# STA 4273H: <br> Statistical Machine Learning 

Russ Salakhutdinov
Department of Statistics
rsalakhu@utstat.toronto.edu
http://www.utstat.utoronto.ca/~rsalakhu/
Sidney Smith Hall, Room 6002

## Lecture 9

## Project Reminder

- Brief 5-minute presentations of projects will take place on Monday, March 23. You need to send me 6-7 slides in pdf formtat describing your project.
- Deadline: Sunday March 22, 2015. Submit your slides by e-mail: rsalakhu@cs.toronto.edu
- You should have your name, and project title on the first slide.
- You will have 5-7 mins to briefly describe your project and what you would want to accomplish in this project.
- Brief presentations will be done in an alphabetical order.


## Sequential Data

- So far we focused on problems that assumed that the data points were independent and identically distributed (i.i.d. assumption).
- Express the likelihood function as a product over all data points of the probability distribution evaluated at each data point.
- Poor assumption when working with sequential data.
- For many applications, e.g. financial forecasting, we want to predict the next value in a time series, given past values.
- Intuitively, the recent observations are likely to be more informative in predicting the future.
- Markov models: future predictions are independent of all but the most recent observations.


## Example of a Spectrogram



- Example of a spectrogram of a spoken word 'Bayes theorem':
- Successive observations are highly correlated.
$|b|$ ey $\mid$ z $\mid$ th $\mid$ ih $\mid$ er $\mid$ em $\mid$

| Bayes' |
| :--- | :--- |

## Markov Models

- The simplest model is the first-order Markov chain:

- The joint distribution for a sequence of $N$ observations under this model is:

$$
p\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{N}\right)=p\left(\mathbf{x}_{1}\right) \prod_{n=2}^{N} p\left(\mathbf{x}_{n} \mid \mathbf{x}_{n-1}\right)
$$

- From the d-separation property, the conditionals are given by:

$$
p\left(\mathbf{x}_{n} \mid \mathbf{x}_{1}, \ldots, \mathbf{x}_{n-1}\right)=p\left(\mathbf{x}_{n} \mid \mathbf{x}_{n-1}\right)
$$

- For many applications, these conditional distributions that define the model will be constrained to be equal.
- This corresponds to the assumption of a stationary time series.
- The model is known as homogenous Markov chain.


## Second-Order Markov Models

- We can also consider a second-order Markov chain:

- The joint distribution for a sequence of N observations under this model is:

$$
p\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{N}\right)=p\left(\mathbf{x}_{1}\right) p\left(\mathbf{x}_{2} \mid \mathbf{x}_{1}\right) \prod_{n=3}^{N} p\left(\mathbf{x}_{n} \mid \mathbf{x}_{n-1}, \mathbf{x}_{n-2}\right)
$$

- We can similarly consider extensions to the $\mathrm{M}^{\text {th }}$ order Markov chain.
- Increased flexibility $\rightarrow$ Exponential growth in the number of parameters.
- Markov models need big orders to remember past "events".


## Learning Markov Models

- The ML parameter estimates for a simple Markov model are easy.

Consider a $\mathrm{K}^{\text {th }}$ order model:

$$
p\left(\mathbf{x}_{N}, \ldots, \mathbf{x}_{k+1} \mid \mathbf{x}_{1}, . ., \mathbf{x}_{k}\right)=\prod_{n=k+1}^{N} p\left(\mathbf{x}_{n} \mid \mathbf{x}_{n-1}, \mathbf{x}_{n-2}, . ., \mathbf{x}_{n-k}\right)
$$

- Each window of $\mathrm{k}+1$ outputs is a training case for the model. $p\left(\mathbf{x}_{n} \mid \mathbf{x}_{n-1}, \mathbf{x}_{n-2}, . ., \mathbf{x}_{n-k}\right)$.
- Example: for discrete outputs (symbols) and a 2nd-order Markov model we can use the multinomial model:

$$
p\left(\mathbf{x}_{n}=m \mid \mathbf{x}_{n-1}=a, \mathbf{x}_{n-2}=b\right)=\alpha_{m a b} .
$$

- The maximum likelihood values for $\alpha$ are:

$$
\alpha_{m a b}^{*}=\frac{\operatorname{num}\left[n, \text { s.t. } \mathbf{x}_{n}=m, \mathbf{x}_{n-1}=a, \mathbf{x}_{n-2}=b\right]}{\operatorname{num}\left[n, \text { s.t. } \mathbf{x}_{n-1}=a, \mathbf{x}_{n-2}=b\right]}
$$

## State Space Models

- How about the model that is not limited by the Markov assumption to any order.
- Solution: Introduce additional latent variables!

- Graphical structure known as the State Space Model.
- For each observation $x_{n}$, we have a latent variable $z_{n}$. Assume that latent variables form a Markov chain.
- If the latent variables are discrete $\rightarrow$ Hidden Markov Models (HMMs).

Observed variables can be discrete or continuous.

- If the latent and observed variables are Gaussian $\rightarrow$ Linear Dynamical System.


## State Space Models

- The joint distribution is given by:

$$
p\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{N}, \mathbf{z}_{1}, . ., \mathbf{z}_{N}\right)=p\left(\mathbf{z}_{1}\right) \prod_{n=2}^{N} p\left(\mathbf{z}_{n} \mid \mathbf{z}_{n-1}\right) \prod_{n=1}^{N} p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right)
$$



- Graphical structure known as the State Space Model.
- There is always a path connecting two observed variables $x_{n}, x_{m}$ via latent variables.
- The predictive distribution:

$$
p\left(\mathbf{x}_{n+1} \mid \mathbf{x}_{1}, \ldots, \mathbf{x}_{N}\right)
$$

does not exhibit any conditional independence properties! And so prediction depends on all previous observations.

- Even though hidden state sequence is first-order Markov, the output process is not Markov of any order!


## Hidden Markov Model

- First order Markov chain generates hidden state sequence (known as transition probabilities):

$$
p\left(\mathbf{z}_{n}=k \mid \mathbf{z}_{n-1}=j\right)=A_{j k}, \quad p\left(\mathbf{z}_{1}=k\right)=\pi_{k}
$$

- A set of output probability distributions (one per state) converts state path into sequence of observable symbols/vectors (known as emission probabilities):

$$
p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}, \phi\right)
$$

Gaussian, if $\mathbf{x}$ is continuous.
Conditional probability table if $\mathbf{x}$ is discrete.


State transition
Observation model

$$
p\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{N}, \mathbf{z}_{1}, . ., \mathbf{z}_{N}\right)=p\left(\mathbf{z}_{1}\right) \prod_{n=2}^{N} p\left(\mathbf{z}_{n} \mid \mathbf{z}_{n-1}\right) \prod_{n=1}^{N} p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right)
$$

## Links to Other Models

- You can view HMM as: A Markov chain with stochastic measurements.

- Or a mixture model with states coupled across time.


We will adopt this view, as we worked with
mixture models before.

## Transition Probabilities

- It will be convenient to use 1 -of-K encoding for the latent variables.
- The matrix of transition probabilities takes form:

$$
p\left(z_{n k}=1 \mid z_{n-1, j}=1\right)=A_{j k}, \sum_{k} A_{j k}=1
$$

- The conditionals can be written as:

$$
p\left(\mathbf{z}_{n} \mid \mathbf{z}_{n-1}, A\right)=\prod_{k=1}^{K} \prod_{j=1}^{K} A_{j k}^{z_{n-1, j}, z_{n k}}, \quad p\left(\mathbf{z}_{1} \mid \pi\right)=\prod_{k=1}^{K} \pi_{k}^{z_{1 k}}
$$

- We will focus on homogenous models: all of the conditional distributions over latent variables share the same parameters $\mathbf{A}$.
- Standard mixture model for i.i.d. data: special case in which all parameters $\mathrm{A}_{\mathrm{jk}}$ are the same for all j .
- Or the conditional distribution $p\left(z_{n} \mid z_{n-1}\right)$ is independent of $z_{n-1}$.


## Emission Probabilities

- The emission probabilities take form:

$$
p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}, \boldsymbol{\phi}\right)=\prod_{k=1}^{K} p\left(\mathbf{x}_{n} \mid \boldsymbol{\phi}_{k}\right)^{z_{n k}}
$$

- For example, for a continuous $\mathbf{x}$, we have

$$
p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}, \boldsymbol{\phi}\right)=\prod_{k=1}^{K} \mathcal{N}\left(\mathbf{x}_{n} \mid \boldsymbol{\mu}_{k}, \boldsymbol{\Sigma}_{k}\right)^{z_{n k}}
$$

- For the discrete, multinomial observed variable $\mathbf{x}$, using 1-of-K encoding, the conditional distribution takes form:

$$
p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}, \boldsymbol{\phi}\right)=\prod_{i=1}^{D} \prod_{k=1}^{K} \mu_{i k}^{x_{n i} z_{n k}}
$$

## HMM Model Equations

- The joint distribution over the observed and latent variables is given by:

$$
p(\mathbf{X}, \mathbf{Z} \mid \theta)=p\left(\mathbf{z}_{1} \mid \boldsymbol{\pi}\right) \prod_{n=2}^{N} p\left(\mathbf{z}_{n} \mid \mathbf{z}_{n-1}, A\right) \prod_{n=1}^{N} p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}, \boldsymbol{\phi}\right)
$$

where $\theta=\{\boldsymbol{\pi}, A, \boldsymbol{\phi}\}$ are the model parameters.


- Data are not i.i.d. Everything is coupled across time.
- Three problems: computing probabilities of observed sequences, inference of hidden state sequences, learning of parameters.


## HMM as a Mixture Through Time

- Sampling from a 3-state HMM with a 2-d Gaussian emission model.

- The transition matrix is fixed: $\mathrm{A}_{\mathrm{kk}}=0.9$ and $\mathrm{A}_{\mathrm{jk}}=0.05$.


## Applications of HMMs

- Speech recognition.
- Language modeling
- Motion video analysis/tracking.
- Protein sequence and genetic sequence alignment and analysis.
- Financial time series prediction.


## Maximum Likelihood for the HMM

- We observe a dataset $\mathbf{X}=\left\{\mathrm{x}_{1}, \ldots, \mathrm{x}_{\mathrm{N}}\right\}$.
- The goal is to determine model parameters $\theta=\{\boldsymbol{\pi}, A, \phi\}$.
- The probability of observed sequence takes form:

$$
\begin{aligned}
p(\mathbf{X} \mid \theta) & =\sum_{\mathbf{Z}} p(\mathbf{X}, \mathbf{Z} \mid \theta) \\
p(\text { observed sequence }) & =\sum_{\text {all paths }} p(\text { observed outputs, state paths }) .
\end{aligned}
$$

- In contrast to mixture models, the joint distribution $p(\mathbf{X}, \mathbf{Z} \mid \theta)$ does not factorize over $n$.
- It looks hard: N variables, each of which has K states. Hence $\mathrm{N}^{K}$ total paths.
- Remember inference problem on a simple chain.


## Probability of an Observed Sequence

- The joint distribution factorizes:

$$
\begin{aligned}
p(\mathbf{X} \mid \theta)=\sum_{\mathbf{Z}} p(\mathbf{X}, \mathbf{Z} \mid \theta) & =\sum_{\mathbf{z}_{1}, \ldots, \mathbf{z}_{n}} p\left(\mathbf{z}_{1}, \mathbf{x}_{1}\right) \prod_{n=2}^{N} p\left(\mathbf{z}_{n} \mid \mathbf{z}_{n-1}\right) p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right) \\
& =\sum_{\mathbf{z}_{1}} p\left(\mathbf{z}_{1}\right) p\left(\mathbf{x}_{1} \mid \mathbf{z}_{1}\right) \sum_{\mathbf{z}_{2}} p\left(\mathbf{z}_{2} \mid \mathbf{z}_{1}\right) p\left(\mathbf{x}_{2} \mid \mathbf{z}_{2}\right) \ldots
\end{aligned}
$$



$$
\sum_{\mathbf{z}_{N}} p\left(\mathbf{z}_{N} \mid \mathbf{z}_{N-1}\right) p\left(\mathbf{x}_{N} \mid \mathbf{z}_{N}\right)
$$

- Dynamic Programming: By moving the summations inside, we can save a lot of work.


## EM algorithm

- We cannot perform direct maximization (no closed form solution):

$$
p(\mathbf{X} \mid \theta)=\sum_{\mathbf{Z}} p(\mathbf{X}, \mathbf{Z} \mid \theta)
$$

- EM algorithm: we will derive efficient algorithm for maximizing the likelihood function in HMMs (and later for linear state-space models).
- E-step: Compute the posterior distribution over latent variables:

$$
p\left(\mathbf{Z} \mid \mathbf{X}, \theta^{o l d}\right)
$$

- M-step: Maximize the expected complete data log-likelihood:

$$
Q\left(\theta, \theta^{\text {old }}\right)=\sum_{\mathbf{Z}} p\left(\mathbf{Z} \mid \mathbf{X}, \theta^{\text {old }}\right) \log p(\mathbf{X}, \mathbf{Z} \mid \theta)
$$

- If we knew the true state path, then ML parameter estimation would be trivial.
- We will first look at the E-step: Computing the true posterior distribution over the state paths.


## Inference of Hidden States

- We want to estimate the hidden states given observations. To start with, let us estimate a single hidden state:

$$
\gamma\left(\mathbf{z}_{n}\right)=p\left(\mathbf{z}_{n} \mid \mathbf{X}\right)=\frac{p\left(\mathbf{X} \mid \mathbf{z}_{n}\right) p\left(\mathbf{z}_{n}\right)}{p(\mathbf{X})}
$$

- Using conditional independence property, we obtain:

$$
\begin{aligned}
p\left(\mathbf{z}_{n} \mid \mathbf{X}\right) & =\frac{p\left(\mathbf{x}_{1}, . ., \mathbf{x}_{n} \mid \mathbf{z}_{n}\right) p\left(\mathbf{x}_{n+1}, \ldots, \mathbf{x}_{N} \mid \mathbf{z}_{n}\right) p\left(\mathbf{z}_{n}\right)}{p(\mathbf{X})} \\
& =\frac{p\left(\mathbf{x}_{1}, ., \mathbf{x}_{n}, \mathbf{z}_{n}\right) p\left(\mathbf{x}_{n+1}, \ldots, \mathbf{x}_{N} \mid \mathbf{z}_{n}\right)}{p(\mathbf{X})}=\frac{\alpha\left(\mathbf{z}_{n}\right) \beta\left(\mathbf{z}_{n}\right)}{p(\mathbf{X})}
\end{aligned}
$$



## Inference of Hidden States

- Hence:

$$
\begin{array}{cl}
\gamma\left(\mathbf{z}_{n}\right)=\frac{p\left(\mathbf{x}_{1}, . ., \mathbf{x}_{n}, \mathbf{z}_{n}\right) p\left(\mathbf{x}_{n+1}, \ldots, \mathbf{x}_{N} \mid \mathbf{z}_{n}\right)}{p(\mathbf{X})}=\frac{\alpha\left(\mathbf{z}_{n}\right) \beta\left(\mathbf{z}_{n}\right)}{p(\mathbf{X})} . \\
\alpha\left(\mathbf{z}_{n}\right) \equiv p\left(\mathbf{x}_{1}, . ., \mathbf{x}_{n}, \mathbf{z}_{n}\right) \quad \begin{array}{l}
\text { The joint probability of observing } \\
\text { of the data up to time } \mathrm{n} \text { and } \mathbf{z}_{\mathrm{n}} .
\end{array} \\
\beta\left(\mathbf{z}_{n}\right) \equiv p\left(\mathbf{x}_{n+1}, . ., \mathbf{x}_{N} \mid \mathbf{z}_{n}\right) . \longleftarrow & \begin{array}{l}
\text { The conditional probability of all } \\
\text { future data from time } \mathrm{n}+1 \text { to } \mathrm{N} .
\end{array}
\end{array}
$$

- Each $\alpha\left(z_{n}\right)$ and $\beta\left(\mathbf{z}_{n}\right)$ represent a set of $K$ numbers, one for each of the possible settings of the 1 -of-K binary vector $z_{n}$.
- We will derive efficient recursive algorithm, known as the alpha-beta recursion, or forward-backward algorithm.
- Relates to the sum-product message passing algorithm for tree-structured graphical models.


## The Forward ( $\alpha$ ) Recursion

- The forward recursion:

$$
\begin{aligned}
\alpha\left(\mathbf{z}_{n}\right) & =p\left(\mathbf{x}_{1}, . ., \mathbf{x}_{n}, \mathbf{z}_{n}\right) \\
& =p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right) p\left(\mathbf{x}_{1}, . ., \mathbf{x}_{n-1}, \mathbf{z}_{n}\right) \\
& =p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right) \sum_{\mathbf{z}_{n-1}} p\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{n-1}, \mathbf{z}_{n-1}, \mathbf{z}_{n}\right) \\
& =p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right) \sum_{\mathbf{z}_{n-1}} p\left(\mathbf{x}_{1}, . ., \mathbf{x}_{n-1}, \mathbf{z}_{n-1}\right) p\left(\mathbf{z}_{n} \mid \mathbf{z}_{n-1}\right) \\
& =p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right) \sum_{\mathbf{z}_{n-1}} \alpha\left(\mathbf{z}_{n-1}\right) p\left(\mathbf{z}_{n} \mid \mathbf{z}_{n-1}\right)
\end{aligned}
$$



- Observe:

$$
p(\mathbf{X})=\sum_{\mathbf{z}_{N}} \alpha\left(\mathbf{z}_{N}\right)
$$

- This enables us to easily (cheaply) compute the desired likelihood.


## The Forward ( $\alpha$ ) Recursion

- The forward recursion:


Exponentially many paths.


At each node, sum up the values of all incoming paths.

- This is exactly dynamic programming.


## The Forward ( $\alpha$ ) Recursion

- Illustration of the forward recursion


Here $\alpha\left(z_{n, 1}\right)$ is obtained by

- Taking the elements $\alpha\left(z_{n-1}, j\right)$
- Summing the up with weights $\mathrm{A}_{\mathrm{j} 1}$, corresponding to $p\left(z_{n} \mid z_{n-1}\right)$
- Multiplying by the data contribution $p\left(x_{n} \mid z_{n 1}\right)$.

$$
\alpha\left(\mathbf{z}_{n}\right)=p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right) \sum_{\mathbf{z}_{n-1}} \alpha\left(\mathbf{z}_{n-1}\right) p\left(\mathbf{z}_{n} \mid \mathbf{z}_{n-1}\right)
$$

- The initial condition is given by:

$$
\alpha\left(\mathbf{z}_{1}\right)=p\left(\mathbf{x}_{1} \mid \mathbf{z}_{1}\right) p\left(\mathbf{z}_{1}\right)=\prod_{k=1}^{K}\left[\pi_{k} p\left(\mathbf{x}_{1} \mid \phi_{k}\right)\right]^{z_{1 k}}
$$

## The Backward ( $\beta$ ) Recursion

- There is also a simple recursion for $\beta\left(\mathrm{z}_{\mathrm{n}}\right)$ :

$$
\begin{aligned}
& \beta\left(\mathbf{z}_{n}\right)=p\left(\mathbf{x}_{n+1}, . ., \mathbf{x}_{N} \mid \mathbf{z}_{n}\right) \\
&=\sum_{\mathbf{z}_{n+1}} p\left(\mathbf{x}_{n+1}, \ldots, \mathbf{x}_{N}, \mathbf{z}_{n+1} \mid \mathbf{z}_{n}\right) \\
&=\sum_{\mathbf{z}_{n+1}} p\left(\mathbf{x}_{n+1}, \ldots, \mathbf{x}_{N} \mid \mathbf{z}_{n+1}, \mathbf{z}_{n}\right) p\left(\mathbf{z}_{n+1} \mid \mathbf{z}_{n}\right) \\
&=\sum_{\mathbf{z}_{n+1}} p\left(\mathbf{x}_{n+1}, \ldots, \mathbf{x}_{N} \mid \mathbf{z}_{n+1}\right) p\left(\mathbf{z}_{n+1} \mid \mathbf{z}_{n}\right) \\
&=\sum_{\mathbf{z}_{n+1}} p\left(\mathbf{x}_{n+2}, \ldots, \mathbf{x}_{N} \mid \mathbf{z}_{n+1}\right) p\left(\mathbf{x}_{n+1} \mid \mathbf{z}_{n+1}\right) p\left(\mathbf{z}_{n+1} \mid \mathbf{z}_{n}\right) \\
&=\sum_{\mathbf{z}_{n+1}} \beta\left(\mathbf{z}_{n+1}\right) p\left(\mathbf{x}_{n+1} \mid \mathbf{z}_{n+1}\right) p\left(\mathbf{z}_{n+1} \mid \mathbf{z}_{n}\right) \\
& \mathbf{z}_{n-1}
\end{aligned}
$$



## The Backward ( $\beta$ ) Recursion

- Illustration of the backward recursion



## The Backward ( $\beta$ ) Recursion

- $\alpha\left(\mathrm{z}_{\mathrm{nk}}\right)$ gives total inflow of probability to node ( $\mathrm{n}, \mathrm{k}$ ).
- $\beta\left(z_{n k}\right)$ gives total outflow of probability.

- In fact, we can do one forward pass to compute all the $\alpha\left(z_{n}\right)$ and one backward pass to compute all the $\beta\left(z_{n}\right)$ and then compute any $\gamma\left(z_{n}\right)$ we want. Total cost is $\mathrm{O}\left(\mathrm{K}^{2} \mathrm{~N}\right)$.


## Computing Likelihood

- Note that

$$
\sum_{\mathbf{z}_{n}} \gamma\left(\mathbf{z}_{n}\right)=\sum_{\mathbf{z}_{n}} p\left(\mathbf{z}_{n} \mid \mathbf{X}\right)=1
$$

- We can compute the likelihood at any time using $\alpha-\beta$ recursion:

$$
p(\mathbf{X} \mid \theta)=\sum_{\mathbf{z}_{n}} \alpha\left(\mathbf{z}_{n}\right) \beta\left(\mathbf{z}_{n}\right)
$$

- In the forward calculation we proposed originally, we did this at the final time step $\mathrm{n}=\mathrm{N}$.

Because $\beta\left(z_{n}\right)=1$.

$$
p(\mathbf{X} \mid \theta)=\sum_{\mathbf{z}_{N}} \alpha\left(\mathbf{z}_{N}\right)
$$

- This is a good way to check your code!


## Two-Frame Inference

- We will also need the cross-time statistics for adjacent time steps:

$$
\begin{aligned}
\xi\left(\mathbf{z}_{n-1}, \mathbf{z}_{n}\right) & =p\left(\mathbf{z}_{n-1}, \mathbf{z}_{n} \mid \mathbf{X}\right) \\
& =\frac{p\left(\mathbf{X} \mid \mathbf{z}_{n-1}, \mathbf{z}_{n}\right) p\left(\mathbf{z}_{n-1}, \mathbf{z}_{n}\right)}{p(\mathbf{X})} \\
& =\frac{p\left(\mathbf{x}_{1}, . ., \mathbf{x}_{n-1} \mid \mathbf{z}_{n-1}\right) p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right) p\left(\mathbf{x}_{n+1}, \ldots, \mathbf{x}_{N} \mid \mathbf{z}_{n}\right) p\left(\mathbf{z}_{n} \mid \mathbf{z}_{n-1}\right) p\left(\mathbf{z}_{n-1}\right)}{p(\mathbf{X})} \\
& =\frac{\alpha\left(\mathbf{z}_{n-1}\right) p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right) p\left(\mathbf{z}_{n} \mid \mathbf{z}_{n-1}\right) \beta\left(\mathbf{z}_{n}\right)}{p(\mathbf{X})}
\end{aligned}
$$

- This is a $\mathrm{K} \times \mathrm{K}$ matrix with elements $\xi(\mathrm{i}, \mathrm{j})$ representing the expected number of transitions from state $i$ to state $j$ that begin at time $n-1$, given all the observations.

- It can be computed with the same $\alpha$ and $\beta$ recursions.


## EM algorithm

- Intuition: if only we knew the true state path then ML parameter estimation would be trivial.
- E-step: Compute the posterior distribution over the state path using $\alpha-\beta$ recursion (dynamic programming):

$$
p\left(\mathbf{Z} \mid \mathbf{X}, \theta^{o l d}\right)
$$

- M-step: Maximize the expected complete data log-likelihood (parameter reestimation):

$$
Q\left(\theta, \theta^{\text {old }}\right)=\sum_{\mathbf{Z}} p\left(\mathbf{Z} \mid \mathbf{X}, \theta^{\text {old }}\right) \log p(\mathbf{X}, \mathbf{Z} \mid \theta) .
$$

- We then iterate. This is also known as a Baum-Welch algorithm (special case of EM).



## Complete Data Log-likelihood

- Complete data log-likelihood takes form:

$$
\begin{aligned}
& \log p(\mathbf{X}, \mathbf{Z} \mid \theta)=\log \left[p\left(\mathbf{z}_{1} \mid \boldsymbol{\pi}\right) \prod_{n=2}^{N} p\left(\mathbf{z}_{n} \mid \mathbf{z}_{n-1}, A\right) \prod_{n=1}^{N} p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}, \boldsymbol{\phi}\right)\right] \\
&= \log \left[\prod_{k=1}^{K} \pi_{k}^{z_{1 k}} \prod_{n=2}^{N} \prod_{k=1}^{K} \prod_{j=1}^{K} A_{j k}^{z_{n-1, j}, z_{n k}} \prod_{n=1}^{N} \prod_{k=1}^{K} p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right)^{z_{n k}}\right] \\
&=\sum_{k=1}^{K} z_{1 k} \log \pi_{k}+\sum_{n=2}^{N} \sum_{k=1}^{K} \sum_{j=1}^{K}\left[z_{n-1, j} z_{n k}\right] \log A_{j k}+\sum_{n=1}^{N} \sum_{k=1}^{K} z_{n k} \log p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right) .
\end{aligned}
$$

- Statistics we need from the E-step are:

$$
\begin{aligned}
& \gamma\left(\mathbf{z}_{n}\right)=p\left(\mathbf{z}_{n} \mid \mathbf{X}\right) \\
& \xi\left(\mathbf{z}_{n-1}, \mathbf{z}_{n}\right)=p\left(\mathbf{z}_{n-1}, \mathbf{z}_{n} \mid \mathbf{X}\right)
\end{aligned}
$$

## Expected Complete Data Log-likelihood

- The complete data log-likelihood takes form:

$$
\begin{aligned}
& Q\left(\theta, \theta^{o l d}\right)=\sum_{\mathbf{Z}} p\left(\mathbf{Z} \mid \mathbf{X}, \theta^{o l d}\right) \log p(\mathbf{X}, \mathbf{Z} \mid \theta) . \\
& \quad=\sum_{k=1}^{K} \gamma\left(z_{1 k}\right) \log \pi_{k}+\sum_{n=2}^{N} \sum_{k=1}^{K} \sum_{j=1}^{K} \xi\left(z_{n-1, j} z_{n k}\right) \log A_{j k}+\sum_{n=1}^{N} \sum_{k=1}^{K} \gamma_{n k} \log p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right) .
\end{aligned}
$$

- Hence in the E-step we evaluate:

$$
\begin{aligned}
& \gamma\left(\mathbf{z}_{n}\right)=p\left(\mathbf{z}_{n} \mid \mathbf{X}\right) \\
& \xi\left(\mathbf{z}_{n-1}, \mathbf{z}_{n}\right)=p\left(\mathbf{z}_{n-1}, \mathbf{z}_{n} \mid \mathbf{X}\right)
\end{aligned}
$$

- In the M -step we optimize Q with respect to parameters: $\theta=\{\pi, A, \phi\}$.


## Parameter Estimation

- Initial state distribution: expected number of times in state k at time 1 :

$$
\pi_{k}^{n e w}=\frac{\gamma\left(z_{1 k}\right)}{\sum_{j=1}^{K} \gamma\left(z_{1 j}\right)}
$$

- Expected \# of transitions from state j to k which begin at time $\mathrm{n}-1$ :

$$
\xi\left(\mathbf{z}_{n-1, j}, \mathbf{z}_{n, k}\right)=p\left(\mathbf{z}_{n-1, j}, \mathbf{z}_{n, k} \mid \mathbf{X}\right)
$$

so the estimated transition probabilities are:

$$
A_{j k}^{n e w}=\frac{\sum_{n=2}^{N} \xi\left(z_{n-1, j}, z_{n k}\right)}{\sum_{l=1}^{K} \sum_{n=2}^{N} \xi\left(z_{n-1, j}, z_{n l}\right)} .
$$

- The EM algorithm must be initialized by choosing starting values for $\pi$ and A.
- Note that any elements of $\pi$ or $\mathbf{A}$ that initially are set to zero will remain zero in subsequent EM updates.


## Parameter Estimation: Emission Model

- For the case of discrete multinomial observed variables, the observation model takes form:

$$
p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}, \boldsymbol{\phi}\right)=\prod_{i=1}^{D} \prod_{k=1}^{K} \mu_{i k}^{x_{n i} z_{n k}}
$$

Same as fitting Bernoulli mixture model.


- And the corresponding M-step update: $\mu_{i k}^{n e w}=\frac{\sum_{n=1}^{N} \gamma\left(z_{n k}\right) x_{n i}}{\sum_{n=1}^{N} \gamma\left(z_{n k}\right)}$.
- For the case of the Gaussian emission model:

$$
p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}, \boldsymbol{\phi}\right)=\prod_{k=1}^{K} \mathcal{N}\left(\mathbf{x}_{n} \mid \boldsymbol{\mu}_{k}, \boldsymbol{\Sigma}_{k}\right)^{z_{n k}}
$$

## Remember:

$$
\gamma\left(\mathbf{z}_{n}\right)=p\left(\mathbf{z}_{n} \mid \mathbf{X}\right)
$$

- And the corresponding M-step updates:

$$
\begin{aligned}
\boldsymbol{\mu}_{k}^{n e w} & =\frac{1}{N_{k}} \sum_{n} \gamma\left(z_{n k}\right) \mathbf{x}_{n}, \quad N_{k}=\sum_{n} \gamma\left(z_{n k}\right) \\
\boldsymbol{\Sigma}_{k}^{n e w} & =\frac{1}{N_{k}} \sum_{n=1}^{N} \gamma\left(y_{n k}\right)\left(\mathbf{x}_{n}-\boldsymbol{\mu}_{k}\right)\left(\mathbf{x}_{n}-\boldsymbol{\mu}_{k}\right)^{T}
\end{aligned}
$$

Same as fitting a Gaussian mixture model.

## Viterbi Decoding

- The numbers $\gamma\left(z_{n}\right)$ above gave the probability distribution over all states at any time.
- By choosing the state $\gamma^{*}\left(Z_{n}\right)$ with the largest probability at each time, we can make an "average" state path. This is the path with the maximum expected number of correct states.
- To find the single best path, we do Viterbi decoding which is Bellman's dynamic programming algorithm applied to this problem.
- The recursions look the same, except with max instead of $\sum$.
- Same dynamic programming trick: instead of summing, we keep the term with the highest value at each node.
- There is also a modified EM (Baum-Welch) training based on the Viterbi decoding. Like K-means instead of mixtures of Gaussians.
- Relates to the max-sum algorithm for tree structured graphical models.


## Viterbi Decoding

- A fragment of the HMM lattice showing two possible paths:

- Viterbi decoding efficiently determines the most probable path from the exponentially many possibilities.
- The probability of each path is given by the product of the elements of the transition matrix $\mathrm{A}_{\mathrm{jk}}$, along with the emission probabilities associated with each node in the path.


## Using HMMs for Recognition

- We can use HMMs for recognition by:
- training one HMM for each class (requires labeled training data)
- evaluating probability of an unknown sequence under each HMM
- classifying unknown sequence by choosing an HMM with highest likelihood

- This requires the solution of two problems:
- Given model, evaluate probability of a sequence. (We can do this exactly and efficiently.)
- Given some training sequences, estimate model parameters. (We can find the local maximum using EM.)


## Autoregressive HMMs

- One limitation of the standard HMM is that it is poor at capturing longrange correlations between observations, as these have to be mediated via the first order Markov chain of hidden states.

- Autoregressive HMM: The distribution over $\mathrm{x}_{\mathrm{n}}$ depends on a subset of previous observations.
- The number of additional links must be limited to avoid an excessive number of free parameters.
- The graphical model framework motivates a number of different models based on HMMs.


## Input-Output HMMs

- Both the emission probabilities and the transition probabilities depend on the values of a sequence of observations $u_{1}, \ldots, u_{N}$.

- Model parameters can be efficiently fit using EM, in which the E-step involves forward-backward recursion.


## Factorial HMMs

- Example of Factorial HMM comprising of two Markov chains of latent variables:

- Motivation: In order to represent 10 bits of information at a given time step, a standard HMM would need $K=2^{10}=1024$ states.
- Factorial HMMs would use 10 binary chains.
- Much more powerful model.
- The key disadvantage: Exact inference is intractable.
- Observing the $\mathbf{x}$ variables introduces dependencies between latent chains.
- Hence E-step for this model does not correspond to running forwardbackward along the M latent chain independently.


## Factorial HMMs

- The conditional independence property: $\mathrm{z}_{\mathrm{n}+1} \perp \mathrm{z}_{\mathrm{n}-1} \mid \mathrm{z}_{\mathrm{n}}$ does not hold for the individual latent chains.

- There is no efficient exact E-step for this model.
- One solution would be to use MCMC techniques to obtain approximate sample from the posterior.
- Another alternative is to resort to variational inference.
- The variational distribution can be described by M separate Markov chains corresponding to the latent chains in the original model (structured meanfield approximation).


## Regularizing HMMs

- There are two problems:
- for high dimensional outputs, lots of parameters in the emission model
- with many states, transition matrix has many (squared) elements
- First problem: full covariance matrices in high dimensions or discrete symbol models with many symbols have lots of parameters. To estimate these accurately requires a lot of training data.

- We can use mixtures of diagonal covariance Gaussians.
- For discrete data, we can use mixtures of base rates.
- We can also tie parameters across states.


## Regularizing Transition Matrices

- One way to regularize large transition matrices is to constrain them to be sparse: instead of being allowed to transition to any other state, each state has only a few possible successor states.
- A very effective way to constrain the transitions is to order the states in the HMM and allow transitions only to states that come later in the ordering.
- Such models are known as "linear HMMs", "chain HMMs" or "left- to-right HMMs". Transition matrix is upper- diagonal (usually only has a few bands).



## Linear Dynamical Systems

- In HMMs, latent variables are discrete but with arbitrary emission probability distributions.

- We now consider a linear-Gaussian state-space model, so that latent variables and observed variables are multivariate Gaussian distributions.
- An HMM can be viewed as an extension of the mixture models to allow for sequential correlations in the data
- Similarly, the linear dynamical system (LDS) can be viewed as a generalization of the continuous latent variable models, such as probabilistic PCA.


## Linear Dynamical Systems

- The model is represented by a tree-structured directed graph, so inference can be solved efficiently using the sum-product algorithm.
- The forward recursions, analogous to the $\alpha$-messages of HMMs are known as the Kalman filter equations.
- The backward recursions, analogous to the $\beta$-messages, are known as the Kalman smoother equations.
- The Kalman filter is used in many real-time tracking applications.
- Because the LDS is a linear-Gaussian model, the joint distribution over all variables, as well as marginals and conditionals, will be Gaussian.
- This leads to tractable inference and learning.


## The Model

- We can write the transition and emission distributions in the general form:

- These can be expressed in terms of noisy linear equations:

$$
\begin{aligned}
\mathbf{z}_{n}=\mathbf{A} \mathbf{z}_{n-1}+\mathbf{w}_{n}, & \mathbf{w} \sim \mathcal{N}(\mathbf{w} \mid \mathbf{0}, \boldsymbol{\Gamma}) \\
\mathbf{x}_{n}=\mathbf{C} \mathbf{z}_{n}+\mathbf{v}_{n}, & \mathbf{v} \sim \mathcal{N}(\mathbf{v} \mid \mathbf{0}, \boldsymbol{\Sigma}), \\
\mathbf{z}_{1}=\boldsymbol{\mu}_{0}+\mathbf{u}, & \mathbf{u} \sim \mathcal{N}\left(\mathbf{u} \mid \mathbf{0}, \mathbf{V}_{0}\right)
\end{aligned}
$$

- Model parameters $\theta=\left\{\mathbf{A}, \boldsymbol{\Gamma}, \mathbf{C}, \boldsymbol{\Sigma}, \boldsymbol{\mu}_{0}, \mathbf{V}_{0}\right\}$ can be learned using EM algorithm (similar to standard HMM case).


## Inference in LDS

- Consider forward equations. The initial message is Gaussian, and since each of the factors is Gaussian, all subsequent messages will also be Gaussians.


$$
\begin{aligned}
\mathbf{z}_{n}=\mathbf{A} \mathbf{z}_{n-1}+\mathbf{w}_{n}, & \mathbf{w} \sim \mathcal{N}(\mathbf{w} \mid \mathbf{0}, \boldsymbol{\Gamma}) \\
\mathbf{x}_{n}=\mathbf{C} \mathbf{z}_{n}+\mathbf{v}_{n}, & \mathbf{v} \sim \mathcal{N}(\mathbf{v} \mid \mathbf{0}, \boldsymbol{\Sigma}) \\
\mathbf{z}_{1}=\boldsymbol{\mu}_{0}+\mathbf{u}, & \mathbf{u} \sim \mathcal{N}\left(\mathbf{u} \mid \mathbf{0}, \mathbf{V}_{0}\right)
\end{aligned}
$$

- Similar to HMMs, let us define the normalized version of $\alpha\left(z_{n}\right)$ :

$$
\hat{\alpha}\left(\mathbf{z}_{n}\right)=p\left(\mathbf{z}_{n} \mid \mathbf{x}_{1}, \ldots, \mathbf{x}_{n}\right)=\frac{\alpha\left(\mathbf{z}_{n}\right)}{p\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{n}\right)}=\mathcal{N}\left(\mathbf{z}_{n} \mid \boldsymbol{\mu}_{n}, \mathbf{V}_{n}\right)
$$

- Using forward recursion, we get:
Remember: for HMMs

$$
\alpha\left(\mathbf{z}_{n}\right)=p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right) \sum_{\mathbf{z}_{n-1}} \alpha\left(\mathbf{z}_{n-1}\right) p\left(\mathbf{z}_{n} \mid \mathbf{z}_{n-1}\right)
$$

$$
c_{n} \hat{\alpha}\left(\mathbf{z}_{n}\right)=p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right) \int \hat{\alpha}\left(\mathbf{z}_{n-1}\right) p\left(\mathbf{z}_{n} \mid \mathbf{z}_{n-1}\right) d \mathbf{z}_{n-1}
$$

$$
=\mathcal{N}\left(\mathbf{x}_{n} \mid \mathbf{C} \mathbf{z}_{n}, \boldsymbol{\Sigma}\right) \int \mathcal{N}\left(\mathbf{z}_{n-1} \mid \boldsymbol{\mu}_{n-1}, \mathbf{V}_{n-1}\right) \mathcal{N}\left(\mathbf{z}_{n} \mid \mathbf{A} \mathbf{z}_{n-1}, \boldsymbol{\Gamma}\right) d \mathbf{z}_{n-1}
$$

## Inference in LDS

- Hence we obtain:

$$
\begin{aligned}
c_{n} \hat{\alpha}\left(\mathbf{z}_{n}\right) & =p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right) \int \hat{\alpha}\left(\mathbf{z}_{n-1}\right) p\left(\mathbf{z}_{n} \mid \mathbf{z}_{n-1}\right) d \mathbf{z}_{n-1} \\
& =\mathcal{N}\left(\mathbf{x}_{n} \mid \mathbf{C} \mathbf{z}_{n}, \mathbf{\Sigma}\right) \int \mathcal{N}\left(\mathbf{z}_{n-1} \mid \boldsymbol{\mu}_{n-1}, \mathbf{V}_{n-1}\right) \mathcal{N}\left(\mathbf{z}_{n} \mid \mathbf{A} \mathbf{z}_{n-1}, \boldsymbol{\Gamma}\right) d \mathbf{z}_{n-1} .
\end{aligned}
$$

in which case $\alpha\left(Z_{n}\right)$ is Gaussian:

$$
\begin{aligned}
c_{n} \hat{\alpha}\left(\mathbf{z}_{n}\right) & =\mathcal{N}\left(\mathbf{z}_{n} \mid \boldsymbol{\mu}_{n}, \mathbf{V}_{n}\right) . \\
\boldsymbol{\mu}_{n} & =\mathbf{A} \boldsymbol{\mu}_{n-1}+\mathbf{K}_{n}\left(\mathbf{x}_{n}-\mathbf{C A} \boldsymbol{\mu}_{n-1}\right) \\
\mathbf{V}_{n} & =\left(\mathbf{I}-\mathbf{K}_{n} \mathbf{C}\right) \mathbf{P}_{n-1} \\
\mathbf{P}_{n-1} & =\mathbf{A} \mathbf{V}_{n-1} \mathbf{A}^{T}+\boldsymbol{\Gamma} .
\end{aligned}
$$

and we have also defined the Kalman gain matrix:

$$
\mathbf{K}_{n}=\mathbf{P}_{n-1} \mathbf{C}^{T}\left(\mathbf{C} \mathbf{P}_{n-1} \mathbf{C}^{T}+\boldsymbol{\Sigma}\right)^{-1}
$$

## Kalman Filter

- Let us examine the evolution of the mean:




Predicted mean plus the correction term controlled by the Kalman gain matrix.

Error between the predicted observation $x_{n}$ and the actual observation $x_{n}$.

- We can view the Kalman filter as a process of making subsequent predictions and then correcting these predictions in the light of the new observations.


## Kalman Filter

- Example:

$p\left(\mathbf{z}_{n-1} \mid \mathbf{x}_{1}, \ldots, \mathbf{x}_{n-1}\right)$
blue curve

$p\left(\mathbf{z}_{n} \mid \mathbf{x}_{1}, \ldots, \mathbf{x}_{n-1}\right)$
red curve
incorporate transition model


$$
p\left(\mathbf{z}_{n} \mid \mathbf{x}_{1}, \ldots, \mathbf{x}_{n}\right)
$$

blue curve
incorporate new observation (density of the new point is given by the green curve)

- The new observation has shifted and narrowed the distribution compared to (see red curve) $p\left(\mathbf{z}_{n} \mid \mathbf{x}_{1}, \ldots, \mathbf{x}_{n-1}\right)$.


## Tracking Example

- LDS that is being used to track a moving object in 2-D space:

- Blue points indicate the true position of the object.
- Green points denote the noisy measurements.
- Red crosses indicate the means of the inferred posterior distribution of the positions inferred by the Kalman filter.


## Particle Filters

- For dynamical systems that are non-Gaussian (e.g. emission densities are nonGaussian), we can use sampling methods to find a tractable solution to the inference problem.
- Consider a class of distributions represented by the graphical model:

- Suppose we observed $X_{n}=$ $\left\{x_{1}, \ldots, x_{n}\right\}$, and we wish to approximate:

$$
\begin{aligned}
\mathbb{E}\left[f\left(\mathbf{z}_{n}\right)\right] & =\int f\left(\mathbf{z}_{n}\right) p\left(\mathbf{z}_{n} \mid \mathbf{X}_{n}\right) d \mathbf{z}_{n} \\
& =\int f\left(\mathbf{z}_{n}\right) p\left(\mathbf{z}_{n} \mid \mathbf{x}_{n}, \mathbf{X}_{n-1}\right) d \mathbf{z}_{n} \\
& =\frac{\int f\left(\mathbf{z}_{n}\right) p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right) p\left(\mathbf{z}_{n} \mid \mathbf{X}_{n-1}\right) d \mathbf{z}_{n}}{\int p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}\right) p\left(\mathbf{z}_{n} \mid \mathbf{X}_{n-1}\right) d \mathbf{z}_{n}} \approx \sum_{l=1}^{L} w_{n}^{(l)} f\left(\mathbf{z}_{n}^{(l)}\right)
\end{aligned}
$$

## Particle Filters

- Hence

$$
\mathbb{E}\left[f\left(\mathbf{z}_{n}\right)\right]=\int f\left(\mathbf{z}_{n}\right) p\left(\mathbf{z}_{n} \mid \mathbf{X}_{n}\right) d \mathbf{z}_{n} \approx \sum_{l=1}^{L} w_{n}^{(l)} f\left(\mathbf{z}_{n}^{(l)}\right), \quad \mathbf{z}_{n}^{(l)} \sim p\left(\mathbf{z}_{n} \mid \mathbf{X}_{n-1}\right)
$$

with importance weights:

$$
w_{n}^{(l)}=\frac{p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}^{(l)}\right)}{\sum_{m=1}^{L} p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}^{(m)}\right)}
$$

- Hence the posterior $p\left(z_{n} \mid X_{n}\right)$ is represented by the set of $L$ samples together with the corresponding importance weights.
- We would like to define a sequential algorithm.
- Suppose that a set of samples and weights have been obtained at time step $n$.
- We wish to find the set of new samples and weights at time step $n+1$.


## Particle Filters

- From our previous result, let $f\left(\mathbf{z}_{n}\right)=p\left(\mathbf{z}_{n+1} \mid \mathbf{z}_{n}\right)$,

$$
\begin{aligned}
& \qquad \begin{array}{ll}
p\left(\mathbf{z}_{n+1} \mid \mathbf{X}_{n}\right) & =\int p\left(\mathbf{z}_{n+1} \mid \mathbf{z}_{n}\right) p\left(\mathbf{z}_{n} \mid \mathbf{X}_{n}\right) d \mathbf{z}_{n} \\
& \approx \sum_{l=1}^{L} w_{n}^{(l)} p\left(\mathbf{z}_{n+1} \mid \mathbf{z}_{n}^{(l)}\right), \quad \mathbf{z}_{n}^{(l)} \sim p\left(\mathbf{z}_{n} \mid \mathbf{X}_{n-1}\right), \\
& w_{n}^{(l)}=\frac{p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}^{(l)}\right)}{\sum_{m=1}^{L} p\left(\mathbf{x}_{n} \mid \mathbf{z}_{n}^{(m)}\right)} .
\end{array} \\
& \text { - Summary of the particle filter algorithm: }
\end{aligned}
$$

- At time n , we have a sample representation of the posterior distribution $p\left(z_{n} \mid X_{n}\right)$ expressed as $L$ samples with corresponding weights.
- We next draw $L$ samples from the mixture distribution (above).
- For each sample, we then use the new observation to re-evaluate the weights:

$$
w_{n+1}^{(l)}=\frac{p\left(\mathbf{x}_{n+1} \mid \mathbf{z}_{n+1}^{(l)}\right)}{\sum_{m=1}^{L} p\left(\mathbf{x}_{n+1} \mid \mathbf{z}_{n+1}^{(m)}\right)}
$$

## Example



- At time $n$, the posterior $p\left(z_{n} \mid X_{n}\right)$ is represented as a mixture distribution.
- We draw a set of $L$ samples from this distribution (incorporating the transition model).
- The new weights evaluated by incorporating the new observation $x_{n+1}$.

