

---

# Markov Decision Processes

Philipp Koehn

3 November 2015



# Outline



- Hidden Markov models
- Inference: filtering, smoothing, best sequence
- Kalman filters (a brief mention)
- Dynamic Bayesian networks
- Speech recognition


# Time and Uncertainty



- The world changes; we need to track and predict it
- Diabetes management vs vehicle diagnosis
- Basic idea: sequence of state and evidence variables
- $\mathbf{X}_t$  = set of unobservable state variables at time  $t$   
e.g., *BloodSugar<sub>t</sub>*, *StomachContents<sub>t</sub>*, etc.
- $\mathbf{E}_t$  = set of observable evidence variables at time  $t$   
e.g., *MeasuredBloodSugar<sub>t</sub>*, *PulseRate<sub>t</sub>*, *FoodEaten<sub>t</sub>*
- This assumes **discrete time**; step size depends on problem
- Notation:  $\mathbf{X}_{a:b} = \mathbf{X}_a, \mathbf{X}_{a+1}, \dots, \mathbf{X}_{b-1}, \mathbf{X}_b$


# Markov Processes (Markov Chains)

- Construct a Bayes net from these variables: parents?
- Markov assumption:  $\mathbf{X}_t$  depends on **bounded** subset of  $\mathbf{X}_{0:t-1}$
- First-order Markov process:  $\mathbf{P}(\mathbf{X}_t | \mathbf{X}_{0:t-1}) = \mathbf{P}(\mathbf{X}_t | \mathbf{X}_{t-1})$   
Second-order Markov process:  $\mathbf{P}(\mathbf{X}_t | \mathbf{X}_{0:t-1}) = \mathbf{P}(\mathbf{X}_t | \mathbf{X}_{t-2}, \mathbf{X}_{t-1})$



- Sensor Markov assumption:  $\mathbf{P}(\mathbf{E}_t | \mathbf{X}_{0:t}, \mathbf{E}_{0:t-1}) = \mathbf{P}(\mathbf{E}_t | \mathbf{X}_t)$
- Stationary process: transition model  $\mathbf{P}(\mathbf{X}_t | \mathbf{X}_{t-1})$  and sensor model  $\mathbf{P}(\mathbf{E}_t | \mathbf{X}_t)$  fixed for all  $t$

# Example



- First-order Markov assumption not exactly true in real world!
- Possible fixes:
  1. **Increase order** of Markov process
  2. **Augment state**, e.g., add  $Temp_t$ ,  $Pressure_t$

# inference

# Inference Tasks



- **Filtering:**  $P(\mathbf{X}_t | \mathbf{e}_{1:t})$   
belief state—input to the decision process of a rational agent■
- **Smoothing:**  $P(\mathbf{X}_k | \mathbf{e}_{1:t})$  for  $0 \leq k < t$   
better estimate of past states, essential for learning■
- **Most likely explanation:**  $\arg \max_{\mathbf{x}_{1:t}} P(\mathbf{x}_{1:t} | \mathbf{e}_{1:t})$   
speech recognition, decoding with a noisy channel

# Filtering



- Aim: devise a **recursive** state estimation algorithm

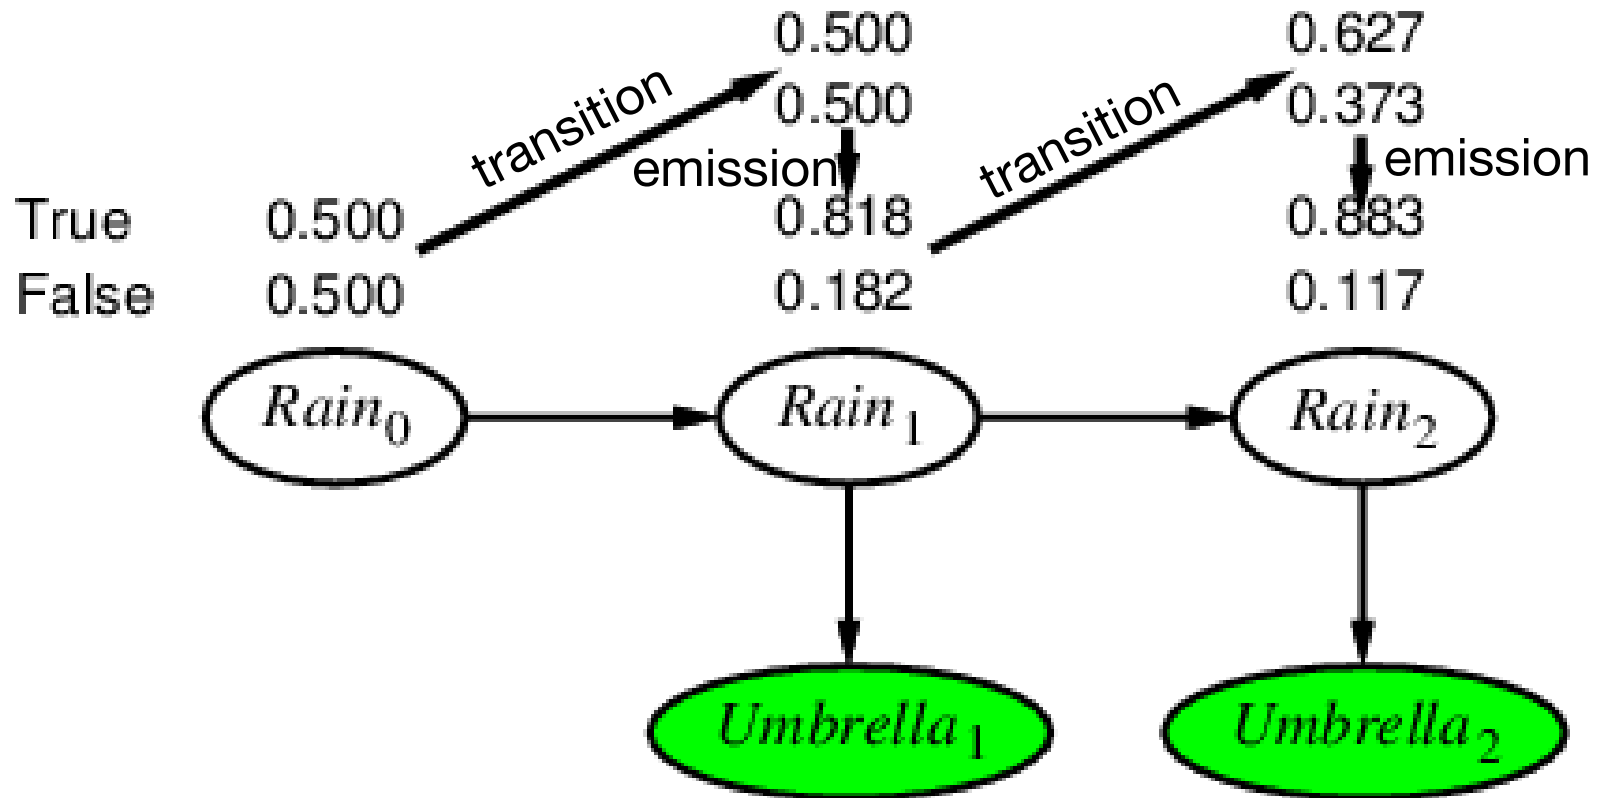
$$\begin{aligned} \mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t+1}) &= \mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t}, \mathbf{e}_{t+1}) \\ &= \alpha \mathbf{P}(\mathbf{e}_{t+1}|\mathbf{X}_{t+1}, \mathbf{e}_{1:t}) \mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t}) \quad (\text{Bayes rule}) \blacksquare \\ &= \alpha \mathbf{P}(\mathbf{e}_{t+1}|\mathbf{X}_{t+1}) \mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t}) \quad (\text{Sensor Markov assumption}) \blacksquare \\ &= \alpha \mathbf{P}(\mathbf{e}_{t+1}|\mathbf{X}_{t+1}) \sum_{\mathbf{x}_t} \mathbf{P}(\mathbf{X}_{t+1}|\mathbf{x}_t, \mathbf{e}_{1:t}) P(\mathbf{x}_t|\mathbf{e}_{1:t}) \quad (\text{multiplying out}) \blacksquare \\ &= \alpha \mathbf{P}(\mathbf{e}_{t+1}|\mathbf{X}_{t+1}) \sum_{\mathbf{x}_t} \mathbf{P}(\mathbf{X}_{t+1}|\mathbf{x}_t) P(\mathbf{x}_t|\mathbf{e}_{1:t}) \quad (\text{first order Markov model}) \blacksquare \end{aligned}$$

- Summary: 
$$\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t+1}) = \alpha \underbrace{\mathbf{P}(\mathbf{e}_{t+1}|\mathbf{X}_{t+1})}_{\text{emission}} \sum_{\mathbf{x}_t} \underbrace{\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{x}_t)}_{\text{transition}} \underbrace{P(\mathbf{x}_t|\mathbf{e}_{1:t})}_{\text{recursive call}}$$


- $\mathbf{f}_{1:t+1} = \text{FORWARD}(\mathbf{f}_{1:t}, \mathbf{e}_{t+1})$  where  $\mathbf{f}_{1:t} = \mathbf{P}(\mathbf{X}_t|\mathbf{e}_{1:t})$   
Time and space **constant** (independent of  $t$ )



# Filtering Example




# Smoothing



- If full sequence is known  
⇒ what is the state probability  $P(X_k | e_{1:t})$  including future evidence?
- Smoothing: sum over all paths

# Smoothing



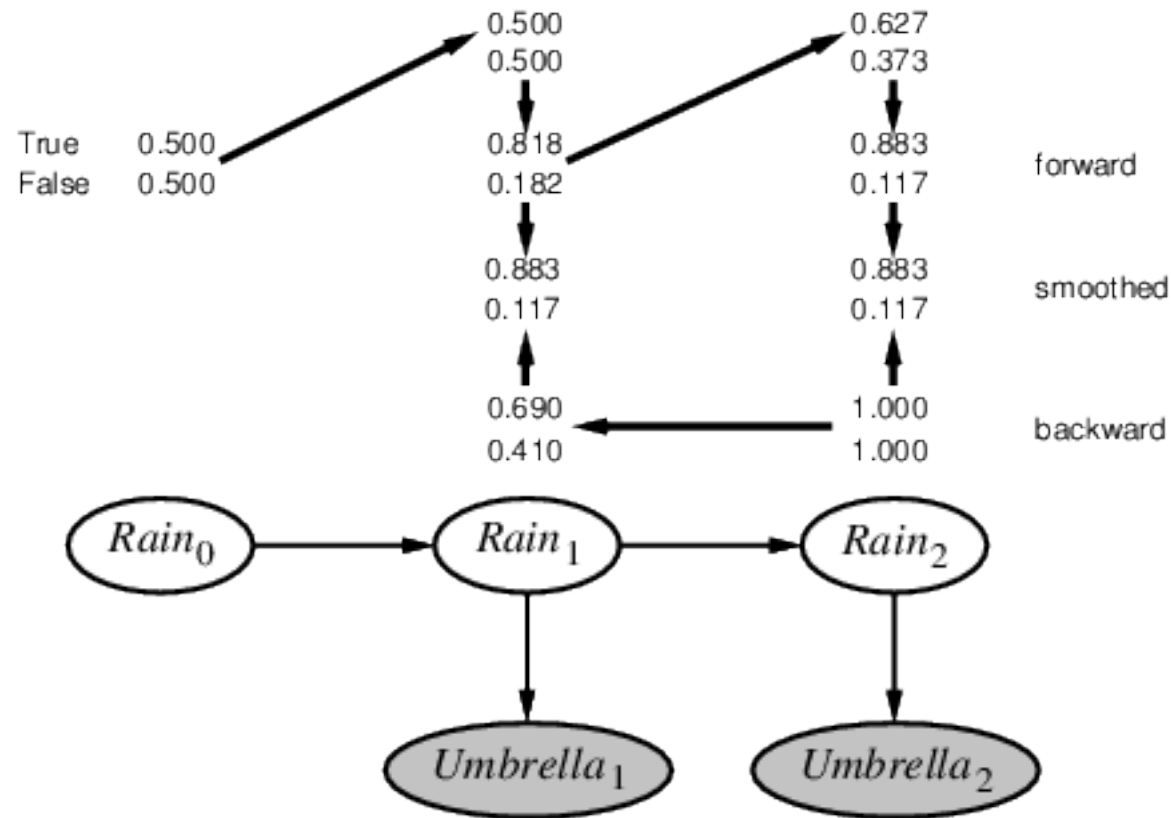
- Divide evidence  $\mathbf{e}_{1:t}$  into  $\mathbf{e}_{1:k}$ ,  $\mathbf{e}_{k+1:t}$ :

$$\begin{aligned}
 \mathbf{P}(\mathbf{X}_k | \mathbf{e}_{1:t}) &= \mathbf{P}(\mathbf{X}_k | \mathbf{e}_{1:k}, \mathbf{e}_{k+1:t}) \\
 &= \alpha \mathbf{P}(\mathbf{X}_k | \mathbf{e}_{1:k}) \mathbf{P}(\mathbf{e}_{k+1:t} | \mathbf{X}_k, \mathbf{e}_{1:k}) \\
 &= \alpha \mathbf{P}(\mathbf{X}_k | \mathbf{e}_{1:k}) \mathbf{P}(\mathbf{e}_{k+1:t} | \mathbf{X}_k) \\
 &= \alpha \mathbf{f}_{1:k} \mathbf{b}_{k+1:t}
 \end{aligned}$$

- Backward message  $\mathbf{b}_{k+1:t}$  computed by a backwards recursion

$$\begin{aligned}
 \mathbf{P}(\mathbf{e}_{k+1:t} | \mathbf{X}_k) &= \sum_{\mathbf{x}_{k+1}} \mathbf{P}(\mathbf{e}_{k+1:t} | \mathbf{X}_k, \mathbf{x}_{k+1}) \mathbf{P}(\mathbf{x}_{k+1} | \mathbf{X}_k) \\
 &= \sum_{\mathbf{x}_{k+1}} P(\mathbf{e}_{k+1:t} | \mathbf{x}_{k+1}) \mathbf{P}(\mathbf{x}_{k+1} | \mathbf{X}_k) \\
 &= \sum_{\mathbf{x}_{k+1}} P(\mathbf{e}_{k+1} | \mathbf{x}_{k+1}) P(\mathbf{e}_{k+2:t} | \mathbf{x}_{k+1}) \mathbf{P}(\mathbf{x}_{k+1} | \mathbf{X}_k)
 \end{aligned}$$

# Smoothing Example



Forward-backward algorithm: cache forward messages along the way  
 Time linear in  $t$  (polytree inference), space  $O(t|f|)$

# Most Likely Explanation

- Most likely sequence  $\neq$  sequence of most likely states
- Most likely path to each  $\mathbf{x}_{t+1}$   
= most likely path to **some**  $\mathbf{x}_t$  plus one more step

$$\begin{aligned} & \max_{\mathbf{x}_1 \dots \mathbf{x}_t} \mathbf{P}(\mathbf{x}_1, \dots, \mathbf{x}_t, \mathbf{X}_{t+1} | \mathbf{e}_{1:t+1}) \\ & = \mathbf{P}(\mathbf{e}_{t+1} | \mathbf{X}_{t+1}) \max_{\mathbf{x}_t} \left( \mathbf{P}(\mathbf{X}_{t+1} | \mathbf{x}_t) \max_{\mathbf{x}_1 \dots \mathbf{x}_{t-1}} P(\mathbf{x}_1, \dots, \mathbf{x}_{t-1}, \mathbf{x}_t | \mathbf{e}_{1:t}) \right) \blacksquare \end{aligned}$$

- Identical to filtering, except  $\mathbf{f}_{1:t}$  replaced by

$$\mathbf{m}_{1:t} = \max_{\mathbf{x}_1 \dots \mathbf{x}_{t-1}} \mathbf{P}(\mathbf{x}_1, \dots, \mathbf{x}_{t-1}, \mathbf{X}_t | \mathbf{e}_{1:t})$$

i.e.,  $\mathbf{m}_{1:t}(i)$  gives the probability of the most likely path to state  $i$ .  $\blacksquare$


- Update has sum replaced by max, giving the **Viterbi algorithm**:

$$\mathbf{m}_{1:t+1} = \mathbf{P}(\mathbf{e}_{t+1} | \mathbf{X}_{t+1}) \max_{\mathbf{x}_t} (\mathbf{P}(\mathbf{X}_{t+1} | \mathbf{x}_t) \mathbf{m}_{1:t})$$

Also requires back-pointers for backward pass to retrieve best sequence

$$\mathbf{b}_{\mathbf{x}_{t+1}, t+1} = \operatorname{argmax}_{\mathbf{x}_t} (\mathbf{P}(\mathbf{X}_{t+1} | \mathbf{x}_t) \mathbf{m}_{1:t})$$

# Viterbi Example



# Hidden Markov Models


- $\mathbf{X}_t$  is a single, discrete variable (usually  $\mathbf{E}_t$  is too)  
Domain of  $X_t$  is  $\{1, \dots, S\}$
- Transition matrix  $\mathbf{T}_{ij} = P(X_t = j | X_{t-1} = i)$ , e.g.,  $\begin{pmatrix} 0.7 & 0.3 \\ 0.3 & 0.7 \end{pmatrix}$
- Sensor matrix  $\mathbf{O}_t$  for each time step, diagonal elements  $P(e_t | X_t = i)$   
e.g., with  $U_1 = \text{true}$ ,  $\mathbf{O}_1 = \begin{pmatrix} 0.9 & 0 \\ 0 & 0.2 \end{pmatrix}$
- Forward and backward messages as column vectors:
$$\mathbf{f}_{1:t+1} = \alpha \mathbf{O}_{t+1} \mathbf{T}^\top \mathbf{f}_{1:t}$$
$$\mathbf{b}_{k+1:t} = \mathbf{T} \mathbf{O}_{k+1} \mathbf{b}_{k+2:t}$$
- Forward-backward algorithm needs time  $O(S^2t)$  and space  $O(St)$

# kalman filters



# Kalman Filters

- Modelling systems described by a set of continuous variables, e.g., tracking a bird flying— $\mathbf{X}_t = X, Y, Z, \dot{X}, \dot{Y}, \dot{Z}$ .  
Airplanes, robots, ecosystems, economies, chemical plants, planets, ...



- Gaussian prior, linear Gaussian transition model and sensor model

# Updating Gaussian Distributions

- Prediction step: if  $\mathbf{P}(\mathbf{X}_t|\mathbf{e}_{1:t})$  is Gaussian, then prediction

$$\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t}) = \int_{\mathbf{x}_t} \mathbf{P}(\mathbf{X}_{t+1}|\mathbf{x}_t)P(\mathbf{x}_t|\mathbf{e}_{1:t}) d\mathbf{x}_t$$

is Gaussian. If  $\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t})$  is Gaussian, then the updated distribution

$$\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t+1}) = \alpha \mathbf{P}(\mathbf{e}_{t+1}|\mathbf{X}_{t+1})\mathbf{P}(\mathbf{X}_{t+1}|\mathbf{e}_{1:t})$$


is Gaussian

- Hence  $\mathbf{P}(\mathbf{X}_t|\mathbf{e}_{1:t})$  is multivariate Gaussian  $N(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$  for all  $t$
- General (nonlinear, non-Gaussian) process: description of posterior grows **unboundedly** as  $t \rightarrow \infty$

# Simple 1-D Example

- Gaussian random walk on  $X$ -axis, s.d.  $\sigma_x$ , sensor s.d.  $\sigma_z$

$$\mu_{t+1} = \frac{(\sigma_t^2 + \sigma_x^2)z_{t+1} + \sigma_z^2\mu_t}{\sigma_t^2 + \sigma_x^2 + \sigma_z^2} \quad \sigma_{t+1}^2 = \frac{(\sigma_t^2 + \sigma_x^2)\sigma_z^2}{\sigma_t^2 + \sigma_x^2 + \sigma_z^2}$$



# General Kalman Update

- Transition and sensor models:

$$\begin{aligned}P(\mathbf{x}_{t+1}|\mathbf{x}_t) &= N(\mathbf{F}\mathbf{x}_t, \Sigma_x)(\mathbf{x}_{t+1}) \\P(\mathbf{z}_t|\mathbf{x}_t) &= N(\mathbf{H}\mathbf{x}_t, \Sigma_z)(\mathbf{z}_t)\end{aligned}$$

$\mathbf{F}$  is the matrix for the transition;  $\Sigma_x$  the transition noise covariance  
 $\mathbf{H}$  is the matrix for the sensors;  $\Sigma_z$  the sensor noise covariance


- Filter computes the following update:

$$\begin{aligned}\mu_{t+1} &= \mathbf{F}\mu_t + \mathbf{K}_{t+1}(\mathbf{z}_{t+1} - \mathbf{H}\mathbf{F}\mu_t) \\ \Sigma_{t+1} &= (\mathbf{I} - \mathbf{K}_{t+1})(\mathbf{F}\Sigma_t\mathbf{F}^\top + \Sigma_x)\end{aligned}$$


where  $\mathbf{K}_{t+1} = (\mathbf{F}\Sigma_t\mathbf{F}^\top + \Sigma_x)\mathbf{H}^\top(\mathbf{H}(\mathbf{F}\Sigma_t\mathbf{F}^\top + \Sigma_x)\mathbf{H}^\top + \Sigma_z)^{-1}$   
is the **Kalman gain matrix**

- $\Sigma_t$  and  $\mathbf{K}_t$  are independent of observation sequence, so compute offline

# 2-D Tracking Example: Filtering





# 2-D Tracking Example: Smoothing



# dynamic bayesian networks

# Dynamic Bayesian Networks


- $\mathbf{X}_t, \mathbf{E}_t$  contain arbitrarily many variables in a sequentialized Bayes net





# DBNs vs. HMMs



- Every HMM is a single-variable DBN; every discrete DBN is an HMM



- Sparse dependencies  $\Rightarrow$  exponentially fewer parameters;  
e.g., 20 state variables, three parents each  
DBN has  $20 \times 2^3 = 160$  parameters, HMM has  $2^{20} \times 2^{20} \approx 10^{12}$


# DBNs vs Kalman Filters

- Every Kalman filter model is a DBN, but few DBNs are KFs; real world requires non-Gaussian posteriors
- E.g., where my keys? What's the battery charge?



# Exact Inference in DBNs


- Naive method: **unroll** the network and run any exact algorithm




- Problem: inference cost for each update grows with  $t$
- **Rollup filtering**: add slice  $t + 1$ , “sum out” slice  $t$  using variable elimination
- Largest factor is  $O(d^{n+1})$ , update cost  $O(d^{n+2})$   
(cf. HMM update cost  $O(d^{2n})$ )

# Likelihood Weighting for DBNs

- Set of weighted samples approximates the belief state




- LW samples pay no attention to the evidence!
  - ⇒ fraction “agreeing” falls exponentially with  $t$
  - ⇒ number of samples required grows exponentially with  $t$



# Particle Filtering

- Basic idea: ensure that the population of samples (“particles”) tracks the high-likelihood regions of the state-space
- Replicate particles proportional to likelihood for  $e_t$



- Widely used for tracking nonlinear systems, esp. in vision
- Also used for simultaneous localization and mapping in mobile robots  
 $10^5$ -dimensional state space

# speech recognition

## It's not easy to wreck a nice beach

- Speech signals are noisy, variable, ambiguous
- What is the **most likely** word sequence, given the speech signal?  
I.e., choose *Words* to maximize  $P(\textit{Words}|\textit{signal})$
- Use Bayes' rule:  
$$P(\textit{Words}|\textit{signal}) = \alpha P(\textit{signal}|\textit{Words})P(\textit{Words})$$
  
i.e., decomposes into **acoustic model** + **language model**
- *Words* are the hidden state sequence, *signal* is the observation sequence

# Phones

- All human speech is composed from 40-50 **phones**, determined by the configuration of **articulators** (lips, teeth, tongue, vocal cords, air flow)
- Form an intermediate level of hidden states between words and signal  
 ⇒ acoustic model = pronunciation model + phone model
- ARPAbet designed for American English


[iy]	b <u>e</u> at	[b]	<u>b</u> et	[p]	<u>p</u> et
[ih]	b <u>i</u> t	[ch]	<u>Ch</u> et	[r]	<u>r</u> at
[ey]	b <u>e</u> t	[d]	<u>d</u> ebt	[s]	<u>s</u> et
[ao]	b <u>ou</u> ght	[hh]	<u>h</u> at	[th]	<u>th</u> ick
[ow]	b <u>o</u> at	[hv]	<u>h</u> igh	[dh]	<u>th</u> at
[er]	<u>B</u> ert	[l]	<u>l</u> et	[w]	<u>w</u> et
[ix]	ros <u>e</u> s	[ng]	si <u>ng</u>	[en]	butt <u>o</u> n
⋮	⋮	⋮	⋮	⋮	⋮

e.g., “ceiling” is [s iy l ih ng] / [s iy l ix ng] / [s iy l en]




# Speech Sounds

- Raw signal is the microphone displacement as a function of time; processed into overlapping 30ms **frames**, each described by **features**



- Frame features are typically **formants**—peaks in the power spectrum


# Speech Spectrogram



- Frame features in  $P(\text{features}|\text{phone})$  summarized by
  - an integer in  $[0 \dots 255]$  (using **vector quantization**); or
  - the parameters of a mixture of Gaussians■
- **Three-state phones**: each phone has three phases (Onset, Mid, End)  
E.g., [t] has silent Onset, explosive Mid, hissing End  
⇒  $P(\text{features}|\text{phone}, \text{phase})$ ■
- **Triphone context**: each phone becomes  $n^2$  distinct phones, depending on the phones to its left and right  
E.g., [t] in “star” is written [t(s,aa)] (different from “tar”!)■
- Triphones useful for handling **coarticulation** effects: the articulators have inertia and cannot switch instantaneously between positions  
E.g., [t] in “eighth” has tongue against front teeth

# Phone Model Example

Phone HMM for [m]:




Output probabilities for the phone HMM:

Onset:	Mid:	End:
C1: 0.5	C3: 0.2	C4: 0.1
C2: 0.2	C4: 0.7	C6: 0.5
C3: 0.3	C5: 0.1	C7: 0.4

# Word Pronunciation Models

- Each word is described as a distribution over phone sequences
- Distribution represented as an HMM transition model



$$P([touwmeytow]|\text{"tomato"}) = P([touwmaatow]|\text{"tomato"}) = 0.1$$

$$P([tahmeytow]|\text{"tomato"}) = P([tahmaatow]|\text{"tomato"}) = 0.4$$

- Structure is created manually, transition probabilities learned from data

# Recognition of Isolated Words

- Phone models + word models fix likelihood  $P(e_{1:t}|word)$  for isolated word

$$P(word|e_{1:t}) = \alpha P(e_{1:t}|word)P(word)$$

- Prior probability  $P(word)$  obtained simply by counting word frequencies

$P(e_{1:t}|word)$  can be computed recursively: define

$$\alpha_{1:t} = \mathbf{P}(\mathbf{X}_t, \mathbf{e}_{1:t})$$

and use the recursive update

$$\alpha_{1:t+1} = \text{FORWARD}(\ell_{1:t}, \mathbf{e}_{t+1})$$

and then  $P(e_{1:t}|word) = \sum_{\mathbf{x}_t} \alpha_{1:t}(\mathbf{x}_t)$

- Isolated-word dictation systems with training reach 95–99% accuracy

# Continuous Speech

- Not just a sequence of isolated-word recognition problems!
  - adjacent words highly correlated
  - sequence of most likely words  $\neq$  most likely sequence of words
  - segmentation: there are few gaps in speech
  - cross-word coarticulation—e.g., “next thing”
- Complications
  - mismatch between speaker in training and test
  - noise
  - crosstalk
  - bad microphone position
- Continuous speech systems manage over 90% accuracy on a good day

# Language Model

- Prior probability of a word sequence is given by chain rule:

$$P(w_1 \cdots w_n) = \prod_{i=1}^n P(w_i | w_1 \cdots w_{i-1})$$

- Bigram model:

$$P(w_i | w_1 \cdots w_{i-1}) \approx P(w_i | w_{i-1})$$

- Train by counting all word pairs in a large text corpus
- More sophisticated models (trigrams, grammars, etc.) help a little bit




# Combined HMM



- States of the combined language+word+phone model are labelled by the word we're in + the phone in that word + the phone state in that phone
- Viterbi algorithm finds the most likely **phone state** sequence
- Does segmentation by considering all possible word sequences and boundaries
- Doesn't always give the most likely word sequence because each word sequence is the sum over many state sequences
- Jelinek invented A\* in 1969 a way to find most likely word sequence where "step cost" is  $-\log P(w_i|w_{i-1})$

# DBNs for Speech Recognition

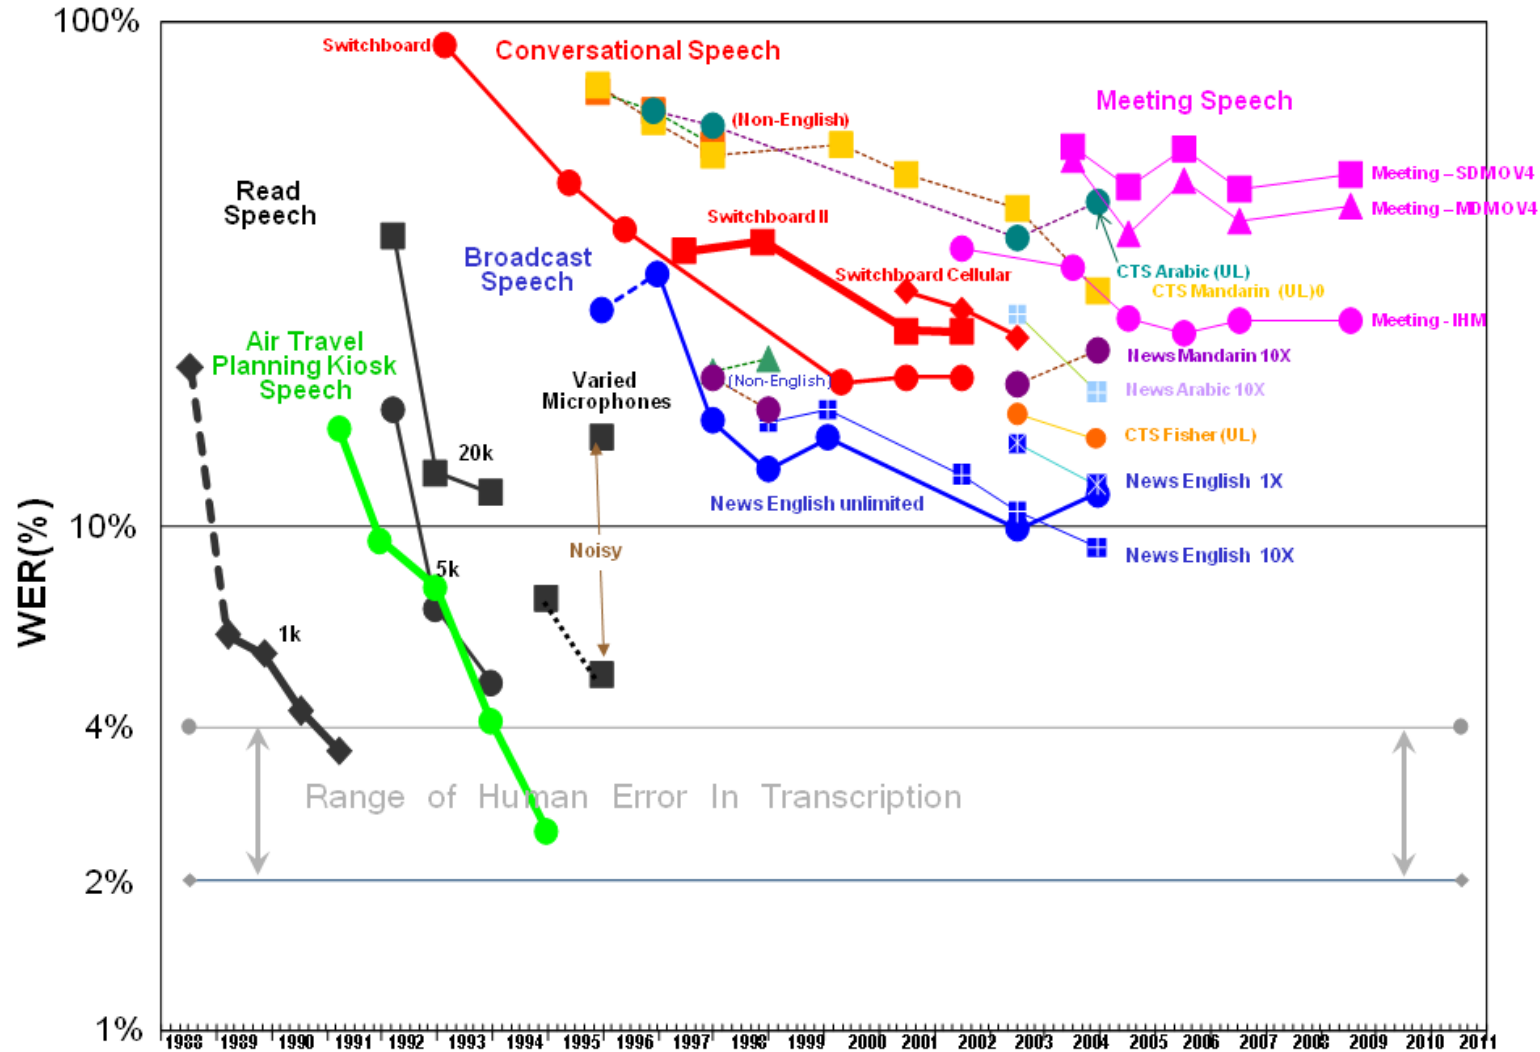
41



- Also easy to add variables for, e.g., gender, accent, speed
- Zweig and Russell (1998) show up to 40% error reduction over HMMs


# Progress

## NIST STT Benchmark Test History – May. '09



# Progress

43



# Summary

- Temporal models use state and sensor variables replicated over time
- Markov assumptions and stationarity assumption, so we need
  - transition model  $\mathbf{P}(\mathbf{X}_t|\mathbf{X}_{t-1})$
  - sensor model  $\mathbf{P}(\mathbf{E}_t|\mathbf{X}_t)$
- Tasks are filtering, smoothing, most likely sequence;  
**all done recursively with constant cost per time step**
- Hidden Markov models have a single discrete state variable; used for speech recognition
- Kalman filters allow  $n$  state variables, linear Gaussian,  $O(n^3)$  update
- Dynamic Bayes nets subsume HMMs, Kalman filters; exact update intractable
- Speech recognition