \cdot \begin{array}{|c|c|c|} \hline U_2 & R_2=\text{no} & R_2=\text{yes} \\ \hline \text{no} & 0.8 & 0.1 \\ \hline \text{yes} & 0.2 & 0.9 \\ \hline \end{array} \right]$$

$$\text{Normalize} \left[ \sum_{R_1} \begin{array}{|c|c|c|} \hline R_2 & R_1=\text{no} & R_1=\text{yes} \\ \hline \text{no} & 0.7 * 0.18 & 0.3 * 0.82 \\ \hline \text{yes} & 0.3 * 0.18 & 0.7 * 0.82 \\ \hline \end{array} \cdot \begin{array}{|c|c|c|} \hline U_2 & R_2=\text{no} & R_2=\text{yes} \\ \hline \text{no} & 0.8 & 0.1 \\ \hline \text{yes} & 0.2 & 0.9 \\ \hline \end{array} \right]$$

$$\text{Normalize} \left[ \sum_{R_1} \begin{array}{|c|c|c|} \hline R_2 & R_1=\text{no} & R_1=\text{yes} \\ \hline \text{no} & 0.127 & 0.245 \\ \hline \text{yes} & 0.055 & 0.573 \\ \hline \end{array} \cdot \begin{array}{|c|c|c|} \hline U_2 & R_2=\text{no} & R_2=\text{yes} \\ \hline \text{no} & 0.8 & 0.1 \\ \hline \text{yes} & 0.2 & 0.9 \\ \hline \end{array} \right] =$$

(c) 2003 Thomas G. Dietterich

13

# Day 2 (continued)

$$\text{Normalize} \left[ \sum_{R_1} \begin{array}{|c|c|c|} \hline R_2 & R_1=\text{no} & R_1=\text{yes} \\ \hline \text{no} & 0.127 & 0.245 \\ \hline \text{yes} & 0.055 & 0.573 \\ \hline \end{array} \cdot \begin{array}{|c|c|c|} \hline U_2 & R_2=\text{no} & R_2=\text{yes} \\ \hline \text{no} & 0.8 & 0.1 \\ \hline \text{yes} & 0.2 & 0.9 \\ \hline \end{array} \right] =$$

$$\text{Normalize} \left[ \begin{array}{|c|c|} \hline R_2 & \\ \hline \text{no} & 0.373 \\ \hline \text{yes} & 0.627 \\ \hline \end{array} \cdot \begin{array}{|c|c|c|} \hline U_2 & R_2=\text{no} & R_2=\text{yes} \\ \hline \text{no} & 0.8 & 0.1 \\ \hline \text{yes} & 0.2 & 0.9 \\ \hline \end{array} \right] =$$

$$\text{Normalize} \left[ \begin{array}{|c|c|} \hline R_2 & \\ \hline \text{no} & 0.075 \\ \hline \text{yes} & 0.565 \\ \hline \end{array} \right] = \begin{array}{|c|c|} \hline R_2 & P(R_2) \\ \hline \text{no} & 0.116 \\ \hline \text{yes} & 0.883 \\ \hline \end{array}$$

(c) 2003 Thomas G. Dietterich

14

## Prediction: Multiply by the Transition Probabilities and Sum Away


- $\underline{P}(R_{t+k} | U_{1:t}) = \sum_{R_{t:t+k-1}} \underline{P}(R_t | U_{1:t}) \cdot \underline{P}(R_{t+1}|R_t) \cdot \underline{P}(R_{t+2}|R_{t+1}) \cdot \dots \cdot \underline{P}(R_{t+k}|R_{t+k-1})$
- $\underline{P}(R_{t+1} | U_{1:t}) = \sum_{R_t} \underline{P}(R_t | U_{1:t}) \cdot \underline{P}(R_{t+1}|R_t)$
- $\underline{P}(R_{t+2} | U_{1:t}) = \sum_{R_{t+1}} \underline{P}(R_{t+1} | U_{1:t}) \cdot \underline{P}(R_{t+2}|R_{t+1})$
- ...

## Question: What Happens if We Predict Far Into the Future?

- Each multiplication by  $\underline{P}(R_{t+1}|R_t)$  makes our predictions “fuzzier”. Eventually, (for this problem) they converge to  $\langle 0.5, 0.5 \rangle$ . This is called the *stationary distribution* of the Markov process. Much is known about the stationary distribution and the rate of convergence. The stationary distribution depends on the transition probability distribution.



# Smoothing: Reconstructing $R_k$ given $U_{1:t}$

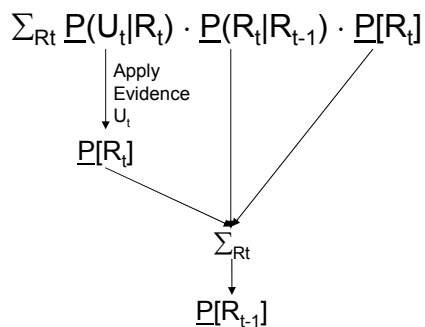


Assume  $k < t$ .

Example:  $k=3, t=7$ :

$$\underline{P}(R_3|U_{1:7}) = \text{Normalize}[ \text{ApplyEvidence}[U_{1:7}, \underline{P}(R_3|U_{1:3}) \cdot \underline{P}(U_{4:7}|R_3) ] ]$$

## The Backward Algorithm




## The Backward Algorithm (2)

$$\text{Backward}(\underline{P}[R_t], U_t) = \sum_{R_t} \text{ApplyEvidence}[U_t, \underline{P}(U_t|R_t)] \cdot \underline{P}(R_t|R_{t-1}) \cdot \underline{P}[R_t]$$

This can then be applied recursively

$$\underline{P}[R_{t-1}] = \text{Backward}(\underline{P}[R_t], U_t)$$

## Forward-Backward Algorithm for Smoothing



$$\underline{P}(R_k|U_{1:t}) = \text{Normalize}[ \text{Forward}(\underline{P}(R_0), U_{1:k}) \cdot \text{Backward}(\underline{1}, U_{k+1:t}) ]$$

# Umbrella Example: $\underline{P}(R_1|U_{1:2})$

Normalize[ Forward( $\underline{P}(R_0)$ ,  $U_1$ ) · Backward( $\underline{1}$ ,  $U_2$ ) ]

Forward( $\underline{P}(R_0)$ ,  $U_1$ ) =

$R_1$	$P(R_1)$
no	0.18
yes	0.82

Backward( $\underline{1}$ ,  $U_1$ ) =  $\sum_{R_2} \underline{1} \cdot \underline{P}(R_2|R_1) \cdot \underline{P}(U_2|R_2)$

## Backward from Day 2

$U_2 = \text{yes}$

$$\sum_{R_2} \begin{array}{|c|c|} \hline R_2 & P[R_2] \\ \hline \text{no} & 1 \\ \hline \text{yes} & 1 \\ \hline \end{array} \cdot \begin{array}{|c|c|c|} \hline R_2 & R_1=\text{no} & R_1=\text{yes} \\ \hline \text{no} & 0.7 & 0.3 \\ \hline \text{yes} & 0.3 & 0.7 \\ \hline \end{array} \cdot \begin{array}{|c|c|c|} \hline U_2 & R_2=\text{no} & R_2=\text{yes} \\ \hline \text{no} & 0.8 & 0.1 \\ \hline \text{yes} & 0.2 & 0.9 \\ \hline \end{array}$$

$$\sum_{R_2} \begin{array}{|c|c|c|} \hline R_2 & R_1=\text{no} & R_1=\text{yes} \\ \hline \text{no} & 0.7 * 1 * 0.2 & 0.3 * 1 * 0.2 \\ \hline \text{yes} & 0.3 * 1 * 0.9 & 0.7 * 1 * 0.9 \\ \hline \end{array}$$

$$\sum_{R_2} \begin{array}{|c|c|c|} \hline R_2 & R_1=\text{no} & R_1=\text{yes} \\ \hline \text{no} & 0.14 & 0.06 \\ \hline \text{yes} & 0.27 & 0.63 \\ \hline \end{array} = \begin{array}{|c|c|} \hline R_1 & P[R_1] \\ \hline \text{no} & 0.41 \\ \hline \text{yes} & 0.69 \\ \hline \end{array}$$

# Forward-Backward:

$$\text{Normalize} \left[ \begin{array}{cc} R_1 & P(R_1) \\ \text{no} & 0.18 \\ \text{yes} & 0.82 \end{array} \right] \cdot \begin{array}{cc} R_1 & P[R_1] \\ \text{no} & 0.41 \\ \text{yes} & 0.69 \end{array} =$$

$$\text{Normalize} \left[ \begin{array}{cc} R_1 & P(R_1) \\ \text{no} & 0.074 \\ \text{yes} & 0.566 \end{array} \right] = \begin{array}{cc} R_1 & P(R_1) \\ \text{no} & 0.115 \\ \text{yes} & 0.885 \end{array}$$

Notice that  $P(R_1=\text{yes}|U_1=\text{yes}) < P(R_1=\text{yes}|U_1=\text{yes}, U_2=\text{yes})$

Evidence from the future allows us to revise our beliefs about the past.

(c) 2003 Thomas G. Dietterich

23

# Most Likely Explanation


- Find  $\text{argmax}_{R_{1:n}} \underline{P}(R_{1:n}|U_{1:n})$ 
  - Note that this is the maximum over all *sequences* of rain states:  $R_{1:n}$
  - There are  $2^n$  such sequences!
  - Fortunately, there is a dynamic programming algorithm: the Viterbi Algorithm

(c) 2003 Thomas G. Dietterich

24

# Viterbi Algorithm

- Suppose we observe  $\langle \text{yes, yes, no, yes, yes} \rangle$  for  $U_{1:5}$
- Our goal is to find the best path through a “trellis” of possible rain states:



(c) 2003 Thomas G. Dietterich

25

## Max distributes over conformal product

$$\max_{A,B,C} \begin{array}{|c|c|c|} \hline & B & \\ \hline A=no & 0.10 & 0.20 \\ \hline A=yes & 0.30 & 0.40 \\ \hline \end{array} \cdot \begin{array}{|c|c|c|} \hline & C & \\ \hline B=no & 0.10 & 0.35 \\ \hline B=yes & 0.40 & 0.15 \\ \hline \end{array} =$$

$$\begin{array}{|c|c|c|c|} \hline B & C & A=no & A=yes \\ \hline no & no & 0.10 \cdot 0.10 & 0.20 \cdot 0.10 \\ \hline no & yes & 0.10 \cdot 0.40 & 0.20 \cdot 0.40 \\ \hline yes & no & 0.30 \cdot 0.35 & 0.40 \cdot 0.35 \\ \hline yes & yes & 0.30 \cdot 0.15 & 0.40 \cdot 0.15 \\ \hline \end{array} = \begin{array}{|c|c|c|c|} \hline B & C & A=no & A=yes \\ \hline no & no & 0.010 & 0.020 \\ \hline no & yes & 0.040 & 0.080 \\ \hline yes & no & 0.105 & 0.140 \\ \hline yes & yes & 0.045 & 0.060 \\ \hline \end{array}$$

(c) 2003 Thomas G. Dietterich

26

# Max propagation

$$\max_{A,B} \begin{array}{|c|c|c|} \hline B & A=no & A=yes \\ \hline no & 0.10 & 0.20 \\ \hline yes & 0.30 & 0.40 \\ \hline \end{array} \cdot \max_C \begin{array}{|c|c|c|} \hline C & B=no & B=yes \\ \hline no & 0.10 & 0.35 \\ \hline yes & 0.40 & 0.15 \\ \hline \end{array} =$$

$$\max_{A,B} \begin{array}{|c|c|c|} \hline B & A=no & A=yes \\ \hline no & 0.10 & 0.20 \\ \hline yes & 0.30 & 0.40 \\ \hline \end{array} \cdot \begin{array}{|c|c|} \hline B=no & B=yes \\ \hline 0.40 & 0.35 \\ \hline \end{array} =$$

$$\begin{array}{|c|c|c|} \hline B & A=no & A=yes \\ \hline no & 0.10 \cdot 0.40 & 0.20 \cdot 0.40 \\ \hline yes & 0.30 \cdot 0.35 & 0.40 \cdot 0.35 \\ \hline \end{array} = \begin{array}{|c|c|c|} \hline B & A=no & A=yes \\ \hline no & 0.040 & 0.080 \\ \hline yes & 0.105 & 0.140 \\ \hline \end{array}$$

(c) 2003 Thomas G. Dietterich

27

# Follow the Maxes

$$\max_{A,B,C} \begin{array}{|c|c|c|} \hline B & A=no & A=yes \\ \hline no & 0.10 & 0.20 \\ \hline yes & 0.30 & 0.40 \\ \hline \end{array} \cdot \begin{array}{|c|c|c|} \hline C & B=no & B=yes \\ \hline no & 0.10 & 0.35 \\ \hline yes & 0.40 & 0.15 \\ \hline \end{array} =$$

$$\begin{array}{|c|c|c|c|} \hline B & C & A=no & A=yes \\ \hline no & no & 0.10 \cdot 0.10 & 0.20 \cdot 0.10 \\ \hline no & yes & 0.10 \cdot 0.40 & 0.20 \cdot 0.40 \\ \hline yes & no & 0.30 \cdot 0.35 & 0.40 \cdot 0.35 \\ \hline yes & yes & 0.30 \cdot 0.15 & 0.40 \cdot 0.15 \\ \hline \end{array} = \begin{array}{|c|c|c|c|} \hline B & C & A=no & A=yes \\ \hline no & no & 0.010 & 0.020 \\ \hline no & yes & 0.040 & 0.080 \\ \hline yes & no & 0.105 & 0.140 \\ \hline yes & yes & 0.045 & 0.060 \\ \hline \end{array}$$

Because the “losers” (0.10 and 0.15) will be multiplied against the same values as the “winners” (0.40 and 0.35), they can never be the overall winners.

(c) 2003 Thomas G. Dietterich

28

# Extracting the Maximum Configuration

- Remember the winning combinations

$$\begin{array}{c}
 \max_{A,B} \begin{array}{|c|c|c|} \hline B & A=no & A=yes \\ \hline no & 0.10 & 0.20 \\ \hline yes & 0.30 & 0.40 \\ \hline \end{array} \cdot \max_C \begin{array}{|c|c|c|} \hline C & B=no & B=yes \\ \hline no & 0.10 & 0.35 \\ \hline yes & 0.40 & 0.15 \\ \hline \end{array} = \\
 \\
 \max_{A,B} \begin{array}{|c|c|c|} \hline B & A=no & A=yes \\ \hline no & 0.10 & 0.20 \\ \hline yes & 0.30 & 0.40 \\ \hline \end{array} \cdot \begin{array}{|c|c|} \hline B=no & B=yes \\ \hline 0.40 & 0.35 \\ \hline C=yes & C=no \\ \hline \end{array} = \\
 \\
 \begin{array}{|c|c|c|} \hline B & A=no & A=yes \\ \hline no & 0.10 \cdot 0.40 & 0.20 \cdot 0.40 \\ \hline yes & 0.30 \cdot 0.35 & 0.40 \cdot 0.35 \\ \hline \end{array} = \begin{array}{|c|c|c|} \hline B & A=no & A=yes \\ \hline no & 0.040 & 0.080 \\ \hline yes & 0.105 & 0.140 \\ \hline \end{array}
 \end{array}$$


(B=yes,A=yes) is winner combination. Corresponding value is C=no

# Viterbi Algorithm

$$\begin{aligned}
 \max_{R_{0:2}} \underline{P}(R_{0:2}|U_{1:2}) &= \\
 \max_{R_{0:2}} P(R_0) \cdot P(R_1|R_0) \cdot P(U_1|R_1) \cdot & \\
 P(R_2|R_1) \cdot P(U_2|R_2) &= \\
 \max_{R_2} P(U_2|R_2) \cdot [\max_{R_1} P(U_1|R_1) \cdot P(R_2|R_1)] & \\
 \cdot [\max_{R_0} P(R_0) \cdot P(R_1|R_0)] &=
 \end{aligned}$$


# Viterbi

- $[\max_{R_0} \underline{P}(R_0) \cdot \underline{P}(R_1|R_0)] \cdot \underline{P}(U_1|R_1)$



# Viterbi


- $[\max_{R_0} \underline{P}(R_0) \cdot \underline{P}(R_1|R_0)] \cdot \underline{P}(U_1|R_1)$






# Viterbi

- $$[\max_{R_0} \underline{P}(R_0) \cdot \underline{P}(R_1|R_0)] \cdot \underline{P}(U_1|R_1)$$




# Viterbi

- $$[\max_{R_1} \underline{P}(R_1) \cdot \underline{P}(R_2|R_1)] \cdot \underline{P}(U_2|R_2)$$



# Viterbi

- $[\max_{R_2} \underline{P}[R_1] \cdot \underline{P}(R_2|R_1)] \cdot \underline{P}(U_2|R_2)$




(c) 2003 Thomas G. Dietterich

35

# Viterbi

- $[\max_{R_2} \underline{P}[R_1] \cdot \underline{P}(R_2|R_1)] \cdot \underline{P}(U_2|R_2)$




(c) 2003 Thomas G. Dietterich

36

# Viterbi

- $[\max_{R_3} \underline{P}[R_2] \cdot \underline{P}(R_3|R_2)] \cdot \underline{P}(U_3|R_3)$




(c) 2003 Thomas G. Dietterich

37

# Viterbi

- $[\max_{R_3} \underline{P}[R_2] \cdot \underline{P}(R_3|R_2)] \cdot \underline{P}(U_3|R_3)$




(c) 2003 Thomas G. Dietterich

38

# Viterbi

- $[\max_{R_3} \underline{P}[R_2] \cdot \underline{P}(R_3|R_2)] \cdot \underline{P}(U_3|R_3)$




(c) 2003 Thomas G. Dietterich

39

# Viterbi

- $[\max_{R_4} \underline{P}[R_3] \cdot \underline{P}(R_4|R_3)] \cdot \underline{P}(U_4|R_4)$




(c) 2003 Thomas G. Dietterich

40

# Viterbi

- $[\max_{R_4} \underline{P}[R_3] \cdot \underline{P}(R_4|R_3)] \cdot \underline{P}(U_4|R_4)$




(c) 2003 Thomas G. Dietterich

41

# Viterbi

- $[\max_{R_4} \underline{P}[R_3] \cdot \underline{P}(R_4|R_3)] \cdot \underline{P}(U_4|R_4)$




(c) 2003 Thomas G. Dietterich

42


# Viterbi

- $[\max_{R_4} \underline{P}[R_4] \cdot \underline{P}(R_5|R_4)] \cdot \underline{P}(U_5|R_5)$




# Viterbi

- $[\max_{R_4} \underline{P}[R_4] \cdot \underline{P}(R_5|R_4)] \cdot \underline{P}(U_5|R_5)$



# Viterbi

- $[\max_{R_4} \underline{P}[R_4] \cdot \underline{P}(R_5|R_4)] \cdot \underline{P}(U_5|R_5)$




(c) 2003 Thomas G. Dietterich

45

# Viterbi

- $\max_{R_5} \underline{P}[R_5]$




(c) 2003 Thomas G. Dietterich

46

# Viterbi

- Traceback




(c) 2003 Thomas G. Dietterich

47

## Dynamic Bayesian Networks

- Multiple State Variables and Multiple Sensors
- Robot state variables:
  - Position  $X_t$
  - Velocity  $\dot{X}_t$
  - Battery power
- Sensors
  - Battery meter
  - GPS sensor
- DBN captures sparseness in the interactions among the variables



(c) 2003 Thomas G. Dietterich

48



## Inference for DBNs

- Problem: The cost of inference for DBNs is generally exponential in the number of state variables.
- Solution: Approximate Inference using Particle Filters

## Particle Filters

- Key idea: Represent  $P(X_t, X_{dot_t}, B_{att_t} | Z_{1:t}, B_{M_{1:t}})$  as a set of points (“particles”)
- Implement the Forward algorithm by simulating the behavior of these points

## Particle Filtering (we will use HMMs for simplicity)

- HMM:  $\underline{P}(X_t|X_{t-1})$ ;  $\underline{P}(Z_t|X_t)$ ;  $\underline{P}(X_1)$
- At each time  $t$ , we will have a set of points  $S_t = \{x_1, \dots, x_N\}$  that represent  $\underline{P}(X_t|Z_{1:t})$ .
- Step 1: Apply  $\underline{P}(X_{t+1}|X_t)$ : Push each point “forward” in time  $x_i \sim \underline{P}(X_{t+1} | x_i)$
- Step 2: Apply evidence. Assign a weight to each point:  $w_i = \underline{P}(Z_{t+1}|x_i)$
- Step 3: Normalize by drawing a new sample according to weight  $w_i$ .
  - Let  $W = \sum_i w_i$  be the total amount of weight.
  - Draw  $N$  points *with replacement* from  $S = \{x_i\}$ , where point  $x_i$  has probability  $w_i/W$  of being chosen.

## More on Particle Filters

- Sebastian Thrun ([cs.stanford.edu](http://cs.stanford.edu))
- Dieter Fox ([cs.washington.edu](http://cs.washington.edu))



# Signal Processing

- Divide speech signal into short chunks (e.g., 10ms) called “frames”
  - Frames overlap by 5ms
- Extract from each frame a vector of real-valued “features”
  - Frequency x Energy features (“Cepstral coefficients”)
  - Changes in these, etc.


# Generative Model of Frames

- $P(\text{frame} \mid \text{phone})$ 
  - Vector Quantization: Discretize frames by clustering them into 256 clusters.
    - Frame becomes single 256-valued variable
  - Model frame as a mixture of multi-variate Gaussian random variables whose mean and variance depends on the phone.

# HMM Models of Phones

- A phone lasts 50-100 ms (= 10-20 frames)
  - Different pronunciations, speaking rates

Phone HMM for [m]:



Here, C<sub>1</sub>, C<sub>2</sub>, etc. are frame cluster numbers

Output probabilities for the phone HMM:


Onset:	Mid:	End:
C <sub>1</sub> : 0.5	C <sub>3</sub> : 0.2	C <sub>4</sub> : 0.1
C <sub>2</sub> : 0.2	C <sub>4</sub> : 0.7	C <sub>6</sub> : 0.5
C <sub>3</sub> : 0.3	C <sub>5</sub> : 0.1	

© 2003 Thomas G. Dietterich


# HMM Models of Words

- A word may produce more than one possible phone sequence
  - Different pronunciations: “[t][ah][m][ey][t][ow]” versus “[t][ah][m][aa][t][ow]”
  - Coarticulation effects: “[t][ah][m][ey][t][ow]” versus “[t][ow][m][ey][t][ow]”

(a) Word model with dialect variation:



(b) Word model with coarticulation and dialect variations



© 2003 Thomas G. Dietterich


# Language Model

- Bigram or Trigram Models

## “Macro Expanding”

- We can combine the language model, word models, and phone models to obtain a very large HMM that contains only phones and frames

## Fragment of the Flattened Phone Model – Each state generates frames



61

## Learning the Model Parameters

- Fully-supervised: Manually label frames with phone states (onset, middle, end)
  - Very time-consuming
- Abstract supervision: Label each sentence with the sequence of words spoken
  - Treat phones as hidden variables
  - Apply EM algorithm for learning Bayesian networks with missing variables

(c) 2003 Thomas G. Dietterich

62

# Speech Recognition

- Viterbi algorithm finds most likely path through the flattened HMM
  - Does not necessarily find the most likely sequence of words. Why not?
- Beam Search
  - Too expensive to compute: Branching factor of 20,000
  - Keep track of the B most likely states in the HMM at each time t
- “It’s hard to wreck a nice beach”