Probabilistic & Unsupervised Learning

Latent Variable Models for Time Series

Maneesh Sahani

maneesh@gatsby.ucl.ac.uk

Gatsby Computational Neuroscience Unit, and MSc ML/CSML, Dept Computer Science University College London

Term 1, Autumn 2016

Thus far, our data have been (marginally) iid. Now consider a sequence of observations:

 $\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \dots, \mathbf{X}_t$

that are not independent.

Thus far, our data have been (marginally) iid. Now consider a sequence of observations:

 $\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \dots, \mathbf{X}_t$

that are *not* independent.

Examples:

- Sequence of images
- Stock prices
- Recorded speech or music, English sentences
- Kinematic variables in a robot
- Sensor readings from an industrial process
- Amino acids, DNA, etc. . .

Thus far, our data have been (marginally) iid. Now consider a sequence of observations:

 $\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \dots, \mathbf{X}_t$

that are *not* independent.

Examples:

- Sequence of images
- Stock prices
- Recorded speech or music, English sentences
- Kinematic variables in a robot
- Sensor readings from an industrial process
- Amino acids, DNA, etc. . .

Goal: To build a joint probabilistic model of the data $p(\mathbf{x}_1, \ldots, \mathbf{x}_t)$.

Thus far, our data have been (marginally) iid. Now consider a sequence of observations:

 $\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \dots, \mathbf{X}_t$

that are *not* independent.

Examples:

- Sequence of images
- Stock prices
- Recorded speech or music, English sentences
- Kinematic variables in a robot
- Sensor readings from an industrial process
- Amino acids, DNA, etc. . .

Goal: To build a joint probabilistic model of the data $p(\mathbf{x}_1, \ldots, \mathbf{x}_t)$.

• Predict $p(\mathbf{x}_t | \mathbf{x}_1, \dots, \mathbf{x}_{t-1})$

Thus far, our data have been (marginally) iid. Now consider a sequence of observations:

 $\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \dots, \mathbf{X}_t$

that are *not* independent.

Examples:

- Sequence of images
- Stock prices
- Recorded speech or music, English sentences
- Kinematic variables in a robot
- Sensor readings from an industrial process
- Amino acids, DNA, etc...

Goal: To build a joint probabilistic model of the data $p(\mathbf{x}_1, \ldots, \mathbf{x}_t)$.

- Predict $p(\mathbf{x}_t | \mathbf{x}_1, \dots, \mathbf{x}_{t-1})$
- ▶ Detect abnormal/changed behaviour (if $p(\mathbf{x}_t, \mathbf{x}_{t+1}, ... | \mathbf{x}_1, ..., \mathbf{x}_{t-1})$ small)

Thus far, our data have been (marginally) iid. Now consider a sequence of observations:

 $\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \dots, \mathbf{X}_t$

that are *not* independent.

Examples:

- Sequence of images
- Stock prices
- Recorded speech or music, English sentences
- Kinematic variables in a robot
- Sensor readings from an industrial process
- Amino acids, DNA, etc...

Goal: To build a joint probabilistic model of the data $p(\mathbf{x}_1, \ldots, \mathbf{x}_t)$.

- Predict $p(\mathbf{x}_t | \mathbf{x}_1, \dots, \mathbf{x}_{t-1})$
- ▶ Detect abnormal/changed behaviour (if $p(\mathbf{x}_t, \mathbf{x}_{t+1}, ..., | \mathbf{x}_1, ..., \mathbf{x}_{t-1})$ small)
- Recover underlying/latent/hidden causes linking entire sequence

Markov models

In general:

 $P(\mathbf{x}_1,\ldots,\mathbf{x}_t)=P(\mathbf{x}_1)P(\mathbf{x}_2|\mathbf{x}_1)P(\mathbf{x}_3|\mathbf{x}_1,\mathbf{x}_2)\cdots P(\mathbf{x}_t|\mathbf{x}_1,\mathbf{x}_2\ldots,\mathbf{x}_{t-1})$

Markov models

In general:

$$P(\mathbf{x}_1,\ldots,\mathbf{x}_t)=P(\mathbf{x}_1)P(\mathbf{x}_2|\mathbf{x}_1)P(\mathbf{x}_3|\mathbf{x}_1,\mathbf{x}_2)\cdots P(\mathbf{x}_t|\mathbf{x}_1,\mathbf{x}_2\ldots,\mathbf{x}_{t-1})$$

First-order Markov model:

$$P(\mathbf{x}_1, \dots, \mathbf{x}_t) = P(\mathbf{x}_1)P(\mathbf{x}_2|\mathbf{x}_1) \cdots P(\mathbf{x}_t|\mathbf{x}_{t-1})$$

$$(\mathbf{x}_1) \longrightarrow (\mathbf{x}_2) \longrightarrow (\mathbf{x}_3) \longrightarrow \bullet \bullet \bullet \longrightarrow (\mathbf{x}_T)$$

The term *Markov* refers to a conditional independence relationship. In this case, the Markov property is that, given the present observation (\mathbf{x}_t) , the future $(\mathbf{x}_{t+1}, ...)$ is independent of the past $(\mathbf{x}_1, ..., \mathbf{x}_{t-1})$.

Markov models

In general:

$$P(\mathbf{x}_1,\ldots,\mathbf{x}_t)=P(\mathbf{x}_1)P(\mathbf{x}_2|\mathbf{x}_1)P(\mathbf{x}_3|\mathbf{x}_1,\mathbf{x}_2)\cdots P(\mathbf{x}_t|\mathbf{x}_1,\mathbf{x}_2\ldots,\mathbf{x}_{t-1})$$

First-order Markov model:

$$P(\mathbf{x}_1, \dots, \mathbf{x}_t) = P(\mathbf{x}_1)P(\mathbf{x}_2|\mathbf{x}_1) \cdots P(\mathbf{x}_t|\mathbf{x}_{t-1})$$

$$(\mathbf{x}_1) \longrightarrow (\mathbf{x}_2) \longrightarrow (\mathbf{x}_3) \longrightarrow \bullet \bullet \bullet \longrightarrow (\mathbf{x}_T)$$

The term *Markov* refers to a conditional independence relationship. In this case, the Markov property is that, given the present observation (\mathbf{x}_t) , the future $(\mathbf{x}_{t+1}, ...)$ is independent of the past $(\mathbf{x}_1, ..., \mathbf{x}_{t-1})$.

Second-order Markov model:

$$P(\mathbf{x}_1,\ldots,\mathbf{x}_t)=P(\mathbf{x}_1)P(\mathbf{x}_2|\mathbf{x}_1)\cdots P(\mathbf{x}_{t-1}|\mathbf{x}_{t-3},\mathbf{x}_{t-2})P(\mathbf{x}_t|\mathbf{x}_{t-2},\mathbf{x}_{t-1})$$



Causal structure and latent variables



Temporal dependence captured by latents, with observations conditionally independent. Speech recognition:

- y underlying phonemes or words
- x acoustic waveform

Vision:

- > y object identities, poses, illumination
- x image pixel values

Industrial Monitoring:

- y current state of molten steel in caster
- x temperature and pressure sensor readings

Latent-Chain models



Joint probability factorizes:

$$P(\mathbf{y}_{1:T}, \mathbf{x}_{1:T}) = P(\mathbf{y}_1)P(\mathbf{x}_1|\mathbf{y}_1)\prod_{t=2}^T P(\mathbf{y}_t|\mathbf{y}_{t-1})P(\mathbf{x}_t|\mathbf{y}_t)$$

where \mathbf{y}_t and \mathbf{x}_t are both real-valued vectors, and $\Box_{1:T} \equiv \Box_1, \ldots, \Box_T$.

Two frequently-used tractable models:

- Linear-Gaussian state-space models
- Hidden Markov models

Linear-Gaussian state-space models (SSMs)



In a linear Gaussian SSM all conditional distributions are linear and Gaussian:

Output equation:	$\mathbf{x}_t = C \mathbf{y}_t + \mathbf{v}_t$
State dynamics equation:	$\mathbf{y}_t = A\mathbf{y}_{t-1} + \mathbf{w}_t$

where \mathbf{v}_t and \mathbf{w}_t are uncorrelated zero-mean multivariate Gaussian noise vectors.

Also assume \mathbf{y}_1 is multivariate Gaussian. The joint distribution over all variables $\mathbf{x}_{1:T}, \mathbf{y}_{1:T}$ is (one big) multivariate Gaussian.

These models are also known as stochastic linear dynamical systems, Kalman filter models.

From factor analysis to state space models





Interpretation 1:

- Observations confined near low-dimensional subspace (as in FA/PCA).
- Successive observations are generated from correlated points in the latent space.

 \leftrightarrow

- However:
 - FA requires K < D and Ψ diagonal; SSMs may have K ≥ D and arbitrary output noise. Why?
 - Thus ML estimates of subspace by FA and SSM may differ.

Linear dynamical systems



Interpretation 2:

- Markov chain with linear dynamics y_t = Ay_t...
- ... perturbed by Gaussian innovations noise may describe stochasticity, unknown control, or model mismatch.
- Observations are a linear projection of the dynamical state, with additive iid Gaussian noise.
- Note:
 - Dynamical process (\mathbf{y}_t) may be higher dimensional than the observations (\mathbf{x}_t) .
 - Observations do not form a Markov chain longer-scale dependence reflects/reveals latent dynamics.

State Space Models with Control Inputs



State space models can be used to model the input–output behaviour of controlled systems. The observed variables are divided into inputs (\mathbf{u}_t) and outputs (\mathbf{x}_t) .

State dynamics equation: $\mathbf{y}_t = A\mathbf{y}_{t-1} + B\mathbf{u}_{t-1} + \mathbf{w}_t$.

Output equation: $\mathbf{x}_t = C\mathbf{y}_t + D\mathbf{u}_t + \mathbf{v}_t$.

Note that we can have many variants, e.g. $\mathbf{y}_t = A\mathbf{y}_{t-1} + B\mathbf{u}_t + \mathbf{w}_t$ or even $\mathbf{y}_t = A\mathbf{y}_{t-1} + B\mathbf{x}_{t-1} + \mathbf{w}_t$.

Hidden Markov models



Discrete hidden states $s_t \in \{1..., K\}$; outputs \mathbf{x}_t can be discrete or continuous. Joint probability factorizes:

$$P(s_{1:T}, \mathbf{x}_{1:T}) = P(s_1)P(\mathbf{x}_1|s_1)\prod_{t=2}^{T} P(s_t|s_{t-1})P(\mathbf{x}_t|s_t)$$

Generative process:

A first-order Markov chain generates the hidden state sequence (path):

initial state probs: $\pi_j = P(s_1 = j)$ transition matrix: $\Phi_{ij} = P(s_{t+1} = j | s_t = i)$

- A set of emission (output) distributions A_i(·) (one per state) converts state path to a sequence of observations x_t.
 - $\begin{aligned} A_j(\mathbf{x}) &= P(\mathbf{x}_t = \mathbf{x} | s_t = j) \\ A_{jk} &= P(\mathbf{x}_t = k | s_t = j) \end{aligned} \qquad (\text{for continuous } \mathbf{x}_t) \end{aligned}$

Hidden Markov models

Two interpretations:

a Markov chain with stochastic measurements:



or a mixture model with states coupled across time:



Even though hidden state sequence is first-order Markov, the output process may not be Markov of any order (for example: 1111121111311121111131...).

Discrete state, discrete output models can approximate any continuous dynamics and observation mapping even if nonlinear; however this is usually not practical.

HMMs are related to stochastic finite state machines/automata.

Input-output hidden Markov models



Hidden Markov models can also be used to model sequential input-output behaviour:

$$P(s_{1:T}, \mathbf{x}_{1:T}|u_{1:T}) = P(s_1|u_1)P(\mathbf{x}_1|s_1, u_1)\prod_{t=2}^T P(s_t|s_{t-1}, u_{t-1})P(\mathbf{x}_t|s_t, u_t)$$

IOHMMs can capture arbitrarily complex input-output relationships, however the number of states required is often impractical.

HMMs and SSMs

(Linear Gaussian) State space models are the continuous state analogue of hidden Markov models.



• A continuous vector state is a very powerful representation.

For an HMM to communicate N bits of information about the past, it needs 2^N states! But a real-valued state vector can store an arbitrary number of bits in principle.



Linear-Gaussian output/dynamics are very weak.

The types of dynamics linear SSMs can capture is very limited. HMMs can in principle represent arbitrary stochastic dynamics and output mappings.



Many Extensions



- Constrained HMMs
- Continuous state models with discrete outputs for time series and static data



- Hierarchical models
- ► Hybrid systems ⇔ Mixed continuous & discrete states, switching state-space models



Richer state representations



- These are hidden Markov models with many state variables (i.e. a distributed representation of the state).
- The state can capture many more bits of information about the sequence (linear in the number of state variables).



The structure of learning and inference for both models is dictated by the factored structure.



The structure of learning and inference for both models is dictated by the factored structure.

$$P(\mathbf{x}_1,\ldots,\mathbf{x}_T,\mathbf{y}_1,\ldots,\mathbf{y}_T)=P(\mathbf{y}_1)\prod_{t=2}^T P(\mathbf{y}_t|\mathbf{y}_{t-1})\prod_{t=1}^T P(\mathbf{x}_t|\mathbf{y}_t)$$



The structure of learning and inference for both models is dictated by the factored structure.

$$P(\mathbf{x}_1,\ldots,\mathbf{x}_T,\mathbf{y}_1,\ldots,\mathbf{y}_T)=P(\mathbf{y}_1)\prod_{t=2}^T P(\mathbf{y}_t|\mathbf{y}_{t-1})\prod_{t=1}^T P(\mathbf{x}_t|\mathbf{y}_t)$$

Learning (M-step):

$$\operatorname{argmax} \left\langle \log P(\mathbf{x}_{1}, \dots, \mathbf{x}_{T}, \mathbf{y}_{1}, \dots, \mathbf{y}_{T}) \right\rangle_{q(\mathbf{y}_{1}, \dots, \mathbf{y}_{T})} = \\ \operatorname{argmax} \left[\left\langle \log P(\mathbf{y}_{1}) \right\rangle_{q(\mathbf{y}_{1})} + \sum_{t=2}^{T} \left\langle \log P(\mathbf{y}_{t} | \mathbf{y}_{t-1}) \right\rangle_{q(\mathbf{y}_{t}, \mathbf{y}_{t-1})} + \sum_{t=1}^{T} \left\langle \log P(\mathbf{x}_{t} | \mathbf{y}_{t}) \right\rangle_{q(\mathbf{y}_{t})} \right]$$



The structure of learning and inference for both models is dictated by the factored structure.

$$P(\mathbf{x}_1,\ldots,\mathbf{x}_T,\mathbf{y}_1,\ldots,\mathbf{y}_T)=P(\mathbf{y}_1)\prod_{t=2}^T P(\mathbf{y}_t|\mathbf{y}_{t-1})\prod_{t=1}^T P(\mathbf{x}_t|\mathbf{y}_t)$$

Learning (M-step):

$$\operatorname{argmax} \left\langle \log P(\mathbf{x}_{1}, \dots, \mathbf{x}_{T}, \mathbf{y}_{1}, \dots, \mathbf{y}_{T}) \right\rangle_{q(\mathbf{y}_{1}, \dots, \mathbf{y}_{T})} = \\ \operatorname{argmax} \left[\left\langle \log P(\mathbf{y}_{1}) \right\rangle_{q(\mathbf{y}_{1})} + \sum_{t=2}^{T} \left\langle \log P(\mathbf{y}_{t} | \mathbf{y}_{t-1}) \right\rangle_{q(\mathbf{y}_{t}, \mathbf{y}_{t-1})} + \sum_{t=1}^{T} \left\langle \log P(\mathbf{x}_{t} | \mathbf{y}_{t}) \right\rangle_{q(\mathbf{y}_{t})} \right]$$

So the expectations needed (in E-step) are derived from singleton and pairwise marginals.

Chain models: Inference

Three general inference problems:

Filtering:	$P(\mathbf{y}_t \mathbf{x}_1,\ldots,\mathbf{x}_t)$
Smoothing:	$P(\mathbf{y}_t \mathbf{x}_1, \dots, \mathbf{x}_T)$
Prediction:	$P(\mathbf{y}_t \mathbf{x}_1, \ldots, \mathbf{x}_{t-\Delta t})$

(also $P(\mathbf{y}_t, \mathbf{y}_{t-1} | \mathbf{x}_1, \dots, \mathbf{x}_T)$ for learning)

Chain models: Inference

Three general inference problems:

Filtering:	$P(\mathbf{y}_t \mathbf{x}_1,\ldots,\mathbf{x}_t)$	
Smoothing:	$P(\mathbf{y}_t \mathbf{x}_1, \dots, \mathbf{x}_T)$	(also $P(\mathbf{y}_t, \mathbf{y}_{t-1} \mathbf{x}_1, \dots, \mathbf{x}_T)$ for learning)
Prediction:	$P(\mathbf{y}_t \mathbf{x}_1, \ldots, \mathbf{x}_{t-\Delta t})$	

Naively, these marginal posteriors seem to require very large integrals (or sums)

$$P(\mathbf{y}_t|\mathbf{x}_1,\ldots,\mathbf{x}_t) = \int \cdots \int d\mathbf{y}_1 \ldots d\mathbf{y}_{t-1} \ P(\mathbf{y}_1,\ldots,\mathbf{y}_t|\mathbf{x}_1,\ldots,\mathbf{x}_t)$$

Chain models: Inference

Three general inference problems:

Filtering:	$P(\mathbf{y}_t \mathbf{x}_1,\ldots,\mathbf{x}_t)$	
Smoothing:	$P(\mathbf{y}_t \mathbf{x}_1, \dots, \mathbf{x}_T)$	(also $P(\mathbf{y}_t, \mathbf{y}_{t-1} \mathbf{x}_1, \dots, \mathbf{x}_T)$ for learning)
Prediction:	$P(\mathbf{y}_t \mathbf{x}_1, \dots, \mathbf{x}_{t-\Delta t})$	

Naively, these marginal posteriors seem to require very large integrals (or sums)

$$\mathcal{P}(\mathbf{y}_t|\mathbf{x}_1,\ldots,\mathbf{x}_t) = \int \cdots \int d\mathbf{y}_1 \ldots d\mathbf{y}_{t-1} \ \mathcal{P}(\mathbf{y}_1,\ldots,\mathbf{y}_t|\mathbf{x}_1,\ldots,\mathbf{x}_t)$$

but again the factored structure of the distributions will help us. The algorithms rely on a form of temporal updating or message passing.



Consider an HMM, where we want to find $P(s_t = k | \mathbf{x}_1 \dots \mathbf{x}_t) = \sum_{k_1,\dots,k_{t-1}} P(s_t = k_1,\dots,s_t = k | \mathbf{x}_1 \dots \mathbf{x}_t) = \sum_{k_1,\dots,k_{t-1}} \pi_{k_1} A_{k_1}(\mathbf{x}_1) \Phi_{k_1,k_2} A_{k_2}(\mathbf{x}_2) \dots \Phi_{k_{t-1},k} A_k(\mathbf{x}_t)$



Consider an HMM, where we want to find $P(s_t=k|\mathbf{x}_1...\mathbf{x}_t) = \sum_{\substack{k_1,...,k_{t-1}}} P(s_t=k_1,\ldots,s_t=k|\mathbf{x}_1\ldots\mathbf{x}_t) = \sum_{\substack{k_1,...,k_{t-1}}} \pi_{k_1}A_{k_1}(\mathbf{x}_1)\Phi_{k_1,k_2}A_{k_2}(\mathbf{x}_2)\ldots\Phi_{k_{t-1},k}A_k(\mathbf{x}_t)$

Naïve algorithm:

start a "bug" at each of the K states at t = 1 holding value 1



Consider an HMM, where we want to find $P(\mathbf{s}_t = k | \mathbf{x}_1 \dots \mathbf{x}_t) = \sum_{\substack{k_1, \dots, k_{t-1}}} P(\mathbf{s}_1 = k_1, \dots, \mathbf{s}_t = k | \mathbf{x}_1 \dots \mathbf{x}_t) = \sum_{\substack{k_1, \dots, k_{t-1}}} \pi_{k_1} A_{k_1}(\mathbf{x}_1) \Phi_{k_1, k_2} A_{k_2}(\mathbf{x}_2) \dots \Phi_{k_{t-1}, k} A_k(\mathbf{x}_t)$

Naïve algorithm:

- start a "bug" at each of the K states at t = 1 holding value 1
- move each bug forward in time: make copies of each bug to each subsequent state and multiply the value of each copy by transition prob. × output emission prob.



Consider an HMM, where we want to find $P(s_t=k|\mathbf{x}_1...\mathbf{x}_t) = \sum_{k_1,...,k_{t-1}} P(s_t=k_1,\ldots,s_t=k|\mathbf{x}_1...\mathbf{x}_t) = \sum_{k_1,\ldots,k_{t-1}} \pi_{k_1}A_{k_1}(\mathbf{x}_1)\Phi_{k_1,k_2}A_{k_2}(\mathbf{x}_2)\ldots\Phi_{k_{t-1},k}A_k(\mathbf{x}_t)$

Naïve algorithm:

- start a "bug" at each of the K states at t = 1 holding value 1
- move each bug forward in time: make copies of each bug to each subsequent state and multiply the value of each copy by transition prob. × output emission prob.
- repeat



Consider an HMM, where we want to find $P(s_t=k|\mathbf{x}_1...\mathbf{x}_t) = \sum_{k_1,...,k_{t-1}} P(s_t=k_1,\ldots,s_t=k|\mathbf{x}_1...\mathbf{x}_t) = \sum_{k_1,\ldots,k_{t-1}} \pi_{k_1}A_{k_1}(\mathbf{x}_1)\Phi_{k_1,k_2}A_{k_2}(\mathbf{x}_2)\ldots\Phi_{k_{t-1},k}A_k(\mathbf{x}_t)$

Naïve algorithm:

- start a "bug" at each of the K states at t = 1 holding value 1
- move each bug forward in time: make copies of each bug to each subsequent state and multiply the value of each copy by transition prob. × output emission prob.
- repeat until all bugs have reached time t
- sum up values on all K^{t-1} bugs that reach state $s_t = k$ (one bug per state path)



Consider an HMM, where we want to find $P(s_t=k|\mathbf{x}_1...\mathbf{x}_t) = \sum_{k_1,...,k_{t-1}} P(s_t=k_1,\ldots,s_t=k|\mathbf{x}_1...\mathbf{x}_t) = \sum_{k_1,...,k_{t-1}} \pi_{k_1}A_{k_1}(\mathbf{x}_1)\Phi_{k_1,k_2}A_{k_2}(\mathbf{x}_2)\ldots\Phi_{k_{t-1},k}A_k(\mathbf{x}_t)$

Naïve algorithm:

- start a "bug" at each of the K states at t = 1 holding value 1
- move each bug forward in time: make copies of each bug to each subsequent state and multiply the value of each copy by transition prob. × output emission prob.
- repeat until all bugs have reached time t
- sum up values on all K^{t-1} bugs that reach state $s_t = k$ (one bug per state path)

Clever recursion:

> at every step, replace bugs at each node with a single bug carrying sum of values

Probability updating: "Bayesian filtering"



$$P(\mathbf{y}_t|\mathbf{x}_{1:t}) = \int P(\mathbf{y}_t, \mathbf{y}_{t-1}|\mathbf{x}_{1:t}) \ d\mathbf{y}_{t-1}$$


$$P(\mathbf{y}_t|\mathbf{x}_{1:t}) = \int P(\mathbf{y}_t, \mathbf{y}_{t-1}|\mathbf{x}_t, \mathbf{x}_{1:t-1}) \ d\mathbf{y}_{t-1}$$



$$\begin{aligned} P(\mathbf{y}_t | \mathbf{x}_{1:t}) &= \int P(\mathbf{y}_t, \mathbf{y}_{t-1} | \mathbf{x}_t, \mathbf{x}_{1:t-1}) \ d\mathbf{y}_{t-1} \\ &= \int \frac{P(\mathbf{x}_t, \mathbf{y}_t, \mathbf{y}_{t-1} | \mathbf{x}_{1:t-1})}{P(\mathbf{x}_t | \mathbf{x}_{1:t-1})} \ d\mathbf{y}_{t-1} \end{aligned}$$



$$P(\mathbf{y}_{t}|\mathbf{x}_{1:t}) = \int P(\mathbf{y}_{t}, \mathbf{y}_{t-1}|\mathbf{x}_{t}, \mathbf{x}_{1:t-1}) d\mathbf{y}_{t-1}$$

$$= \int \frac{P(\mathbf{x}_{t}, \mathbf{y}_{t}, \mathbf{y}_{t-1}|\mathbf{x}_{1:t-1})}{P(\mathbf{x}_{t}|\mathbf{x}_{1:t-1})} d\mathbf{y}_{t-1}$$

$$\propto \int P(\mathbf{x}_{t}|\mathbf{y}_{t}, \mathbf{y}_{t-1}, \mathbf{x}_{1:t-1}) P(\mathbf{y}_{t}|\mathbf{y}_{t-1}, \mathbf{x}_{1:t-1}) P(\mathbf{y}_{t-1}|\mathbf{x}_{1:t-1}) d\mathbf{y}_{t-1}$$



$$P(\mathbf{y}_{t}|\mathbf{x}_{1:t}) = \int P(\mathbf{y}_{t}, \mathbf{y}_{t-1}|\mathbf{x}_{t}, \mathbf{x}_{1:t-1}) d\mathbf{y}_{t-1}$$

$$= \int \frac{P(\mathbf{x}_{t}, \mathbf{y}_{t}, \mathbf{y}_{t-1}|\mathbf{x}_{1:t-1})}{P(\mathbf{x}_{t}|\mathbf{x}_{1:t-1})} d\mathbf{y}_{t-1}$$

$$\propto \int P(\mathbf{x}_{t}|\mathbf{y}_{t}, \mathbf{y}_{t-1}, \mathbf{x}_{1:t-1}) P(\mathbf{y}_{t}|\mathbf{y}_{t-1}, \mathbf{x}_{1:t-1}) P(\mathbf{y}_{t-1}|\mathbf{x}_{1:t-1}) d\mathbf{y}_{t-1}$$

$$= \int P(\mathbf{x}_t|\mathbf{y}_t)P(\mathbf{y}_t|\mathbf{y}_{t-1})P(\mathbf{y}_{t-1}|\mathbf{x}_{1:t-1}) d\mathbf{y}_{t-1}$$

Markov property



$$\begin{split} P(\mathbf{y}_{t}|\mathbf{x}_{1:t}) &= \int P(\mathbf{y}_{t}, \mathbf{y}_{t-1}|\mathbf{x}_{t}, \mathbf{x}_{1:t-1}) \, d\mathbf{y}_{t-1} \\ &= \int \frac{P(\mathbf{x}_{t}, \mathbf{y}_{t}, \mathbf{y}_{t-1}|\mathbf{x}_{1:t-1})}{P(\mathbf{x}_{t}|\mathbf{x}_{1:t-1})} \, d\mathbf{y}_{t-1} \\ &\propto \int P(\mathbf{x}_{t}|\mathbf{y}_{t}, \mathbf{y}_{t-1}, \mathbf{x}_{1:t-1}) P(\mathbf{y}_{t}|\mathbf{y}_{t-1}, \mathbf{x}_{1:t-1}) P(\mathbf{y}_{t-1}|\mathbf{x}_{1:t-1}) \, d\mathbf{y}_{t-1} \\ &= \int P(\mathbf{x}_{t}|\mathbf{y}_{t}) P(\mathbf{y}_{t}|\mathbf{y}_{t-1}) P(\mathbf{y}_{t-1}|\mathbf{x}_{1:t-1}) \, d\mathbf{y}_{t-1} \\ &\text{Markov property} \end{split}$$

This is a forward recursion based on Bayes rule.

The forward recursion for the HMM is a form of dynamic programming.

The forward recursion for the HMM is a form of dynamic programming. Define:

 $\alpha_t(i) = P(\mathbf{x}_1, \ldots, \mathbf{x}_t, \mathbf{s}_t = i|\theta)$

The forward recursion for the HMM is a form of dynamic programming. Define:

 $\alpha_t(i) = P(\mathbf{x}_1, \ldots, \mathbf{x}_t, \mathbf{s}_t = i|\theta)$

Then much like the Bayesian filtering updates, we have:

$$\alpha_1(i) = \pi_i A_i(\mathbf{x}_1) \qquad \qquad \alpha_{t+1}(i) = \left(\sum_{j=1}^K \alpha_t(j) \Phi_{ji}\right) A_i(\mathbf{x}_{t+1})$$

The forward recursion for the HMM is a form of dynamic programming. Define:

 $\alpha_t(i) = P(\mathbf{x}_1, \ldots, \mathbf{x}_t, \mathbf{s}_t = i|\theta)$

Then much like the Bayesian filtering updates, we have:

$$\alpha_1(i) = \pi_i A_i(\mathbf{x}_1) \qquad \qquad \alpha_{t+1}(i) = \left(\sum_{j=1}^K \alpha_t(j) \Phi_{ji}\right) A_i(\mathbf{x}_{t+1})$$

We've defined $\alpha_t(i)$ to be a joint rather than a posterior.

The forward recursion for the HMM is a form of dynamic programming. Define:

 $\alpha_t(i) = P(\mathbf{x}_1, \ldots, \mathbf{x}_t, \mathbf{s}_t = i|\theta)$

Then much like the Bayesian filtering updates, we have:

$$\alpha_1(i) = \pi_i A_i(\mathbf{x}_1) \qquad \qquad \alpha_{t+1}(i) = \left(\sum_{j=1}^K \alpha_t(j) \Phi_{ji}\right) A_i(\mathbf{x}_{t+1})$$

We've defined $\alpha_t(i)$ to be a joint rather than a posterior. It's easy to obtain the posterior by normalisation:

$$P(\mathbf{s}_t = i | \mathbf{x}_1, \dots, \mathbf{x}_t, \theta) = \frac{\alpha_t(i)}{\sum_k \alpha_t(k)}$$

The forward recursion for the HMM is a form of dynamic programming. Define:

 $\alpha_t(i) = P(\mathbf{x}_1, \ldots, \mathbf{x}_t, \mathbf{s}_t = i|\theta)$

Then much like the Bayesian filtering updates, we have:

$$\alpha_1(i) = \pi_i A_i(\mathbf{x}_1) \qquad \qquad \alpha_{t+1}(i) = \left(\sum_{j=1}^K \alpha_t(j) \Phi_{ji}\right) A_i(\mathbf{x}_{t+1})$$

We've defined $\alpha_t(i)$ to be a joint rather than a posterior. It's easy to obtain the posterior by normalisation:

$$\mathsf{P}(\mathbf{s}_t = i | \mathbf{x}_1, \dots, \mathbf{x}_t, \theta) = \frac{\alpha_t(i)}{\sum_k \alpha_t(k)}$$

This form enables us to compute the likelihood for $\theta = \{A, \Phi, \pi\}$ efficiently in $\mathcal{O}(TK^2)$ time:

$$P(\mathbf{x}_1 \dots \mathbf{x}_T | \theta) = \sum_{s_1, \dots, s_T} P(\mathbf{x}_1, \dots, \mathbf{x}_T, s_1, \dots, s_T, \theta) = \sum_{k=1}^K \alpha_T(k)$$

avoiding the exponential number of paths in the naïve sum (number of paths = K^{T}).



$$egin{aligned} \mathbf{y}_1 &\sim \mathcal{N}(oldsymbol{\mu}_0, oldsymbol{Q}_0) \ \mathbf{y}_t | \mathbf{y}_{t-1} &\sim \mathcal{N}(A \mathbf{y}_{t-1}, oldsymbol{Q}) \ \mathbf{x}_t | \mathbf{y}_t &\sim \mathcal{N}(C \mathbf{y}_t, R) \end{aligned}$$

For the SSM, the sums become integrals.



For the SSM, the sums become integrals. Let $\hat{\mathbf{y}}_1^0 = \mu_0$ and $\hat{V}_1^0 = Q_0$; then (cf. FA) $P(\mathbf{y}_1|\mathbf{x}_1) = \mathcal{N}(\hat{\mathbf{y}}_1^0 + \mathcal{K}_1(\mathbf{x}_1 - C\hat{\mathbf{y}}_i^0), \hat{V}_1^0 - \mathcal{K}_1C\hat{V}_1^0)$





For the SSM, the sums become integrals. Let $\hat{\mathbf{y}}_1^0 = \mu_0$ and $\hat{V}_1^0 = Q_0$; then (cf. FA)

$$P(\mathbf{y}_{1}|\mathbf{x}_{1}) = \mathcal{N}(\underbrace{\hat{\mathbf{y}}_{1}^{0} + \mathcal{K}_{1}(\mathbf{x}_{1} - C\hat{\mathbf{y}}_{i}^{0})}_{\hat{\mathbf{y}}_{1}^{1}}, \underbrace{\hat{\mathcal{V}}_{1}^{0} - \mathcal{K}_{1}C\hat{\mathcal{V}}_{1}^{0}}_{\hat{\mathcal{V}}_{1}^{1}}) \qquad \qquad \mathcal{K}_{1} = \hat{\mathcal{V}}_{1}^{0}C^{\mathsf{T}}(C\hat{\mathcal{V}}_{1}^{0}C^{\mathsf{T}} + R)^{-1}$$



For the SSM, the sums become integrals. Let $\hat{\mathbf{y}}_1^0 = \boldsymbol{\mu}_0$ and $\hat{V}_1^0 = Q_0$; then (cf. FA)

$$P(\mathbf{y}_{1}|\mathbf{x}_{1}) = \mathcal{N}(\underbrace{\hat{\mathbf{y}}_{1}^{0} + K_{1}(\mathbf{x}_{1} - C\hat{\mathbf{y}}_{i}^{0})}_{\hat{\mathbf{y}}_{1}^{1}}, \underbrace{\hat{V}_{1}^{0} - K_{1}C\hat{V}_{1}^{0}}_{\hat{V}_{1}^{1}}) \qquad K_{1} = \hat{V}_{1}^{0}C^{\mathsf{T}}(C\hat{V}_{1}^{0}C^{\mathsf{T}} + R)^{-1}$$



For the SSM, the sums become integrals. Let $\hat{\mathbf{y}}_1^0 = \boldsymbol{\mu}_0$ and $\hat{V}_1^0 = Q_0$; then (cf. FA)

$$P(\mathbf{y}_{1}|\mathbf{x}_{1}) = \mathcal{N}(\underbrace{\hat{\mathbf{y}}_{1}^{0} + K_{1}(\mathbf{x}_{1} - C\hat{\mathbf{y}}_{i}^{0})}_{\hat{\mathbf{y}}_{1}^{1}}, \underbrace{\hat{V}_{1}^{0} - K_{1}C\hat{V}_{1}^{0}}_{\hat{V}_{1}^{1}}) \qquad K_{1} = \hat{V}_{1}^{0}C^{\mathsf{T}}(C\hat{V}_{1}^{0}C^{\mathsf{T}} + R)^{-1}$$

$$P(\mathbf{y}_t|\mathbf{x}_{1:t-1}) = \int d\mathbf{y}_{t-1} P(\mathbf{y}_t|\mathbf{y}_{t-1}) P(\mathbf{y}_{t-1}|\mathbf{x}_{1:t-1})$$



For the SSM, the sums become integrals. Let $\hat{\mathbf{y}}_1^0 = \boldsymbol{\mu}_0$ and $\hat{V}_1^0 = Q_0$; then (cf. FA)

$$P(\mathbf{y}_{1}|\mathbf{x}_{1}) = \mathcal{N}(\underbrace{\hat{\mathbf{y}}_{1}^{0} + K_{1}(\mathbf{x}_{1} - C\hat{\mathbf{y}}_{i}^{0})}_{\hat{\mathbf{y}}_{1}^{1}}, \underbrace{\hat{V}_{1}^{0} - K_{1}C\hat{V}_{1}^{0}}_{\hat{V}_{1}^{1}}) \qquad K_{1} = \hat{V}_{1}^{0}C^{\mathsf{T}}(C\hat{V}_{1}^{0}C^{\mathsf{T}} + R)^{-1}$$

$$P(\mathbf{y}_{t}|\mathbf{x}_{1:t-1}) = \int d\mathbf{y}_{t-1} P(\mathbf{y}_{t}|\mathbf{y}_{t-1}) P(\mathbf{y}_{t-1}|\mathbf{x}_{1:t-1}) = \mathcal{N}(\underbrace{A\hat{\mathbf{y}}_{t-1}^{t-1}}_{\hat{\mathbf{y}}_{t}^{t-1}}, \underbrace{A\hat{V}_{t-1}^{t-1}A^{\mathsf{T}} + Q}_{\hat{V}_{t-1}^{t-1}})$$



For the SSM, the sums become integrals. Let $\hat{\mathbf{y}}_1^0 = \boldsymbol{\mu}_0$ and $\hat{V}_1^0 = Q_0$; then (cf. FA)

$$P(\mathbf{y}_{1}|\mathbf{x}_{1}) = \mathcal{N}(\underbrace{\hat{\mathbf{y}}_{1}^{0} + \mathcal{K}_{1}(\mathbf{x}_{1} - C\hat{\mathbf{y}}_{i}^{0})}_{\hat{\mathbf{y}}_{1}^{1}}, \underbrace{\hat{\mathcal{V}}_{1}^{0} - \mathcal{K}_{1}C\hat{\mathcal{V}}_{1}^{0}}_{\hat{\mathcal{V}}_{1}^{1}}) \qquad \qquad \mathcal{K}_{1} = \hat{\mathcal{V}}_{1}^{0}C^{\mathsf{T}}(C\hat{\mathcal{V}}_{1}^{0}C^{\mathsf{T}} + R)^{-1}$$

$$P(\mathbf{y}_{t}|\mathbf{x}_{1:t-1}) = \int d\mathbf{y}_{t-1} P(\mathbf{y}_{t}|\mathbf{y}_{t-1}) P(\mathbf{y}_{t-1}|\mathbf{x}_{1:t-1}) = \mathcal{N}(\underbrace{A\hat{\mathbf{y}}_{t-1}^{t-1}}_{\hat{\mathbf{y}}_{t}^{t-1}}, \underbrace{A\hat{V}_{t-1}^{t-1}A^{\mathsf{T}} + Q}_{\hat{V}_{t}^{t-1}})$$

$$P(\mathbf{y}_{t}|\mathbf{x}_{1:t}) = \mathcal{N}(\underbrace{\hat{\mathbf{y}}_{t}^{t-1} + \mathbf{K}_{t}(\mathbf{x}_{t} - C\hat{\mathbf{y}}_{t}^{t-1})}_{\hat{\mathbf{y}}_{t}^{t}}, \underbrace{\hat{V}_{t}^{t-1} - \mathbf{K}_{t}C\hat{V}_{t}^{t-1}}_{\hat{V}_{t}^{t}})$$

$$\underbrace{K_{t} = \hat{V}_{t}^{t-1}C^{\mathsf{T}}(C\hat{V}_{t}^{t-1}C^{\mathsf{T}} + R)^{-1}}_{\mathsf{Kalman gain}}$$



For the SSM, the sums become integrals. Let $\hat{\mathbf{y}}_1^0 = \boldsymbol{\mu}_0$ and $\hat{V}_1^0 = Q_0$; then (cf. FA)

$$P(\mathbf{y}_{1}|\mathbf{x}_{1}) = \mathcal{N}(\underbrace{\hat{\mathbf{y}}_{1}^{0} + \mathcal{K}_{1}(\mathbf{x}_{1} - C\hat{\mathbf{y}}_{i}^{0})}_{\hat{\mathbf{y}}_{1}^{1}}, \underbrace{\hat{\mathcal{V}}_{1}^{0} - \mathcal{K}_{1}C\hat{\mathcal{V}}_{1}^{0}}_{\hat{\mathcal{V}}_{1}^{1}}) \qquad \qquad \mathcal{K}_{1} = \hat{\mathcal{V}}_{1}^{0}C^{\mathsf{T}}(C\hat{\mathcal{V}}_{1}^{0}C^{\mathsf{T}} + R)^{-1}$$

$$P(\mathbf{y}_{t}|\mathbf{x}_{1:t-1}) = \int d\mathbf{y}_{t-1} P(\mathbf{y}_{t}|\mathbf{y}_{t-1}) P(\mathbf{y}_{t-1}|\mathbf{x}_{1:t-1}) = \mathcal{N}(\underbrace{A\hat{\mathbf{y}}_{t-1}^{t-1}}_{\hat{\mathbf{y}}_{t}^{t-1}}, \underbrace{A\hat{V}_{t-1}^{t-1}A^{\mathsf{T}} + Q}_{\hat{V}_{t}^{t-1}})$$

$$P(\mathbf{y}_{t}|\mathbf{x}_{1:t}) = \mathcal{N}(\underbrace{\hat{\mathbf{y}}_{t}^{t-1} + \mathcal{K}_{t}(\mathbf{x}_{t} - C\hat{\mathbf{y}}_{t}^{t-1})}_{\hat{\mathbf{y}}_{t}^{t}}, \underbrace{\hat{V}_{t}^{t-1} - \mathcal{K}_{t}C\hat{V}_{t}^{t-1}}_{\hat{V}_{t}^{t}})$$

$$\underbrace{K_{t} = \underbrace{\hat{V}_{t}^{t-1}C^{\mathsf{T}}(C\hat{V}_{t}^{t-1}C^{\mathsf{T}} + R)^{-1}}_{\mathsf{Kalman qain}}$$



For the SSM, the sums become integrals. Let $\hat{\mathbf{y}}_1^0 = \boldsymbol{\mu}_0$ and $\hat{V}_1^0 = Q_0$; then (cf. FA)

$$P(\mathbf{y}_{1}|\mathbf{x}_{1}) = \mathcal{N}(\underbrace{\hat{\mathbf{y}}_{1}^{0} + \mathcal{K}_{1}(\mathbf{x}_{1} - C\hat{\mathbf{y}}_{i}^{0})}_{\hat{\mathbf{y}}_{1}^{1}}, \underbrace{\hat{\mathcal{V}}_{1}^{0} - \mathcal{K}_{1}C\hat{\mathcal{V}}_{1}^{0}}_{\hat{\mathcal{V}}_{1}^{1}}) \qquad \qquad \mathcal{K}_{1} = \hat{\mathcal{V}}_{1}^{0}C^{\mathsf{T}}(C\hat{\mathcal{V}}_{1}^{0}C^{\mathsf{T}} + R)^{-1}$$

$$P(\mathbf{y}_{t}|\mathbf{x}_{1:t-1}) = \int d\mathbf{y}_{t-1} P(\mathbf{y}_{t}|\mathbf{y}_{t-1}) P(\mathbf{y}_{t-1}|\mathbf{x}_{1:t-1}) = \mathcal{N}(\underbrace{A\hat{\mathbf{y}}_{t-1}^{t-1}}_{\hat{\mathbf{y}}_{t}^{t-1}}, \underbrace{A\hat{V}_{t-1}^{t-1}A^{\mathsf{T}} + Q}_{\hat{V}_{t}^{t-1}})$$

$$P(\mathbf{y}_{t}|\mathbf{x}_{1:t}) = \mathcal{N}(\underbrace{\hat{\mathbf{y}}_{t}^{t-1} + \mathbf{K}_{t}(\mathbf{x}_{t} - C\hat{\mathbf{y}}_{t}^{t-1})}_{\hat{\mathbf{y}}_{t}^{t}}, \underbrace{\underbrace{\hat{V}_{t}^{t-1} - \mathbf{K}_{t}C\hat{V}_{t}^{t-1}}_{\hat{V}_{t}^{t}}, \underbrace{\frac{\langle \mathbf{y}\mathbf{x}^{\mathsf{T}} \rangle}{\hat{V}_{t}^{t}}}_{\mathbf{K}_{t} = \underbrace{\hat{V}_{t}^{t-1}C^{\mathsf{T}}(C\hat{V}_{t}^{t-1}C^{\mathsf{T}} + R)^{-1}}_{\mathbf{K}alman gain}$$

$$\mathsf{FA:}\ \beta = (I + \Lambda^{\mathsf{T}} \Psi^{-1} \Lambda)^{-1} \Lambda^{\mathsf{T}} \Psi^{-1} \stackrel{\text{mat. inv. lem.}}{=} \Lambda^{\mathsf{T}} (\Lambda \Lambda^{\mathsf{T}} + \Psi)^{-1}; \ \mu = \beta \mathbf{x}_n; \ \Sigma = I - \beta \Lambda.$$

The marginal posterior: "Bayesian smoothing"



 $P(\mathbf{y}_t | \mathbf{x}_{1:T})$

The marginal posterior: "Bayesian smoothing"



$$P(\mathbf{y}_t|\mathbf{x}_{1:T}) = \frac{P(\mathbf{y}_t, \mathbf{x}_{t+1:T}|\mathbf{x}_{1:t})}{P(\mathbf{x}_{t+1:T}|\mathbf{x}_{1:t})}$$

The marginal posterior: "Bayesian smoothing"



$$P(\mathbf{y}_{t}|\mathbf{x}_{1:T}) = \frac{P(\mathbf{y}_{t}, \mathbf{x}_{t+1:T}|\mathbf{x}_{1:t})}{P(\mathbf{x}_{t+1:T}|\mathbf{x}_{1:t})}$$
$$= \frac{P(\mathbf{x}_{t+1:T}|\mathbf{y}_{t})P(\mathbf{y}_{t}|\mathbf{x}_{1:t})}{P(\mathbf{x}_{t+1:T}|\mathbf{x}_{1:t})}$$

The marginal combines a backward message with the forward message found by filtering.

The HMM: Forward–Backward Algorithm

State estimation: compute marginal posterior distribution over state at time *t*:

$$\gamma_t(i) \equiv P(\mathbf{s}_{t}=i|\mathbf{x}_{1:T}) = \frac{P(\mathbf{s}_{t}=i,\mathbf{x}_{1:t})P(\mathbf{x}_{t+1:T}|\mathbf{s}_{t}=i)}{P(\mathbf{x}_{1:T})} = \frac{\alpha_t(i)\beta_t(i)}{\sum_j \alpha_t(j)\beta_t(j)}$$

where there is a simple backward recursion for

$$\beta_{t}(i) \equiv P(\mathbf{x}_{t+1:T}|s_{t}=i) = \sum_{j=1}^{K} P(s_{t+1}=j, \mathbf{x}_{t+1}, \mathbf{x}_{t+2:T}|s_{t}=i)$$
$$= \sum_{j=1}^{K} P(s_{t+1}=j|s_{t}=i) P(\mathbf{x}_{t+1}|s_{t+1}=j) P(\mathbf{x}_{t+2:T}|s_{t+1}=j) = \sum_{j=1}^{K} \Phi_{ij} A_{j}(\mathbf{x}_{t+1}) \beta_{t+1}(j)$$

 $\alpha_t(i)$ gives total *inflow* of probabilities to node (t, i); $\beta_t(i)$ gives total *outflow* of probabilities.



Bugs again: the bugs run forward from time 0 to t and backward from time T to t.

• The numbers $\gamma_t(i)$ computed by forward-backward give the marginal posterior distribution over states at each time.

- The numbers $\gamma_t(i)$ computed by forward-backward give the marginal posterior distribution over states at each time.
- By choosing the state i_t^* with the largest $\gamma_t(i)$ at each time, we can make a "best" state path. This is the path with the maximum expected number of correct states.

- The numbers $\gamma_t(i)$ computed by forward-backward give the marginal posterior distribution over states at each time.
- By choosing the state i_t^* with the largest $\gamma_t(i)$ at each time, we can make a "best" state path. This is the path with the maximum expected number of correct states.
- But it is not the single path with the highest probability of generating the data. In fact it may be a path of probability zero!

- > The numbers $\gamma_t(i)$ computed by forward-backward give the marginal posterior distribution over states at each time.
- By choosing the state i_t^* with the largest $\gamma_t(i)$ at each time, we can make a "best" state path. This is the path with the maximum expected number of correct states.
- But it is not the single path with the highest probability of generating the data. In fact it may be a path of probability zero!
- To find the single best path, we use the Viterbi decoding algorithm which is just Bellman's dynamic programming algorithm applied to this problem. This is an inference algorithm which computes the most probable state sequences: argmax P(s_{1:T} | x_{1:T}, θ)

- The numbers $\gamma_t(i)$ computed by forward-backward give the marginal posterior distribution over states at each time.
- By choosing the state i_t^* with the largest $\gamma_t(i)$ at each time, we can make a "best" state path. This is the path with the maximum expected number of correct states.
- But it is not the single path with the highest probability of generating the data. In fact it may be a path of probability zero!
- To find the single best path, we use the Viterbi decoding algorithm which is just Bellman's dynamic programming algorithm applied to this problem. This is an inference algorithm which computes the most probable state sequences: argmaxP(s_{1:T}|x_{1:T}, θ)
- The recursions look the same as forward-backward, except with max instead of \sum .

- The numbers $\gamma_t(i)$ computed by forward-backward give the marginal posterior distribution over states at each time.
- By choosing the state i_t^* with the largest $\gamma_t(i)$ at each time, we can make a "best" state path. This is the path with the maximum expected number of correct states.
- But it is not the single path with the highest probability of generating the data. In fact it may be a path of probability zero!
- To find the single best path, we use the Viterbi decoding algorithm which is just Bellman's dynamic programming algorithm applied to this problem. This is an inference algorithm which computes the most probable state sequences: argmaxP(s_{1:T} | x_{1:T}, θ) s_{1:T}
- ► The recursions look the same as forward-backward, except with max instead of ∑.
- Bugs once more: same trick except at each step kill all bugs but the one with the highest value at the node.

- The numbers $\gamma_t(i)$ computed by forward-backward give the marginal posterior distribution over states at each time.
- By choosing the state i_t^* with the largest $\gamma_t(i)$ at each time, we can make a "best" state path. This is the path with the maximum expected number of correct states.
- But it is not the single path with the highest probability of generating the data. In fact it may be a path of probability zero!
- To find the single best path, we use the Viterbi decoding algorithm which is just Bellman's dynamic programming algorithm applied to this problem. This is an inference algorithm which computes the most probable state sequences: argmaxP(s_{1:T} | x_{1:T}, θ) s_{1:T}
- ► The recursions look the same as forward-backward, except with max instead of ∑.
- Bugs once more: same trick except at each step kill all bugs but the one with the highest value at the node.
- There is also a modified EM training based on the Viterbi decoder (assignment).

The LGSSM: Kalman smoothing



We use a slightly different decomposition:

$$P(\mathbf{y}_{t}|\mathbf{x}_{1:T}) = \int P(\mathbf{y}_{t}, \mathbf{y}_{t+1}|\mathbf{x}_{1:T}) d\mathbf{y}_{t+1}$$
$$= \int P(\mathbf{y}_{t}|\mathbf{y}_{t+1}, \mathbf{x}_{1:T}) P(\mathbf{y}_{t+1}|\mathbf{x}_{1:T}) d\mathbf{y}_{t+1}$$
$$= \int P(\mathbf{y}_{t}|\mathbf{y}_{t+1}, \mathbf{x}_{1:T}) P(\mathbf{y}_{t+1}|\mathbf{x}_{1:T}) d\mathbf{y}_{t+1}$$
Markov property

This gives the additional backward recursion:

$$\begin{aligned} \mathbf{J}_t &= \hat{\mathbf{V}}_t^t \mathbf{A}^\mathsf{T} (\hat{\mathbf{V}}_{t+1}^t)^{-1} \\ \hat{\mathbf{y}}_t^\mathsf{T} &= \hat{\mathbf{y}}_t^t + J_t (\hat{\mathbf{y}}_{t+1}^\mathsf{T} - \mathbf{A} \hat{\mathbf{y}}_t^t) \\ \hat{\mathbf{V}}_t^\mathsf{T} &= \hat{\mathbf{V}}_t^t + J_t (\hat{\mathbf{V}}_{t+1}^\mathsf{T} - \hat{\mathbf{V}}_{t+1}^t) J_t^\mathsf{T} \end{aligned}$$

ML Learning for SSMs using batch EM



Parameters: $\theta = \{\mu_0, Q_0, A, Q, C, R\}$ Free energy:

$$\mathcal{F}(q,\theta) = \int d\mathbf{y}_{1:T} \ q(\mathbf{y}_{1:T})(\log P(\mathbf{x}_{1:T}, \mathbf{y}_{1:T}|\theta) - \log q(\mathbf{y}_{1:T}))$$

E-step: Maximise \mathcal{F} w.r.t. q with θ fixed: $q^*(\mathbf{y}) = p(\mathbf{y}|\mathbf{x}, \theta)$ This can be achieved with a two-state extension of the Kalman smoother.

M-step: Maximize \mathcal{F} w.r.t. θ with q fixed.

This boils down to solving a few weighted least squares problems, since all the variables in:

$$\rho(\mathbf{y}, \mathbf{x}|\theta) = \rho(\mathbf{y}_1)\rho(\mathbf{x}_1|\mathbf{y}_1) \prod_{t=2}^{T} \rho(\mathbf{y}_t|\mathbf{y}_{t-1})\rho(\mathbf{x}_t|\mathbf{y}_t)$$

form a multivariate Gaussian.

The M step for C

$$p(\mathbf{x}_t|\mathbf{y}_t) \propto \exp\left[-\frac{1}{2}(\mathbf{x}_t - C\mathbf{y}_t)^{\mathsf{T}} \mathbf{R}^{-1}(\mathbf{x}_t - C\mathbf{y}_t)\right] \Rightarrow$$

The M step for C

$$p(\mathbf{x}_t|\mathbf{y}_t) \propto \exp\left[-\frac{1}{2}(\mathbf{x}_t - C\mathbf{y}_t)^{\mathsf{T}}R^{-1}(\mathbf{x}_t - C\mathbf{y}_t)
ight] \Rightarrow$$

$$C_{\text{new}} = \operatorname*{argmax}_{C} \left\langle \sum_{t} \ln p(\mathbf{x}_{t} | \mathbf{y}_{t}) \right\rangle_{q}$$
$$p(\mathbf{x}_t|\mathbf{y}_t) \propto \exp\left[-\frac{1}{2}(\mathbf{x}_t - C\mathbf{y}_t)^{\mathsf{T}} \mathbf{R}^{-1}(\mathbf{x}_t - C\mathbf{y}_t)\right] \Rightarrow$$

$$\begin{split} C_{\text{new}} &= \operatorname*{argmax}_{\mathcal{C}} \left\langle \sum_{t} \ln p(\mathbf{x}_{t} | \mathbf{y}_{t}) \right\rangle_{q} \\ &= \operatorname*{argmax}_{\mathcal{C}} \left\langle -\frac{1}{2} \sum_{t} (\mathbf{x}_{t} - \mathbf{C} \mathbf{y}_{t})^{\mathsf{T}} \mathbf{R}^{-1} (\mathbf{x}_{t} - \mathbf{C} \mathbf{y}_{t}) \right\rangle_{q} + \text{const} \end{split}$$

$$p(\mathbf{x}_{t}|\mathbf{y}_{t}) \propto \exp\left[-\frac{1}{2}(\mathbf{x}_{t} - C\mathbf{y}_{t})^{\mathsf{T}}R^{-1}(\mathbf{x}_{t} - C\mathbf{y}_{t})\right] \Rightarrow$$

$$C_{\mathsf{new}} = \operatorname*{argmax}_{C} \left\langle \sum_{t} \ln p(\mathbf{x}_{t}|\mathbf{y}_{t}) \right\rangle_{q}$$

$$= \operatorname*{argmax}_{C} \left\langle -\frac{1}{2} \sum_{t} (\mathbf{x}_{t} - C\mathbf{y}_{t})^{\mathsf{T}}R^{-1}(\mathbf{x}_{t} - C\mathbf{y}_{t}) \right\rangle_{q} + \operatorname{const}$$

$$= \operatorname*{argmax}_{C} \left\{ -\frac{1}{2} \sum_{t} \mathbf{x}_{t}^{\mathsf{T}}R^{-1}\mathbf{x}_{t} - 2\mathbf{x}_{t}^{\mathsf{T}}R^{-1}C\langle\mathbf{y}_{t}\rangle + \langle\mathbf{y}_{t}^{\mathsf{T}}C^{\mathsf{T}}R^{-1}C\mathbf{y}_{t}\rangle \right\}$$

$$p(\mathbf{x}_{t}|\mathbf{y}_{t}) \propto \exp\left[-\frac{1}{2}(\mathbf{x}_{t} - C\mathbf{y}_{t})^{\mathsf{T}}R^{-1}(\mathbf{x}_{t} - C\mathbf{y}_{t})\right] \Rightarrow$$

$$C_{\mathsf{new}} = \operatorname*{argmax}_{C} \left\langle \sum_{t} \ln p(\mathbf{x}_{t}|\mathbf{y}_{t}) \right\rangle_{q}$$

$$= \operatorname*{argmax}_{C} \left\langle -\frac{1}{2} \sum_{t} (\mathbf{x}_{t} - C\mathbf{y}_{t})^{\mathsf{T}}R^{-1}(\mathbf{x}_{t} - C\mathbf{y}_{t}) \right\rangle_{q} + \operatorname{const}$$

$$= \operatorname*{argmax}_{C} \left\{ -\frac{1}{2} \sum_{t} \mathbf{x}_{t}^{\mathsf{T}}R^{-1}\mathbf{x}_{t} - 2\mathbf{x}_{t}^{\mathsf{T}}R^{-1}C\langle\mathbf{y}_{t}\rangle + \langle\mathbf{y}_{t}^{\mathsf{T}}C^{\mathsf{T}}R^{-1}C\mathbf{y}_{t}\rangle \right\}$$

$$= \operatorname*{argmax}_{C} \left\{ \operatorname{Tr}\left[C \sum_{t} \langle\mathbf{y}_{t}\rangle\mathbf{x}_{t}^{\mathsf{T}}R^{-1}\right] - \frac{1}{2}\operatorname{Tr}\left[C^{\mathsf{T}}R^{-1}C\left\langle\sum_{t} \mathbf{y}_{t}\mathbf{y}_{t}^{\mathsf{T}}\right\rangle \right] \right\}$$

$$\begin{split} p(\mathbf{x}_{t}|\mathbf{y}_{t}) &\propto \exp\left[-\frac{1}{2}(\mathbf{x}_{t}-C\mathbf{y}_{t})^{\mathsf{T}}R^{-1}(\mathbf{x}_{t}-C\mathbf{y}_{t})\right] &\Rightarrow \\ C_{\mathsf{new}} &= \operatorname*{argmax}_{C} \left\langle \sum_{t} \ln p(\mathbf{x}_{t}|\mathbf{y}_{t}) \right\rangle_{q} \\ &= \operatorname*{argmax}_{C} \left\langle -\frac{1}{2} \sum_{t} (\mathbf{x}_{t}-C\mathbf{y}_{t})^{\mathsf{T}}R^{-1}(\mathbf{x}_{t}-C\mathbf{y}_{t}) \right\rangle_{q} + \operatorname{const} \\ &= \operatorname*{argmax}_{C} \left\{ -\frac{1}{2} \sum_{t} \mathbf{x}_{t}^{\mathsf{T}}R^{-1}\mathbf{x}_{t} - 2\mathbf{x}_{t}^{\mathsf{T}}R^{-1}C\langle\mathbf{y}_{t}\rangle + \langle\mathbf{y}_{t}^{\mathsf{T}}C^{\mathsf{T}}R^{-1}C\mathbf{y}_{t}\rangle \right\} \\ &= \operatorname*{argmax}_{C} \left\{ \operatorname{Tr} \left[C \sum_{t} \langle\mathbf{y}_{t}\rangle\mathbf{x}_{t}^{\mathsf{T}}R^{-1} \right] - \frac{1}{2}\operatorname{Tr} \left[C^{\mathsf{T}}R^{-1}C\left\langle \sum_{t} \mathbf{y}_{t}\mathbf{y}_{t}^{\mathsf{T}}\right\rangle \right] \right\} \\ &\operatorname{using} \frac{\partial \operatorname{Tr}[AB]}{\partial A} = B^{\mathsf{T}}, \text{ we have } \frac{\partial\{\cdot\}}{\partial C} = R^{-1} \sum_{t} \mathbf{x}_{t}\langle\mathbf{y}_{t}\rangle^{\mathsf{T}} - R^{-1}C\left\langle \sum_{t} \mathbf{y}_{t}\mathbf{y}_{t}^{\mathsf{T}}\right\rangle \end{split}$$

$$\begin{split} \rho(\mathbf{x}_{t}|\mathbf{y}_{t}) &\propto \exp\left[-\frac{1}{2}(\mathbf{x}_{t}-C\mathbf{y}_{t})^{\mathsf{T}}R^{-1}(\mathbf{x}_{t}-C\mathbf{y}_{t})\right] &\Rightarrow \\ C_{\mathsf{new}} &= \operatorname*{argmax}_{C} \left\langle \sum_{t} \ln p(\mathbf{x}_{t}|\mathbf{y}_{t}) \right\rangle_{q} \\ &= \operatorname*{argmax}_{C} \left\langle -\frac{1}{2} \sum_{t} (\mathbf{x}_{t}-C\mathbf{y}_{t})^{\mathsf{T}}R^{-1}(\mathbf{x}_{t}-C\mathbf{y}_{t}) \right\rangle_{q} + \operatorname{const} \\ &= \operatorname*{argmax}_{C} \left\{ -\frac{1}{2} \sum_{t} \mathbf{x}_{t}^{\mathsf{T}}R^{-1}\mathbf{x}_{t} - 2\mathbf{x}_{t}^{\mathsf{T}}R^{-1}C\langle\mathbf{y}_{t}\rangle + \langle\mathbf{y}_{t}^{\mathsf{T}}C^{\mathsf{T}}R^{-1}C\mathbf{y}_{t}\rangle \right\} \\ &= \operatorname*{argmax}_{C} \left\{ \operatorname{Tr}\left[C \sum_{t} \langle\mathbf{y}_{t}\rangle\mathbf{x}_{t}^{\mathsf{T}}R^{-1}\right] - \frac{1}{2}\operatorname{Tr}\left[C^{\mathsf{T}}R^{-1}C\left\langle\sum_{t} \mathbf{y}_{t}\mathbf{y}_{t}^{\mathsf{T}}\right\rangle\right] \right\} \\ \operatorname{using} \frac{\partial \operatorname{Tr}[AB]}{\partial A} &= B^{\mathsf{T}}, \text{ we have } \frac{\partial\{\cdot\}}{\partial C} = R^{-1} \sum_{t} \mathbf{x}_{t}\langle\mathbf{y}_{t}\rangle^{\mathsf{T}} - R^{-1}C\left\langle\sum_{t} \mathbf{y}_{t}\mathbf{y}_{t}^{\mathsf{T}}\right\rangle \\ &\Rightarrow C_{\mathsf{new}} &= \left(\sum_{t} \mathbf{x}_{t}\langle\mathbf{y}_{t}\rangle^{\mathsf{T}}\right) \left(\sum_{t} \langle\mathbf{y}_{t}\mathbf{y}_{t}^{\mathsf{T}}\right\rangle^{-1} \end{split}$$

Notice that this is exactly the same equation as in factor analysis and linear regression!

$$p(\mathbf{y}_{t+1}|\mathbf{y}_t) \propto \exp\left\{-\frac{1}{2}(\mathbf{y}_{t+1} - A\mathbf{y}_t)^{\mathsf{T}}Q^{-1}(\mathbf{y}_{t+1} - A\mathbf{y}_t)\right\} \Rightarrow$$

$$p(\mathbf{y}_{t+1}|\mathbf{y}_t) \propto \exp\left\{-\frac{1}{2}(\mathbf{y}_{t+1} - A\mathbf{y}_t)^{\mathsf{T}} Q^{-1}(\mathbf{y}_{t+1} - A\mathbf{y}_t)\right\} \Rightarrow$$
$$A_{\mathsf{new}} = \operatorname*{argmax}_{A} \left\langle \sum_{t} \ln p(\mathbf{y}_{t+1}|\mathbf{y}_t) \right\rangle_{q}$$

$$p(\mathbf{y}_{t+1}|\mathbf{y}_t) \propto \exp\left\{-\frac{1}{2}(\mathbf{y}_{t+1} - A\mathbf{y}_t)^{\mathsf{T}} Q^{-1}(\mathbf{y}_{t+1} - A\mathbf{y}_t)\right\} \Rightarrow$$

$$A_{\text{new}} = \operatorname*{argmax}_{A} \left\langle\sum_{t} \ln p(\mathbf{y}_{t+1}|\mathbf{y}_t)\right\rangle_{q}$$

$$= \operatorname*{argmax}_{A} \left\langle-\frac{1}{2} \sum_{t} (\mathbf{y}_{t+1} - A\mathbf{y}_t)^{\mathsf{T}} Q^{-1}(\mathbf{y}_{t+1} - A\mathbf{y}_t)\right\rangle_{q} + \text{const}$$

$$\begin{split} \rho(\mathbf{y}_{t+1}|\mathbf{y}_t) &\propto \exp\left\{-\frac{1}{2}(\mathbf{y}_{t+1} - A\mathbf{y}_t)^{\mathsf{T}}Q^{-1}(\mathbf{y}_{t+1} - A\mathbf{y}_t)\right\} &\Rightarrow \\ A_{\mathsf{new}} &= \operatorname*{argmax}_{A} \left\langle\sum_{t} \ln \rho(\mathbf{y}_{t+1}|\mathbf{y}_t)\right\rangle_{q} \\ &= \operatorname*{argmax}_{A} \left\langle-\frac{1}{2}\sum_{t}(\mathbf{y}_{t+1} - A\mathbf{y}_t)^{\mathsf{T}}Q^{-1}(\mathbf{y}_{t+1} - A\mathbf{y}_t)\right\rangle_{q} + \operatorname{const} \\ &= \operatorname*{argmax}_{A} \left\{-\frac{1}{2}\sum_{t}\mathbf{y}_{t+1}^{\mathsf{T}}Q^{-1}\mathbf{y}_{t+1} - 2\left\langle\mathbf{y}_{t+1}^{\mathsf{T}}Q^{-1}A\mathbf{y}_t\right\rangle + \left\langle\mathbf{y}_t^{\mathsf{T}}A^{\mathsf{T}}Q^{-1}A\mathbf{y}_t\right\rangle\right\} \end{split}$$

$$\begin{split} \rho(\mathbf{y}_{t+1}|\mathbf{y}_t) &\propto \exp\left\{-\frac{1}{2}(\mathbf{y}_{t+1} - A\mathbf{y}_t)^{\mathsf{T}}Q^{-1}(\mathbf{y}_{t+1} - A\mathbf{y}_t)\right\} &\Rightarrow \\ A_{\mathsf{new}} &= \operatorname*{argmax}_{A} \left\langle \sum_{t} \ln \rho(\mathbf{y}_{t+1}|\mathbf{y}_t) \right\rangle_{q} \\ &= \operatorname*{argmax}_{A} \left\langle -\frac{1}{2} \sum_{t} (\mathbf{y}_{t+1} - A\mathbf{y}_t)^{\mathsf{T}}Q^{-1}(\mathbf{y}_{t+1} - A\mathbf{y}_t) \right\rangle_{q} + \operatorname{const} \\ &= \operatorname*{argmax}_{A} \left\{ -\frac{1}{2} \sum_{t} \mathbf{y}_{t+1}^{\mathsf{T}}Q^{-1}\mathbf{y}_{t+1} - 2\left\langle \mathbf{y}_{t+1}^{\mathsf{T}}Q^{-1}A\mathbf{y}_t \right\rangle + \left\langle \mathbf{y}_t^{\mathsf{T}}A^{\mathsf{T}}Q^{-1}A\mathbf{y}_t \right\rangle \right\} \\ &= \operatorname*{argmax}_{A} \left\{ \operatorname{Tr}\left[A \sum_{t} \left\langle \mathbf{y}_t \mathbf{y}_{t+1}^{\mathsf{T}} \right\rangle Q^{-1}\right] - \frac{1}{2} \operatorname{Tr}\left[A^{\mathsf{T}}Q^{-1}A \sum_{t} \left\langle \mathbf{y}_t \mathbf{y}_t^{\mathsf{T}} \right\rangle \right] \right\} \end{split}$$

$$\begin{split} p(\mathbf{y}_{t+1}|\mathbf{y}_{t}) &\propto \exp\left\{-\frac{1}{2}(\mathbf{y}_{t+1} - A\mathbf{y}_{t})^{\mathsf{T}}Q^{-1}(\mathbf{y}_{t+1} - A\mathbf{y}_{t})\right\} &\Rightarrow \\ A_{\mathsf{new}} &= \operatorname*{argmax}_{A} \left\langle\sum_{t} \ln p(\mathbf{y}_{t+1}|\mathbf{y}_{t})\right\rangle_{q} \\ &= \operatorname*{argmax}_{A} \left\langle-\frac{1}{2}\sum_{t}(\mathbf{y}_{t+1} - A\mathbf{y}_{t})^{\mathsf{T}}Q^{-1}(\mathbf{y}_{t+1} - A\mathbf{y}_{t})\right\rangle_{q} + \operatorname{const} \\ &= \operatorname*{argmax}_{A} \left\{-\frac{1}{2}\sum_{t}\mathbf{y}_{t+1}^{\mathsf{T}}Q^{-1}\mathbf{y}_{t+1} - 2\left\langle\mathbf{y}_{t+1}^{\mathsf{T}}Q^{-1}A\mathbf{y}_{t}\right\rangle + \left\langle\mathbf{y}_{t}^{\mathsf{T}}A^{\mathsf{T}}Q^{-1}A\mathbf{y}_{t}\right\rangle\right\} \\ &= \operatorname*{argmax}_{A} \left\{\operatorname{Tr}\left[A\sum_{t}\left\langle\mathbf{y}_{t}\mathbf{y}_{t+1}^{\mathsf{T}}\right\rangle Q^{-1}\right] - \frac{1}{2}\operatorname{Tr}\left[A^{\mathsf{T}}Q^{-1}A\sum_{t}\left\langle\mathbf{y}_{t}\mathbf{y}_{t}^{\mathsf{T}}\right\rangle\right]\right\} \\ &\operatorname{using} \frac{\partial \operatorname{Tr}[AB]}{\partial A} = B^{\mathsf{T}}, \text{ we have } \frac{\partial\{\cdot\}}{\partial A} = Q^{-1}\sum_{t}\left\langle\mathbf{y}_{t+1}\mathbf{y}_{t}^{\mathsf{T}}\right\rangle - Q^{-1}A\sum_{t}\left\langle\mathbf{y}_{t}\mathbf{y}_{t}^{\mathsf{T}}\right\rangle \end{split}$$

$$p(\mathbf{y}_{t+1}|\mathbf{y}_{t}) \propto \exp\left\{-\frac{1}{2}(\mathbf{y}_{t+1} - A\mathbf{y}_{t})^{\mathsf{T}}Q^{-1}(\mathbf{y}_{t+1} - A\mathbf{y}_{t})\right\} \Rightarrow$$

$$A_{\mathsf{new}} = \operatorname*{argmax}_{A} \left\langle\sum_{t} \ln p(\mathbf{y}_{t+1}|\mathbf{y}_{t})\right\rangle_{q}$$

$$= \operatorname*{argmax}_{A} \left\langle-\frac{1}{2}\sum_{t}(\mathbf{y}_{t+1} - A\mathbf{y}_{t})^{\mathsf{T}}Q^{-1}(\mathbf{y}_{t+1} - A\mathbf{y}_{t})\right\rangle_{q} + \operatorname{const}$$

$$= \operatorname*{argmax}_{A} \left\{-\frac{1}{2}\sum_{t}\mathbf{y}_{t+1}^{\mathsf{T}}Q^{-1}\mathbf{y}_{t+1} - 2\left\langle\mathbf{y}_{t+1}^{\mathsf{T}}Q^{-1}A\mathbf{y}_{t}\right\rangle + \left\langle\mathbf{y}_{t}^{\mathsf{T}}A^{\mathsf{T}}Q^{-1}A\mathbf{y}_{t}\right\rangle\right\}$$

$$= \operatorname*{argmax}_{A} \left\{\operatorname{Tr}\left[A\sum_{t}\left\langle\mathbf{y}_{t}\mathbf{y}_{t+1}^{\mathsf{T}}\right\rangle Q^{-1}\right] - \frac{1}{2}\operatorname{Tr}\left[A^{\mathsf{T}}Q^{-1}A\sum_{t}\left\langle\mathbf{y}_{t}\mathbf{y}_{t}^{\mathsf{T}}\right\rangle\right]\right\}$$

$$\operatorname{using} \frac{\partial \operatorname{Tr}_{A}B}{\partial A} = B^{\mathsf{T}}, \text{ we have } \frac{\partial\left\{\cdot\right\}}{\partial A} = Q^{-1}\sum_{t}\left\langle\mathbf{y}_{t+1}\mathbf{y}_{t}^{\mathsf{T}}\right\rangle - Q^{-1}A\sum_{t}\left\langle\mathbf{y}_{t}\mathbf{y}_{t}^{\mathsf{T}}\right\rangle$$

$$\Rightarrow A_{\mathsf{new}} = \left(\sum_{t}\left\langle\mathbf{y}_{t+1}\mathbf{y}_{t}^{\mathsf{T}}\right\rangle\right)\left(\sum_{t}\left\langle\mathbf{y}_{t}\mathbf{y}_{t}^{\mathsf{T}}\right\rangle\right)^{-1}$$

This is still analagous to factor analysis and linear regression, with expected correlations.

Learning (online gradient)

Time series data must often be processed in real-time, and we may want to update parameters online as observations arrive. We can do so by updating a local version of the likelihood based on the Kalman filter estimates.

Consider the log likelihood contributed by each data point (ℓ_t) :

$$\ell = \sum_{t=1}^{T} \ln p(\mathbf{x}_t | \mathbf{x}_1, \dots, \mathbf{x}_{t-1}) = \sum_{t=1}^{T} \ell_t$$

Then,

$$\ell_t = -\frac{D}{2} \ln 2\pi - \frac{1}{2} \ln |\boldsymbol{\Sigma}| - \frac{1}{2} (\mathbf{x}_t - C \hat{\mathbf{y}}_t^{t-1})^{\mathsf{T}} \boldsymbol{\Sigma}^{-1} (\mathbf{x}_t - C \hat{\mathbf{y}}_t^{t-1})$$

where D is dimension of x, and:

$$\begin{split} \hat{\mathbf{y}}_{t}^{t-1} &= A \hat{\mathbf{y}}_{t-1}^{t-1} \\ \boldsymbol{\Sigma} &= C \hat{V}_{t}^{t-1} C^{\mathsf{T}} + R \\ \hat{V}_{t}^{t-1} &= A \hat{V}_{t-1}^{t-1} A^{\mathsf{T}} + Q \end{split}$$

We differentiate ℓ_t to obtain gradient rules for *A*, *C*, *Q*, *R*. The size of the gradient step (learning rate) reflects our expectation about nonstationarity.

Learning HMMs using EM



Parameters: $\theta = \{\pi, \Phi, A\}$

Free energy:

$$\mathcal{F}(q,\theta) = \sum_{s_{1:T}} q(s_{1:T}) (\log P(x_{1:T}, s_{1:T}|\theta) - \log q(s_{1:T}))$$

E-step: Maximise \mathcal{F} w.r.t. q with θ fixed: $q^*(s_{1:T}) = P(s_{1:T}|\mathbf{x}_{1:T}, \theta)$

We will only need the marginal probabilities $q(s_t, s_{t+1})$, which can also be obtained from the forward–backward algorithm.

M-step: Maximize \mathcal{F} w.r.t. θ with q fixed.

We can re-estimate the parameters by computing the expected number of times the HMM was in state i, emitted symbol k and transitioned to state j.

This is the **Baum-Welch algorithm** and it predates the (more general) EM algorithm.

We can derive the following updates by taking derivatives of \mathcal{F} w.r.t. θ .

We can derive the following updates by taking derivatives of \mathcal{F} w.r.t. θ .

• The initial state distribution is the expected number of times in state *i* at t = 1:

 $\hat{\pi}_i = \gamma_1(i)$

We can derive the following updates by taking derivatives of \mathcal{F} w.r.t. θ .

• The initial state distribution is the expected number of times in state *i* at t = 1:

 $\hat{\pi}_i = \gamma_1(i)$

The expected number of transitions from state i to j which begin at time t is:

 $\xi_t(i \to j) \equiv P(\mathbf{s}_t = i, \mathbf{s}_{t+1} = j | \mathbf{x}_{1:T}) = \alpha_t(i) \Phi_{ij} A_j(\mathbf{x}_{t+1}) \beta_{t+1}(j) / P(\mathbf{x}_{1:T})$

so the estimated transition probabilities are:

$$\widehat{\Phi}_{ij} = \sum_{t=1}^{T-1} \xi_t(i \rightarrow j) \left/ \sum_{t=1}^{T-1} \gamma_t(i) \right.$$

We can derive the following updates by taking derivatives of \mathcal{F} w.r.t. θ .

• The initial state distribution is the expected number of times in state *i* at t = 1:

 $\hat{\pi}_i = \gamma_1(i)$

The expected number of transitions from state i to j which begin at time t is:

 $\xi_t(i \to j) \equiv P(s_t = i, s_{t+1} = j | \mathbf{x}_{1:T}) = \alpha_t(i) \Phi_{ij} A_j(\mathbf{x}_{t+1}) \beta_{t+1}(j) / P(x_{1:T})$

so the estimated transition probabilities are:

$$\widehat{\Phi}_{ij} = \sum_{t=1}^{T-1} \xi_t(i \to j) \left/ \sum_{t=1}^{T-1} \gamma_t(i) \right.$$

The output distributions are the expected number of times we observe a particular symbol in a particular state:

$$\widehat{A}_{ik} = \sum_{t:\mathbf{x}_t=k} \gamma_t(i) \left/ \sum_{t=1}^T \gamma_t(i) \right|$$

(or the state-probability-weighted mean and variance for a Gaussian output model).

HMM practicalities

Numerical scaling: the conventional message definition is in terms of a large joint:

 $\alpha_t(i) = P(\mathbf{x}_{1:t}, \mathbf{s}_t = i) \rightarrow 0$ as *t* grows, and so can easily underflow.

Rescale:

$$\overline{lpha}_t(i) = A_i(\mathbf{x}_t) \sum_j \widetilde{lpha}_{t-1}(j) \Phi_{ji}$$
 $ho_t = \sum_{i=1}^K \overline{lpha}_t(i)$ $\widetilde{lpha}_t(i) = \overline{lpha}_t(i) /
ho_t$

Exercise: show that:

$$\rho_t = P(\mathbf{x}_t | \mathbf{x}_{1:t-1}, \theta) \qquad \qquad \prod_{t=1}^T \rho_t = P(\mathbf{x}_{1:T} | \theta)$$

What does this make $\tilde{\alpha}_t(i)$?

HMM practicalities

Numerical scaling: the conventional message definition is in terms of a large joint:

 $\alpha_t(i) = P(\mathbf{x}_{1:t}, \mathbf{s}_t = i) \rightarrow 0$ as *t* grows, and so can easily underflow.

Rescale:

$$\overline{\alpha}_t(i) = A_i(\mathbf{x}_t) \sum_j \widetilde{\alpha}_{t-1}(j) \Phi_{ji} \qquad \rho_t = \sum_{i=1}^n \overline{\alpha}_t(i) \qquad \widetilde{\alpha}_t(i) = \overline{\alpha}_t(i)/\rho_t$$

v

Exercise: show that:

$$\rho_t = P(\mathbf{x}_t | \mathbf{x}_{1:t-1}, \theta) \qquad \qquad \prod_{t=1}^T \rho_t = P(\mathbf{x}_{1:T} | \theta)$$

What does this make $\tilde{\alpha}_t(i)$?

 Multiple observed sequences: average numerators and denominators in the ratios of updates.

HMM practicalities

Numerical scaling: the conventional message definition is in terms of a large joint:

 $\alpha_t(i) = P(\mathbf{x}_{1:t}, \mathbf{s}_t = i) \rightarrow 0$ as *t* grows, and so can easily underflow.

Rescale:

$$\overline{\alpha}_t(i) = A_i(\mathbf{x}_t) \sum_j \widetilde{\alpha}_{t-1}(j) \Phi_{ji} \qquad \qquad \rho_t = \sum_{i=1}^n \overline{\alpha}_t(i) \qquad \qquad \widetilde{\alpha}_t(i) = \overline{\alpha}_t(i) / \rho_t$$

v

Exercise: show that:

$$\rho_t = P(\mathbf{x}_t | \mathbf{x}_{1:t-1}, \theta) \qquad \qquad \prod_{t=1}^T \rho_t = P(\mathbf{x}_{1:T} | \theta)$$

What does this make $\tilde{\alpha}_t(i)$?

- Multiple observed sequences: average numerators and denominators in the ratios of updates.
- Local optima (random restarts, annealing; see discussion later).

HMM pseudocode: inference (E step)

Forward-backward including scaling tricks. [o is the element-by-element (Hadamard/Schur) product: '.*' in matlab.]

for
$$t = 1:T$$
, $i = 1:K$ $p_t(i) = A_i(\mathbf{x}_t)$
 $\alpha_1 = \boldsymbol{\pi} \circ p_1$ $\rho_1 = \sum_{i=1}^{K} \alpha_1(i)$ $\alpha_1 = \alpha_1/\rho_1$
for $t = 2:T$ $\alpha_t = (\Phi^T * \alpha_{t-1}) \circ p_t$ $\rho_t = \sum_{i=1}^{K} \alpha_t(i)$ $\alpha_t = \alpha_t/\rho_t$
 $\beta_T = 1$
for $t = T - 1:1$ $\beta_t = \Phi * (\beta_{t+1} \circ p_{t+1})/\rho_{t+1}$
 $\log P(\mathbf{x}_{1:T}) = \sum_{t=1}^{T} \log(\rho_t)$
for $t = 1:T$ $\gamma_t = \alpha_t \circ \beta_t$
for $t = 1:T - 1$ $\xi_t = \Phi \circ (\alpha_t * (\beta_{t+1} \circ p_{t+1})^T)/\rho_{t+1}$

HMM pseudocode: parameter re-estimation (M step)

Baum-Welch parameter updates:

For each sequence I = 1 : L, run forward–backward to get $\gamma^{(I)}$ and $\xi^{(I)}$, then

$$\pi_{i} = \frac{1}{L} \sum_{l=1}^{L} \gamma_{1}^{(l)}(i)$$

$$\Phi_{ij} = \frac{\sum_{l=1}^{L} \sum_{t=1}^{T^{(l)}-1} \xi_{t}^{(l)}(ij)}{\sum_{l=1}^{L} \sum_{t=1}^{T^{(l)}-1} \gamma_{t}^{(l)}(i)}$$

$$A_{ik} = \frac{\sum_{l=1}^{L} \sum_{t=1}^{T^{(l)}} \delta(\mathbf{x}_{t} = k) \gamma_{t}^{(l)}(i)}{\sum_{l=1}^{L} \sum_{t=1}^{T^{(l)}} \gamma_{t}^{(l)}(i)}$$

Recall that the FA likelihood is conserved with respect to orthogonal transformations of y:

$$\begin{split} P(\mathbf{y}) &= \mathcal{N}\left(\mathbf{0}, l\right) \\ P(\mathbf{x} | \mathbf{y}) &= \mathcal{N}\left(\Lambda \mathbf{y}, \Psi\right) \end{split}$$

Recall that the FA likelihood is conserved with respect to orthogonal transformations of y:

$$P(\mathbf{y}) = \mathcal{N} (\mathbf{0}, I) \qquad \qquad \tilde{\mathbf{y}} = U\mathbf{y}$$
$$P(\mathbf{x}|\mathbf{y}) = \mathcal{N} (\Lambda \mathbf{y}, \Psi) \qquad \qquad \tilde{\Lambda} = \Lambda U^{\mathsf{T}}$$

Recall that the FA likelihood is conserved with respect to orthogonal transformations of y:

$$P(\mathbf{y}) = \mathcal{N}(\mathbf{0}, I) \qquad \qquad \tilde{\mathbf{y}} = U\mathbf{y} \qquad \qquad P(\tilde{\mathbf{y}}) = \mathcal{N}\left(U\mathbf{0}, UIU^{\mathsf{T}}\right) = \mathcal{N}(\mathbf{0}, I) \\ P(\mathbf{x}|\mathbf{y}) = \mathcal{N}\left(\Lambda\mathbf{y}, \Psi\right) \qquad \qquad \tilde{\Lambda} = \Lambda U^{\mathsf{T}} \qquad \Rightarrow \qquad P(\tilde{\mathbf{y}}) = \mathcal{N}\left(\Lambda U^{\mathsf{T}}U\mathbf{y}, \Psi\right) = \mathcal{N}\left(\tilde{\Lambda}\tilde{\mathbf{y}}, \Psi\right)$$

Recall that the FA likelihood is conserved with respect to orthogonal transformations of y:

$$P(\mathbf{y}) = \mathcal{N}(\mathbf{0}, l) \qquad \qquad \tilde{\mathbf{y}} = U\mathbf{y} \qquad \qquad P(\tilde{\mathbf{y}}) = \mathcal{N}\left(U\mathbf{0}, UlU^{\mathsf{T}}\right) = \mathcal{N}(\mathbf{0}, l) \\ P(\mathbf{x}|\mathbf{y}) = \mathcal{N}\left(\Lambda\mathbf{y}, \Psi\right) \qquad \qquad \tilde{\Lambda} = \Lambda U^{\mathsf{T}} \qquad \Rightarrow \qquad P(\tilde{\mathbf{y}}) = \mathcal{N}\left(\Lambda U^{\mathsf{T}}U\mathbf{y}, \Psi\right) = \mathcal{N}\left(\tilde{\Lambda}\tilde{\mathbf{y}}, \Psi\right)$$

Similarly, a mixture model is invariant to permutations of the latent.

Recall that the FA likelihood is conserved with respect to orthogonal transformations of y:

$$\begin{array}{ccc} P(\mathbf{y}) = \mathcal{N}\left(\mathbf{0}, l\right) \\ P(\mathbf{x}|\mathbf{y}) = \mathcal{N}\left(\Lambda\mathbf{y}, \Psi\right) \end{array} & \begin{array}{c} \tilde{\mathbf{y}} = U\mathbf{y} \\ \tilde{\boldsymbol{\Lambda}} = \Lambda \boldsymbol{U}^{\mathsf{T}} \end{array} \Rightarrow \begin{array}{c} P(\tilde{\mathbf{y}}) = \mathcal{N}\left(U\mathbf{0}, UIU^{\mathsf{T}}\right) = \mathcal{N}\left(\mathbf{0}, l\right) \\ P(\mathbf{x}|\tilde{\mathbf{y}}) = \mathcal{N}\left(\Lambda \boldsymbol{U}^{\mathsf{T}}U\mathbf{y}, \Psi\right) = \mathcal{N}\left(\tilde{\boldsymbol{\Lambda}}\tilde{\mathbf{y}}, \Psi\right) \end{array}$$

Similarly, a mixture model is invariant to permutations of the latent.

The LGSSM likelihood is conserved with respect to any invertible transform of the latent:

$$P(\mathbf{y}_{t+1}|\mathbf{y}_t) = \mathcal{N}(A\mathbf{y}_t, Q)$$
$$P(\mathbf{x}_t|\mathbf{y}_t) = \mathcal{N}(C\mathbf{y}, R)$$

Recall that the FA likelihood is conserved with respect to orthogonal transformations of y:

$$P(\mathbf{y}) = \mathcal{N}(\mathbf{0}, l) \qquad \qquad \tilde{\mathbf{y}} = U\mathbf{y} \qquad \qquad P(\tilde{\mathbf{y}}) = \mathcal{N}\left(U\mathbf{0}, UlU^{\mathsf{T}}\right) = \mathcal{N}(\mathbf{0}, l) \\ P(\mathbf{x}|\mathbf{y}) = \mathcal{N}\left(\Lambda\mathbf{y}, \Psi\right) \qquad \qquad \tilde{\Lambda} = \Lambda U^{\mathsf{T}} \qquad \Rightarrow \qquad P(\tilde{\mathbf{y}}) = \mathcal{N}\left(\Lambda U^{\mathsf{T}}U\mathbf{y}, \Psi\right) = \mathcal{N}\left(\tilde{\Lambda}\tilde{\mathbf{y}}, \Psi\right)$$

Similarly, a mixture model is invariant to permutations of the latent.

The LGSSM likelihood is conserved with respect to any invertible transform of the latent:

$$\begin{aligned} P(\mathbf{y}_{t+1}|\mathbf{y}_t) &= \mathcal{N}\left(A\mathbf{y}_t, Q\right) \\ P(\mathbf{x}_t|\mathbf{y}_t) &= \mathcal{N}\left(C\mathbf{y}, R\right) \end{aligned} \qquad \begin{array}{l} \tilde{\mathbf{y}} &= G\mathbf{y} \\ \tilde{\mathbf{Q}} &= G\mathbf{Q}G^{\mathsf{T}} \\ \tilde{\mathbf{C}} &= CG^{-1} \end{aligned}$$

Recall that the FA likelihood is conserved with respect to orthogonal transformations of y:

$$P(\mathbf{y}) = \mathcal{N}(\mathbf{0}, l) \qquad \qquad \tilde{\mathbf{y}} = U\mathbf{y} \qquad \qquad P(\tilde{\mathbf{y}}) = \mathcal{N}\left(U\mathbf{0}, UlU^{\mathsf{T}}\right) = \mathcal{N}(\mathbf{0}, l) \\ P(\mathbf{x}|\mathbf{y}) = \mathcal{N}\left(\Lambda\mathbf{y}, \Psi\right) \qquad \qquad \tilde{\Lambda} = \Lambda U^{\mathsf{T}} \qquad \Rightarrow \qquad P(\tilde{\mathbf{y}}) = \mathcal{N}\left(\Lambda U^{\mathsf{T}}U\mathbf{y}, \Psi\right) = \mathcal{N}\left(\tilde{\Lambda}\tilde{\mathbf{y}}, \Psi\right)$$

Similarly, a mixture model is invariant to permutations of the latent.

The LGSSM likelihood is conserved with respect to any invertible transform of the latent:

$$P(\mathbf{y}_{t+1}|\mathbf{y}_t) = \mathcal{N}(A\mathbf{y}_t, Q) \qquad \qquad \mathbf{\tilde{y}} = G\mathbf{y} \qquad \mathbf{\tilde{A}} = GAG^{-1}$$

$$P(\mathbf{x}_t|\mathbf{y}_t) = \mathcal{N}(C\mathbf{y}, R) \qquad \qquad \mathbf{\tilde{Q}} = GQG^{\mathsf{T}} \qquad \mathbf{\tilde{C}} = CG^{-1}$$

$$\Rightarrow \qquad P(\mathbf{\tilde{y}}_{t+1}|\mathbf{\tilde{y}}_t) = \mathcal{N}\left(GAG^{-1}G\mathbf{y}_t, GQG^{\mathsf{T}}\right) = \mathcal{N}\left(\mathbf{\tilde{A}}\mathbf{\tilde{y}}_t, \mathbf{\tilde{Q}}\right)$$

$$P(\mathbf{x}_t|\mathbf{\tilde{y}}_t) = \mathcal{N}\left(CG^{-1}G\mathbf{y}, R\right) = \mathcal{N}\left(\mathbf{\tilde{C}}\mathbf{\tilde{y}}, R\right)$$

Recall that the FA likelihood is conserved with respect to orthogonal transformations of y:

$$P(\mathbf{y}) = \mathcal{N} (\mathbf{0}, l) \qquad \qquad \tilde{\mathbf{y}} = U\mathbf{y} \qquad \qquad P(\tilde{\mathbf{y}}) = \mathcal{N} \left(U\mathbf{0}, UlU^{\mathsf{T}}\right) = \mathcal{N} (\mathbf{0}, l) \\ P(\mathbf{x}|\mathbf{y}) = \mathcal{N} \left(\Lambda \mathbf{y}, \Psi\right) \qquad \qquad \tilde{\Lambda} = \Lambda U^{\mathsf{T}} \qquad \Rightarrow \qquad P(\tilde{\mathbf{y}}) = \mathcal{N} \left(\Lambda U^{\mathsf{T}} U\mathbf{y}, \Psi\right) = \mathcal{N} \left(\tilde{\Lambda} \tilde{\mathbf{y}}, \Psi\right)$$

Similarly, a mixture model is invariant to permutations of the latent.

The LGSSM likelihood is conserved with respect to any invertible transform of the latent:

and the HMM is invariant to permutations (and to relaxations into something called an observable operator model).

Besides the degenerate solutions, both LGSSMs and HMMs likelihood functions tend to have multiple local maxima. This complicates any ML-based approach to learning (including EM and gradient methods).

Besides the degenerate solutions, both LGSSMs and HMMs likelihood functions tend to have multiple local maxima. This complicates any ML-based approach to learning (including EM and gradient methods).

Strategies:

Restart EM/gradient ascent from different (often random) initial parameter values.

Besides the degenerate solutions, both LGSSMs and HMMs likelihood functions tend to have multiple local maxima. This complicates any ML-based approach to learning (including EM and gradient methods).

Strategies:

- Restart EM/gradient ascent from different (often random) initial parameter values.
- Initialise output parameters with a stationary model:

$$\label{eq:pca} \begin{split} \text{PCA} &\to \text{FA} \to \text{LGSSM} \\ \text{k-means} \to \text{mixture} \to \text{HMM} \end{split}$$

Besides the degenerate solutions, both LGSSMs and HMMs likelihood functions tend to have multiple local maxima. This complicates any ML-based approach to learning (including EM and gradient methods).

Strategies:

- Restart EM/gradient ascent from different (often random) initial parameter values.
- Initialise output parameters with a stationary model:

 $PCA \rightarrow FA \rightarrow LGSSM$ k-means $\rightarrow mixture \rightarrow HMM$

Stochastic gradient methods and momentum. "Overstepping" EM.

Besides the degenerate solutions, both LGSSMs and HMMs likelihood functions tend to have multiple local maxima. This complicates any ML-based approach to learning (including EM and gradient methods).

Strategies:

- Restart EM/gradient ascent from different (often random) initial parameter values.
- Initialise output parameters with a stationary model:

 $PCA \rightarrow FA \rightarrow LGSSM$ k-means $\rightarrow mixture \rightarrow HMM$

- Stochastic gradient methods and momentum. "Overstepping" EM.
- Deterministic annealing.
The likelihood landscape

Besides the degenerate solutions, both LGSSMs and HMMs likelihood functions tend to have multiple local maxima. This complicates any ML-based approach to learning (including EM and gradient methods).

Strategies:

- Restart EM/gradient ascent from different (often random) initial parameter values.
- Initialise output parameters with a stationary model:

 $PCA \rightarrow FA \rightarrow LGSSM$ k-means $\rightarrow mixture \rightarrow HMM$

- Stochastic gradient methods and momentum. "Overstepping" EM.
- Deterministic annealing.
- Non-ML learning (spectral methods).

Is there a zero-noise limit for an LGSSM analagous to PCA?

Is there a zero-noise limit for an LGSSM analagous to PCA? We can take:

$$\begin{aligned} A &= \text{diag} \left[a_1 \dots a_K \right] \to I \quad (\text{with } a_1 \leq a_2 \leq \dots \leq a_K \leq 1) \\ Q &= I - AA^T \to 0 \quad (\text{setting the stationary latent covariance to } I) \\ R \to 0 \,. \end{aligned}$$

This limit yields a spectral algorithm called Slow Feature Analysis (or SFA).

Is there a zero-noise limit for an LGSSM analagous to PCA? We can take:

$$\begin{split} A &= \text{diag} \left[a_1 \dots a_K \right] \to I \quad (\text{with } a_1 \leq a_2 \leq \dots \leq a_K \leq 1) \\ Q &= I - AA^T \to 0 \quad (\text{setting the stationary latent covariance to } I) \\ R \to 0 \,. \end{split}$$

This limit yields a spectral algorithm called Slow Feature Analysis (or SFA).

SFA is conventionally defined as slowness pursuit (cf PCA and variance pursuit): Given zero-mean series $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ find a $K \times D$ matrix $W(=C^T)$ such that $\mathbf{y}_t = W\mathbf{x}_t$ changes as slowly as possible.

Is there a zero-noise limit for an LGSSM analagous to PCA? We can take:

$$\begin{aligned} A &= \text{diag} \left[a_1 \dots a_K \right] \to I \quad (\text{with } a_1 \leq a_2 \leq \dots \leq a_K \leq 1) \\ Q &= I - AA^T \to 0 \quad (\text{setting the stationary latent covariance to } I) \\ R \to 0 \,. \end{aligned}$$

This limit yields a spectral algorithm called Slow Feature Analysis (or SFA).

SFA is conventionally defined as slowness pursuit (cf PCA and variance pursuit): Given zero-mean series $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ find a $K \times D$ matrix $W(=C^T)$ such that $\mathbf{y}_t = W\mathbf{x}_t$ changes as slowly as possible.

Specifically, find

$$W = \operatorname*{argmin}_{W} \sum_{t} \|\mathbf{y}_t - \mathbf{y}_{t-1}\|^2 \quad ext{subject to} \quad \sum_{t} \mathbf{y}_t \mathbf{y}_t^\mathsf{T} = I.$$

The variance constraint prevents the trivial solutions W = 0 and $W = \mathbf{1}\mathbf{w}^{T}$.

Is there a zero-noise limit for an LGSSM analagous to PCA? We can take:

$$\begin{aligned} A &= \text{diag} \left[a_1 \dots a_K \right] \to I \quad (\text{with } a_1 \leq a_2 \leq \dots \leq a_K \leq 1) \\ Q &= I - AA^T \to 0 \quad (\text{setting the stationary latent covariance to } I) \\ R \to 0 \,. \end{aligned}$$

This limit yields a spectral algorithm called Slow Feature Analysis (or SFA).

SFA is conventionally defined as slowness pursuit (cf PCA and variance pursuit): Given zero-mean series $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ find a $K \times D$ matrix $W(=C^T)$ such that $\mathbf{y}_t = W\mathbf{x}_t$ changes as slowly as possible.

Specifically, find

$$W = \underset{W}{\operatorname{argmin}} \sum_{t} \|\mathbf{y}_{t} - \mathbf{y}_{t-1}\|^{2} \text{ subject to } \sum_{t} \mathbf{y}_{t} \mathbf{y}_{t}^{\mathsf{T}} = I.$$

The variance constraint prevents the trivial solutions W = 0 and $W = \mathbf{1}\mathbf{w}^{\mathsf{T}}$. *W* can be found by solving the generalised eigenvalue problem

$$W\!A = \Omega W\!B$$
 where $A = \sum_t (\mathbf{x}_t - \mathbf{x}_{t-1}) (\mathbf{x}_t - \mathbf{x}_{t-1})^{\mathsf{T}}$ and $B = \sum_t \mathbf{x}_t \mathbf{x}_t^{\mathsf{T}}$.

See http://www.gatsby.ucl.ac.uk/~maneesh/papers/turner-sahani-2007-sfa.pdf.

What if we don't want the $A \rightarrow I$ limit?

What if we don't want the $A \rightarrow I$ limit?

The limit is necessary to make the spectral estimate and the (limiting) ML estimates coincide. However, if we give up on ML, then there are general spectral estimates available.

What if we don't want the $A \rightarrow I$ limit?

The limit is necessary to make the spectral estimate and the (limiting) ML estimates coincide. However, if we give up on ML, then there are general spectral estimates available.

The ML parameter estimates are defined:

$$heta^{\mathsf{ML}} = \operatorname*{argmax}_{ heta} \log P(\mathcal{X}| heta)$$

and as you found, if $P \in ExpFam$ with sufficient statistic T then

$$\langle T(\mathbf{x}) \rangle_{\theta^{ML}} = \frac{1}{N} \sum_{i} T(\mathbf{x}_{i})$$
 (moment matching).

What if we don't want the $A \rightarrow I$ limit?

The limit is necessary to make the spectral estimate and the (limiting) ML estimates coincide. However, if we give up on ML, then there are general spectral estimates available.

The ML parameter estimates are defined:

$$heta^{\mathsf{ML}} = \operatorname*{argmax}_{ heta} \log P(\mathcal{X}| heta)$$

and as you found, if $P \in ExpFam$ with sufficient statistic T then

$$\langle T(\mathbf{x}) \rangle_{\theta^{ML}} = \frac{1}{N} \sum_{i} T(\mathbf{x}_{i})$$
 (moment matching).

We can generalise this condition to arbitrary *T* even if $P \notin ExpFam$.

That is, we find estimate θ^* with

$$\langle T(\mathbf{x}) \rangle_{\theta^*} = \frac{1}{N} \sum_i T(\mathbf{x}_i) \text{ or } \theta^* = \underset{\theta}{\operatorname{argmin}} \| \langle T(\mathbf{x}) \rangle_{\theta} - \frac{1}{N} \sum_i T(\mathbf{x}_i) \|_{\mathcal{C}}$$

What if we don't want the $A \rightarrow I$ limit?

The limit is necessary to make the spectral estimate and the (limiting) ML estimates coincide. However, if we give up on ML, then there are general spectral estimates available.

The ML parameter estimates are defined:

$$heta^{\mathsf{ML}} = \operatorname*{argmax}_{ heta} \log P(\mathcal{X}| heta)$$

and as you found, if $P \in ExpFam$ with sufficient statistic T then

$$\langle T(\mathbf{x}) \rangle_{\theta^{ML}} = \frac{1}{N} \sum_{i} T(\mathbf{x}_{i})$$
 (moment matching).

We can generalise this condition to arbitrary *T* even if $P \notin ExpFam$.

That is, we find estimate θ^* with

$$\langle T(\mathbf{x}) \rangle_{\theta^*} = \frac{1}{N} \sum_i T(\mathbf{x}_i) \text{ or } \theta^* = \underset{\theta}{\operatorname{argmin}} \| \langle T(\mathbf{x}) \rangle_{\theta} - \frac{1}{N} \sum_i T(\mathbf{x}_i) \|_{\mathcal{C}}$$

Judicious choice of T and metric C might make solution unique (no local optima) and consistent (correct given infinite within-model data).





$$\boldsymbol{M}_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle$$



$$\boldsymbol{M}_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left\langle \mathbf{x}_{t+\tau} (\boldsymbol{C} \mathbf{y}_{t} + \boldsymbol{\eta})^{\mathsf{T}} \right\rangle$$



$$\boldsymbol{M}_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left\langle \mathbf{x}_{t+\tau} \mathbf{y}_{t}^{\mathsf{T}} \right\rangle \boldsymbol{C}^{\mathsf{T}}$$



$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left\langle \mathbf{x}_{t+\tau} \mathbf{y}_{t}^{\mathsf{T}} \right\rangle C^{\mathsf{T}} = \left\langle (CA^{\tau} \mathbf{y}_{t} + \eta) \mathbf{y}_{t}^{\mathsf{T}} \right\rangle C^{\mathsf{T}}$$



$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left\langle \mathbf{x}_{t+\tau} \mathbf{y}_{t}^{\mathsf{T}} \right\rangle C^{\mathsf{T}} = C A^{\tau} \left\langle \mathbf{y}_{t} \mathbf{y}_{t}^{\mathsf{T}} \right\rangle C^{\mathsf{T}}$$



$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left\langle \mathbf{x}_{t+\tau} \mathbf{y}_{t}^{\mathsf{T}} \right\rangle C^{\mathsf{T}} = C A^{\tau} \left\langle \mathbf{y}_{t} \mathbf{y}_{t}^{\mathsf{T}} \right\rangle C^{\mathsf{T}} = C A^{\tau} \Pi C^{\mathsf{T}}$$



$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left\langle \mathbf{x}_{t+\tau} \mathbf{y}_{t}^{\mathsf{T}} \right\rangle C^{\mathsf{T}} = C A^{\tau} \left\langle \mathbf{y}_{t} \mathbf{y}_{t}^{\mathsf{T}} \right\rangle C^{\mathsf{T}} = C A^{\tau} \Pi C^{\mathsf{T}}$$

 $H \equiv \left\langle \mathbf{x}_t^+, \mathbf{x}_t^{-\mathsf{T}} \right\rangle$



$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left\langle \mathbf{x}_{t+\tau} \mathbf{y}_{t}^{\mathsf{T}} \right\rangle C^{\mathsf{T}} = C A^{\tau} \left\langle \mathbf{y}_{t} \mathbf{y}_{t}^{\mathsf{T}} \right\rangle C^{\mathsf{T}} = C A^{\tau} \Pi C^{\mathsf{T}}$$

$$H \equiv \left\langle \mathbf{x}_{t}^{+}, \mathbf{x}_{t}^{-\mathsf{T}} \right\rangle = \begin{bmatrix} M_{1} & M_{2} & \cdots & M_{L} \\ M_{2} & M_{3} & & & \\ \vdots & & \vdots & & \vdots \\ M_{L} & & \cdots & M_{2L-1} \end{bmatrix} = \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{L-1} \end{bmatrix} \begin{bmatrix} A \Pi C^{\mathsf{T}} & \cdots & A^{L} \Pi C^{\mathsf{T}} \end{bmatrix}$$



$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left\langle \mathbf{x}_{t+\tau} \mathbf{y}_{t}^{\mathsf{T}} \right\rangle C^{\mathsf{T}} = C A^{\tau} \left\langle \mathbf{y}_{t} \mathbf{y}_{t}^{\mathsf{T}} \right\rangle C^{\mathsf{T}} = C A^{\tau} \Pi C^{\mathsf{T}}$$

$$H \equiv \left\langle \mathbf{x}_{t}^{+}, \mathbf{x}_{t}^{-\mathsf{T}} \right\rangle = \begin{bmatrix} M_{1} & M_{2} & \cdots & M_{L} \\ M_{2} & M_{3} & & \\ \vdots & & \vdots \\ M_{L} & \cdots & M_{2L-1} \end{bmatrix} = \begin{bmatrix} C \\ CA \\ \Xi \\ \vdots \\ CA^{L-1} \end{bmatrix} \begin{bmatrix} A \Pi C^{\mathsf{T}} & \widehat{\Lambda} & A^{L} \Pi C^{\mathsf{T}} \end{bmatrix}$$
$$K \times LD$$



$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left\langle \mathbf{x}_{t+\tau} \mathbf{y}_{t}^{\mathsf{T}} \right\rangle C^{\mathsf{T}} = C A^{\tau} \left\langle \mathbf{y}_{t} \mathbf{y}_{t}^{\mathsf{T}} \right\rangle C^{\mathsf{T}} = C A^{\tau} \Pi C^{\mathsf{T}}$$

$$H \equiv \left\langle \mathbf{x}_{t}^{+}, \mathbf{x}_{t}^{-\mathsf{T}} \right\rangle = \begin{bmatrix} M_{1} & M_{2} & \cdots & M_{L} \\ M_{2} & M_{3} & & & \\ \vdots & & \vdots \\ M_{L} & \cdots & M_{2L-1} \end{bmatrix} = \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{L-1} \end{bmatrix} \begin{bmatrix} A \Box C^{\mathsf{T}} & \ddots & A^{L} \Box C^{\mathsf{T}} \end{bmatrix}$$
$$K \times LD$$

Off-diagonal correlation unaffected by noise. $SVD(\frac{1}{T} \sum \mathbf{x}_t^+ \mathbf{x}_t^{-T})$ yields least-squares estimates of Ξ and Υ . Regression between blocks of Ξ yields \widehat{A} and \widehat{C} .



Now consider an HMM with discrete output symbols.



Now consider an HMM with discrete output symbols. The likelihood can be written:

$$P(x_{1:T}|\pi, \Phi, A) = \sum_{i} \pi_i A_i(x_1) \sum_{j} \Phi_{ji} A_j(x_2) \sum_{k} \Phi_{kj} A_k(x_3) \dots$$



Now consider an HMM with discrete output symbols. The likelihood can be written:

$$P(x_{1:T}|\pi, \Phi, A) = \sum_{i} \pi_{i} A_{i}(x_{1}) \sum_{j} \Phi_{ji} A_{j}(x_{2}) \sum_{k} \Phi_{kj} A_{k}(x_{3}) \dots$$
$$= \pi^{\mathsf{T}} A_{x_{1}} \Phi^{\mathsf{T}} A_{x_{2}} \Phi^{\mathsf{T}} A_{x_{3}} \dots \mathbf{1} \qquad [A_{x_{1}} = \operatorname{diag} [A_{1}(x_{1}), A_{2}(x_{1}), \dots]]$$



Now consider an HMM with discrete output symbols. The likelihood can be written:

$$P(x_{1:T}|\pi, \Phi, A) = \sum_{i} \pi_{i} A_{i}(x_{1}) \sum_{j} \Phi_{ji} A_{j}(x_{2}) \sum_{k} \Phi_{kj} A_{k}(x_{3}) \dots$$

= $\pi^{T} A_{x_{1}} \Phi^{T} A_{x_{2}} \Phi^{T} A_{x_{3}} \dots \mathbf{1} \qquad [A_{x_{1}} = \text{diag} [A_{1}(x_{1}), A_{2}(x_{1}), \dots]]$
= $\mathbf{1}^{T} O_{x_{T}} O_{x_{T-1}} \dots O_{x_{1}} \pi$

where $O_a = \Phi A_a$ is a "propagation operator" on the latent belief that depends on observation.



Now consider an HMM with discrete output symbols. The likelihood can be written:

$$P(x_{1:T}|\pi, \Phi, A) = \sum_{i} \pi_{i} A_{i}(x_{1}) \sum_{j} \Phi_{ji} A_{j}(x_{2}) \sum_{k} \Phi_{kj} A_{K}(x_{3}) \dots$$

= $\pi^{T} A_{x_{1}} \Phi^{T} A_{x_{2}} \Phi^{T} A_{x_{3}} \dots \mathbf{1} \qquad [A_{x_{1}} = \text{diag} [A_{1}(x_{1}), A_{2}(x_{1}), \dots]]$
= $\mathbf{1}^{T} O_{x_{T}} O_{x_{T-1}} \dots O_{x_{1}} \pi$

where $O_a = \Phi A_a$ is a "propagation operator" on the latent belief that depends on observation.

Observable operator model (OOM) representation.



Now consider an HMM with discrete output symbols. The likelihood can be written:

$$P(x_{1:T}|\pi, \Phi, A) = \sum_{i} \pi_{i} A_{i}(x_{1}) \sum_{j} \Phi_{ji} A_{j}(x_{2}) \sum_{k} \Phi_{kj} A_{K}(x_{3}) \dots$$

= $\pi^{T} A_{x_{1}} \Phi^{T} A_{x_{2}} \Phi^{T} A_{x_{3}} \dots \mathbf{1} \qquad [A_{x_{1}} = \text{diag} [A_{1}(x_{1}), A_{2}(x_{1}), \dots]]$
= $\mathbf{1}^{T} O_{x_{T}} O_{x_{T-1}} \dots O_{x_{1}} \pi$

where $O_a = \Phi A_a$ is a "propagation operator" on the latent belief that depends on observation.

Observable operator model (OOM) representation.

OOMs with arbitrary O matrices describe a larger class of distributions that HMMs.



Now consider an HMM with discrete output symbols. The likelihood can be written:

$$P(x_{1:T}|\pi, \Phi, A) = \sum_{i} \pi_{i} A_{i}(x_{1}) \sum_{j} \Phi_{ji} A_{j}(x_{2}) \sum_{k} \Phi_{kj} A_{K}(x_{3}) \dots$$

= $\pi^{T} A_{x_{1}} \Phi^{T} A_{x_{2}} \Phi^{T} A_{x_{3}} \dots \mathbf{1} \qquad [A_{x_{1}} = \text{diag} [A_{1}(x_{1}), A_{2}(x_{1}), \dots]]$
= $\mathbf{1}^{T} O_{x_{T}} O_{x_{T-1}} \dots O_{x_{1}} \pi$

where $O_a = \Phi A_a$ is a "propagation operator" on the latent belief that depends on observation.

Observable operator model (OOM) representation.

- OOMs with arbitrary O matrices describe a larger class of distributions that HMMs.
- Not easy to normalise or even guarantee all assigned "probabilities" are positive.



Now consider an HMM with discrete output symbols. The likelihood can be written:

$$P(x_{1:T}|\pi, \Phi, A) = \sum_{i} \pi_{i} A_{i}(x_{1}) \sum_{j} \Phi_{ji} A_{j}(x_{2}) \sum_{k} \Phi_{kj} A_{K}(x_{3}) \dots$$

= $\pi^{T} A_{x_{1}} \Phi^{T} A_{x_{2}} \Phi^{T} A_{x_{3}} \dots \mathbf{1} \qquad [A_{x_{1}} = \text{diag} [A_{1}(x_{1}), A_{2}(x_{1}), \dots]]$
= $\mathbf{1}^{T} O_{x_{T}} O_{x_{T-1}} \dots O_{x_{1}} \pi$

where $O_a = \Phi A_a$ is a "propagation operator" on the latent belief that depends on observation.

Observable operator model (OOM) representation.

- OOMs with arbitrary O matrices describe a larger class of distributions that HMMs.
- Not easy to normalise or even guarantee all assigned "probabilities" are positive.
- Degenerate with respect to similarity transform $\tilde{O} = GOG^{-1}$.





$$\boldsymbol{M}_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle$$



$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = [P(x_{t+\tau}=i, x_{t}=j)]$$



$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left[\left\langle P(x_{t+\tau}=i, x_{t}=j|s_{t}) \right\rangle_{s_{t}} \right]$$



$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left[\left\langle P(x_{t+\tau} = i | s_{t}) P(x_{t} = j | s_{t}) \right\rangle_{s_{t}} \right]$$



$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left[\left\langle P(x_{t+\tau} = i | s_{t}) P(x_{t} = j | s_{t}) \right\rangle_{s_{t}} \right] = \left\langle A \Phi^{\tau} \mathbf{s}_{t} (A \mathbf{s}_{t})^{\mathsf{T}} \right\rangle_{s_{t}}$$


Write $\mathbf{x}_t = [\delta(x_t=i)]; \ \mathbf{s}_t = [\delta(s_t=i)]; \ A = [A_{ij}] = [P(x_t=i|s_t=j)]$

$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left[\left\langle P(x_{t+\tau} = i | s_{t}) P(x_{t} = j | s_{t}) \right\rangle_{s_{t}} \right] = A \Phi^{\tau} \left\langle \mathbf{s}_{t} \mathbf{s}_{t}^{\mathsf{T}} \right\rangle A^{\mathsf{T}}$$



Write $\mathbf{x}_t = [\delta(x_t=i)]; \ \mathbf{s}_t = [\delta(s_t=i)]; \ A = [A_{ij}] = [P(x_t=i|s_t=j)]$

$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left[\left\langle P(x_{t+\tau} = i | s_{t}) P(x_{t} = j | s_{t}) \right\rangle_{s_{t}} \right] = A \Phi^{\tau} \left\langle \mathbf{s}_{t} \mathbf{s}_{t}^{\mathsf{T}} \right\rangle A^{\mathsf{T}} = A \Phi^{\tau} \Pi A^{\mathsf{T}}$$



Write $\mathbf{x}_t = [\delta(x_t=i)]; \ \mathbf{s}_t = [\delta(s_t=i)]; \ A = [A_{ij}] = [P(x_t=i|s_t=j)]$

$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left[\left\langle P(x_{t+\tau} = i | s_{t}) P(x_{t} = j | s_{t}) \right\rangle_{s_{t}} \right] = A \Phi^{\tau} \left\langle \mathbf{s}_{t} \mathbf{s}_{t}^{\mathsf{T}} \right\rangle A^{\mathsf{T}} = A \Phi^{\tau} \Pi A^{\mathsf{T}}$$
$$H \equiv \left\langle \mathbf{x}_{t}^{+}, \mathbf{x}_{t}^{-\mathsf{T}} \right\rangle = \begin{bmatrix} M_{1} & M_{2} & \cdots & M_{L} \\ M_{2} & M_{3} & & \\ \vdots & & \vdots \\ M_{L} & \cdots & M_{2L-1} \end{bmatrix} = \begin{bmatrix} A \\ A \Phi \\ \vdots \\ A \Phi^{L-1} \end{bmatrix} \left[\Phi \Pi A^{\mathsf{T}} & \cdots & \Phi^{L} \Pi A^{\mathsf{T}} \right]$$

(The conventional algorithm is written slightly differently, but follows similar logic).



Write $\mathbf{x}_t = [\delta(x_t=i)]; \ \mathbf{s}_t = [\delta(s_t=i)]; \ A = [A_{ij}] = [P(x_t=i|s_t=j)]$

$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left[\left\langle P(x_{t+\tau} = i | s_{t}) P(x_{t} = j | s_{t}) \right\rangle_{s_{t}} \right] = A \Phi^{\tau} \left\langle \mathbf{s}_{t} \mathbf{s}_{t}^{\mathsf{T}} \right\rangle A^{\mathsf{T}} = A \Phi^{\tau} \Pi A^{\mathsf{T}}$$
$$H \equiv \left\langle \mathbf{x}_{t}^{+}, \mathbf{x}_{t}^{-\mathsf{T}} \right\rangle = \begin{bmatrix} M_{1} & M_{2} & \cdots & M_{L} \\ M_{2} & M_{3} & & \\ \vdots & & \vdots \\ M_{L} & \cdots & M_{2L-1} \end{bmatrix} = \begin{bmatrix} A \\ A \Phi \\ \vdots \\ A \Phi^{L-1} \end{bmatrix} \begin{bmatrix} \Phi \Pi A^{\mathsf{T}} & \cdots & \Phi^{L} \Pi A^{\mathsf{T}} \end{bmatrix}$$
$$LD \times LD \qquad LD \times K$$

(The conventional algorithm is written slightly differently, but follows similar logic). $\widehat{\Phi}$ and \widehat{A} can be recovered as before ...



Write $\mathbf{x}_t = [\delta(x_t=i)]; \ \mathbf{s}_t = [\delta(s_t=i)]; \ A = [A_{ij}] = [P(x_t=i|s_t=j)]$

$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left[\left\langle P(x_{t+\tau} = i | s_{t}) P(x_{t} = j | s_{t}) \right\rangle_{s_{t}} \right] = A \Phi^{\tau} \left\langle \mathbf{s}_{t} \mathbf{s}_{t}^{\mathsf{T}} \right\rangle A^{\mathsf{T}} = A \Phi^{\tau} \Pi A^{\mathsf{T}}$$
$$H \equiv \left\langle \mathbf{x}_{t}^{+}, \mathbf{x}_{t}^{-\mathsf{T}} \right\rangle = \begin{bmatrix} M_{1} & M_{2} & \cdots & M_{L} \\ M_{2} & M_{3} & & \\ \vdots & & \vdots \\ M_{L} & \cdots & M_{2L-1} \end{bmatrix} = \begin{bmatrix} A \\ A \Phi \\ \vdots \\ A \Phi^{L-1} \end{bmatrix} \begin{bmatrix} \Phi \Pi A^{\mathsf{T}} & \cdots & \Phi^{L} \Pi A^{\mathsf{T}} \end{bmatrix}$$
$$LD \times LD \qquad LD \times K$$

(The conventional algorithm is written slightly differently, but follows similar logic). $\widehat{\Phi}$ and \widehat{A} can be recovered as before upto arbitrary invertible transform of $\mathbf{s} \to OOM$.



Write $\mathbf{x}_t = [\delta(x_t=i)]; \ \mathbf{s}_t = [\delta(s_t=i)]; \ A = [A_{ij}] = [P(x_t=i|s_t=j)]$

$$M_{\tau} \equiv \left\langle \mathbf{x}_{t+\tau} \mathbf{x}_{t}^{\mathsf{T}} \right\rangle = \left[\left\langle P(x_{t+\tau}=i|s_{t})P(x_{t}=j|s_{t}) \right\rangle_{s_{t}} \right] = A\Phi^{\tau} \left\langle \mathbf{s}_{t} \mathbf{s}_{t}^{\mathsf{T}} \right\rangle A^{\mathsf{T}} = A\Phi^{\tau} \Pi A^{\mathsf{T}}$$
$$H \equiv \left\langle \mathbf{x}_{t}^{+}, \mathbf{x}_{t}^{-\mathsf{T}} \right\rangle = \begin{bmatrix} M_{1} & M_{2} & \cdots & M_{L} \\ M_{2} & M_{3} & & \\ \vdots & & \vdots \\ M_{L} & \cdots & M_{2L-1} \end{bmatrix} = \begin{bmatrix} A \\ A\Phi \\ \vdots \\ A\Phi^{L-1} \end{bmatrix} \begin{bmatrix} \Phi \Pi A^{\mathsf{T}} & \cdots & \Phi^{L} \Pi A^{\mathsf{T}} \end{bmatrix}$$
$$LD \times LD \qquad LD \times K$$

(The conventional algorithm is written slightly differently, but follows similar logic). $\widehat{\Phi}$ and \widehat{A} can be recovered as before upto arbitrary invertible transform of $\mathbf{s} \to OOM$. 'Projection' to HHM space possible, but amplifies estimation errors.

Spectral learning:

Spectral learning:

Efficient closed-form solution.

Maximum likelihood learning:

Requires iterative maximisation.

Spectral learning:

Efficient closed-form solution finds global optimum.

- Requires iterative maximisation.
- Many local optima.

Spectral learning:

- Efficient closed-form solution finds global optimum.
- Consistent (recovers true parameters upto degeneracies from infinite within-model data).

- Requires iterative maximisation.
- Many local optima.
- Consistent and asymptotically efficient (if the global maximum can be found).

Spectral learning:

- Efficient closed-form solution finds global optimum.
- Consistent (recovers true parameters upto degeneracies from infinite within-model data).
- Eigen-/singular-value spectrum clue to latent dimensionality.

- Requires iterative maximisation.
- Many local optima.
- Consistent and asymptotically efficient (if the global maximum can be found).

Spectral learning:

- Efficient closed-form solution finds global optimum.
- Consistent (recovers true parameters upto degeneracies from infinite within-model data).
- Eigen-/singular-value spectrum clue to latent dimensionality.
- Assumes stationarity. May be inappropriate for short sequences.

- Requires iterative maximisation.
- Many local optima.
- Consistent and asymptotically efficient (if the global maximum can be found).

Spectral learning:

- Efficient closed-form solution finds global optimum.
- Consistent (recovers true parameters upto degeneracies from infinite within-model data).
- Eigen-/singular-value spectrum clue to latent dimensionality.
- Assumes stationarity. May be inappropriate for short sequences.
- HMM learning returns OOM parameters may not correspond to any HMM or indeed to proper probabilistic model.

- Requires iterative maximisation.
- Many local optima.
- Consistent and asymptotically efficient (if the global maximum can be found).

Spectral learning:

- Efficient closed-form solution finds global optimum.
- Consistent (recovers true parameters upto degeneracies from infinite within-model data).
- Eigen-/singular-value spectrum clue to latent dimensionality.
- Assumes stationarity. May be inappropriate for short sequences.
- HMM learning returns OOM parameters may not correspond to any HMM or indeed to proper probabilistic model.
- Not easily generalised to nonlinear or more complex models (but see http://www.gatsby.ucl.ac.uk/~maneesh/papers/buesing-etal-2012-nips.pdf).

- Requires iterative maximisation.
- Many local optima.
- Consistent and asymptotically efficient (if the global maximum can be found).

Spectral learning:

- Efficient closed-form solution finds global optimum.
- Consistent (recovers true parameters upto degeneracies from infinite within-model data).
- Eigen-/singular-value spectrum clue to latent dimensionality.
- Assumes stationarity. May be inappropriate for short sequences.
- HMM learning returns OOM parameters may not correspond to any HMM or indeed to proper probabilistic model.
- Not easily generalised to nonlinear or more complex models (but see http://www.gatsby.ucl.ac.uk/~maneesh/papers/buesing-etal-2012-nips.pdf).
- In practice, error in recovered parameters is often large.

- Requires iterative maximisation.
- Many local optima.
- Consistent and asymptotically efficient (if the global maximum can be found).
- Generalises to "principled" approximate algorithms for nonlinear or complex models.

Spectral learning:

- Efficient closed-form solution finds global optimum.
- Consistent (recovers true parameters upto degeneracies from infinite within-model data).
- Eigen-/singular-value spectrum clue to latent dimensionality.
- Assumes stationarity. May be inappropriate for short sequences.
- HMM learning returns OOM parameters may not correspond to any HMM or indeed to proper probabilistic model.
- Not easily generalised to nonlinear or more complex models (but see http://www.gatsby.ucl.ac.uk/~maneesh/papers/buesing-etal-2012-nips.pdf).
- In practice, error in recovered parameters is often large.
- Often valuable as initialisation for ML methods.

- Requires iterative maximisation.
- Many local optima.
- Consistent and asymptotically efficient (if the global maximum can be found).
- Generalises to "principled" approximate algorithms for nonlinear or complex models.

Recognition (classification) with HMMs

Multiple HMM models:

- 1. train one HMM for each class (requires each sequence to be labelled by the class)
- 2. evaluate the probability of an unknown sequence under each HMM
- 3. classify the unknown sequence by the HMM which gave it the highest likelihood



Recognition (labelling) with HMMs



Use a single HMM to label sequences:

- 1. train a single HMM on sequences of data $\mathbf{x}_1, \ldots, \mathbf{x}_T$ and corresponding labels s_1, \ldots, s_T .
- 2. On an unlabelled test sequence, compute the posterior distribution over label sequences $P(s_1, \ldots, s_T | \mathbf{x}_1, \ldots, \mathbf{x}_T)$.
- 3. Return the label sequence either with highest expected number of correct states, or highest probability under the posterior (Viterbi).

Recognition (labelling) with HMMs



Use a single HMM to label sequences:

- 1. train a single HMM on sequences of data $\mathbf{x}_1, \ldots, \mathbf{x}_T$ and corresponding labels s_1, \ldots, s_T .
- 2. On an unlabelled test sequence, compute the posterior distribution over label sequences $P(s_1, \ldots, s_T | \mathbf{x}_1, \ldots, \mathbf{x}_T)$.
- 3. Return the label sequence either with highest expected number of correct states, or highest probability under the posterior (Viterbi).

Models the whole joint distribution $P(\mathbf{x}_{1:T}, \mathbf{s}_{1:T})$, but only uses $P(\mathbf{s}_{1:T} | \mathbf{x}_{1:T})$.

May be more accurate and more efficient use of data to model P(s₁:T|x₁:T) directly.

Recognition (labelling) with HMMs



Use a single HMM to label sequences:

- 1. train a single HMM on sequences of data $\mathbf{x}_1, \ldots, \mathbf{x}_T$ and corresponding labels s_1, \ldots, s_T .
- 2. On an unlabelled test sequence, compute the posterior distribution over label sequences $P(s_1, \ldots, s_T | \mathbf{x}_1, \ldots, \mathbf{x}_T)$.
- 3. Return the label sequence either with highest expected number of correct states, or highest probability under the posterior (Viterbi).

Models the whole joint distribution $P(\mathbf{x}_{1:T}, \mathbf{s}_{1:T})$, but only uses $P(\mathbf{s}_{1:T} | \mathbf{x}_{1:T})$.

- May be more accurate and more efficient use of data to model $P(s_{1:T}|x_{1:T})$ directly.
- This leads to a model called a Conditional Random Field.

Conditional distribution in a HMM

Conditional distribution over label sequences of a HMM:

$$\begin{split} \mathcal{P}(\mathbf{s}_{1:T}|\mathbf{x}_{1:T},\theta) &= \frac{\mathcal{P}(\mathbf{s}_{1:T},\mathbf{x}_{1:T}|\theta)}{\sum_{s_{1:T}}\mathcal{P}(\mathbf{s}_{1:T},\mathbf{x}_{1:T}|\theta)} \\ &\propto \mathcal{P}(\mathbf{s}_{1}|\pi)\prod_{t=1}^{T-1}\mathcal{P}(\mathbf{s}_{t+1}|s_{t},\mathsf{T})\prod_{t=1}^{T}\mathcal{P}(\mathbf{x}_{t}|s_{t},\mathsf{A}) \\ &= \exp\left(\sum_{i}\delta(s_{1}=i)\log\pi_{i}+\sum_{t=1}^{T-1}\sum_{ij}\delta(s_{t}=i,s_{t+1}=j)\log\Phi_{ij}\right) \\ &+\sum_{t=1}^{T}\sum_{ik}\delta(s_{t}=i,\mathbf{x}_{t}=k)\log A_{ik}\right). \end{split}$$

This functional form gives a well-defined conditional distribution, even if we do not enforce the constraints

$$\Phi_{ij} \ge 0$$
 $\sum_{j} \Phi_{ij} = 1$

or the similar ones for π and A (cf. OOMs). The forward-backward algorithm can still be applied to compute the conditional distribution.

This is an example of a conditional random field.

Conditional random fields

Define two sets of functions: single label and label-pair functions. Single label functions:

$$f_i(\mathbf{s}_t, \mathbf{x}_t)$$
 for $i = 1, \dots, N$

Label-pair functions:

$$g_i(s_t, s_{t+1}, \mathbf{x}_t, \mathbf{x}_{t+1})$$
 for $j = 1, \dots, J$

Each function is associated with a real-valued parameter: λ_i , κ_i . A conditional random field defines a conditional distribution over $\mathbf{x}_{1:\tau}$ given $\mathbf{x}_{1:\tau}$ as follows:

$$P(\boldsymbol{s}_{1:T}|\boldsymbol{x}_{1:T}, \boldsymbol{\lambda}, \boldsymbol{\kappa}) \propto \exp\left(\sum_{t=1}^{T} \sum_{i} \lambda_{i} f_{i}(\boldsymbol{s}_{t}, \boldsymbol{x}_{t}) + \sum_{t=1}^{T-1} \sum_{j} \kappa_{j} g_{j}(\boldsymbol{s}_{t}, \boldsymbol{s}_{t+1}, \boldsymbol{x}_{t}, \boldsymbol{x}_{t+1})\right)$$

The forward-backward algorithm can be used to compute:

$$P(s_{t}|\mathbf{x}_{1:T}, \lambda, \kappa) \qquad P(s_{t}, s_{t+1}|\mathbf{x}_{1:T}, \lambda, \kappa) \qquad \underset{s_{1:T}}{\operatorname{argmax}} P(s_{1:T}|\mathbf{x}_{1:T}, \lambda, \kappa)$$

Factor graph notation for CRFs



Discriminative vs generative modelling

Labelled training data comes from a true underlying distribution $\tilde{P}(s_{1:T}, \mathbf{x}_{1:T})$. Generative modelling: train a HMM by maximizing likelihood:

$$heta_{\mathsf{Joint}} = \operatorname*{argmax}_{ heta} E_{\check{P}}[\log P(s_{1:T}, \mathbf{x}_{1:T} | \theta)]$$

(note do not need EM here, since no latent variables)

Discriminative modelling: train another HMM by maximizing conditional likelihood:

$$\theta_{\text{Cond}} = \operatorname*{argmax}_{\theta} E_{\tilde{P}}[\log P(s_{1:T} | \mathbf{x}_{1:T}, \theta)]$$

By construction:

$$\mathsf{E}_{\tilde{P}}[\log \mathsf{P}(s_{1:T}|\mathbf{x}_{1:T}, \theta_{\mathsf{Cond}})] \geq \mathsf{E}_{\tilde{P}}[\log \mathsf{P}(s_{1:T}|\mathbf{x}_{1:T}, \theta_{\mathsf{Joint}})]$$

If \tilde{P} belongs to model class, $P(\cdot|\theta_{\text{Joint}}) = \tilde{P}$ and equality holds.



Caveats:

- Underlying distribution P
 not usually in model class.
- training set differs from \tilde{P} .
- Overfitting easier in discriminative setting.
- Generative modelling often much simpler (fits each conditional probability separately, not iterative).

Major point of debate in machine learning.

Structured generalized linear models

$$P(\mathbf{s}_{1:T}|\mathbf{x}_{1:T}, \boldsymbol{\lambda}, \boldsymbol{\kappa}) \propto \exp\left(\sum_{t=1}^{T} \sum_{i} \lambda_{i} f_{i}(\mathbf{s}_{t}, \mathbf{x}_{t}) + \sum_{t=1}^{T-1} \sum_{j} \kappa_{j} g_{j}(\mathbf{s}_{t}, \mathbf{s}_{t+1}, \mathbf{x}_{t}, \mathbf{x}_{t+1})\right)$$

The conditional distribution over $s_{1:T}$ forms an exponential family parameterized by λ , κ and dependent on $\mathbf{x}_{1:T}$.

CRFs are a multivariate generalization of generalized linear models (GLMs).

The labels s_t in a CRF are not independently predicted, but they have a Markov property: $s_{1:t-1}$ is independent of $s_{t+1:T}$ given s_t and $\mathbf{x}_{1:T}$.

This allows efficient inference using the forward-backward algorithm.

CRFs are models for structured prediction (another major machine learning frontier).

CRFs are very flexible.

CRFs have found wide spread applications across a number of fields: natural language processing (part-of-speech tagging, named-entity recognition, coreference resolution), information retrieval (information extraction), computer vision (image segmentation, object recognition, depth perception), bioinformatics (protein structure prediction, gene finding)...

Learning CRFs

$$P(\mathbf{s}_{1:T}|\mathbf{x}_{1:T}, \boldsymbol{\lambda}, \boldsymbol{\kappa}) \propto \exp\left(\sum_{t=1}^{T} \sum_{i} \lambda_{i} f_{i}(\mathbf{s}_{t}, \mathbf{x}_{t}) + \sum_{t=1}^{T-1} \sum_{j} \kappa_{j} g_{j}(\mathbf{s}_{t}, \mathbf{s}_{t+1}, \mathbf{x}_{t}, \mathbf{x}_{t+1})\right)$$

Given labelled data $\{s_{1:T}^{(c)}, \mathbf{x}_{1:T}^{(c)}\}_{c=1}^{N}$, we train CRFs by maximum likelihood:

$$\frac{\partial \sum_{c} \log P(\mathbf{s}_{1:T}^{(c)} | \mathbf{x}_{1:T}^{(c)}, \boldsymbol{\lambda}, \boldsymbol{\kappa})}{\partial \lambda_{i}} = \sum_{c=1}^{N} \sum_{t=1}^{T} f_{i}(\mathbf{s}_{t}^{(c)}, \mathbf{x}_{t}^{(c)}) - E_{P(\mathbf{s}_{1:T} | \mathbf{x}_{1:T}^{(c)})}[f_{i}(\mathbf{s}_{t}^{(c)}, \mathbf{x}_{t}^{(c)})]$$
$$\frac{\partial \sum_{c} \log P(\mathbf{s}_{1:T}^{(c)} | \mathbf{x}_{1:T}^{(c)}, \boldsymbol{\lambda}, \boldsymbol{\kappa})}{\partial \kappa_{j}} = \sum_{c=1}^{N} \sum_{t=1}^{T-1} g_{j}(\mathbf{s}_{t:t+1}^{(c)}, \mathbf{x}_{t:t+1}^{(c)}) - E_{P(\mathbf{s}_{1:T} | \mathbf{x}_{1:T}^{(c)})}[g_{j}(\mathbf{s}_{t:t+1}, \mathbf{x}_{t:t+1}^{(c)})]$$

There is no closed-form solution for the parameters, so we use gradient ascent instead.

Note: expectations are computed using the forward-backward algorithm.

The log likelihood is concave, so unlike EM we will get to global optimum (another major frontier in machine learning).

References

- Yu, B.M, Shenoy K.V. and Sahani M. (2004) Derivation of Kalman Filtering and Smoothing Equations. Working note. http://www.gatsby.ucl.ac.uk/~byron/derive_ks.pdf
- Ghahramani, Z. and Hinton, G.E. (1996) Parameter estimation for linear dynamical systems. University of Toronto Technical Report CRG-TR-96-2, 6 pages.
- Rabiner, L.R. (1989) A tutorial on Hidden Markov Models and selected applications in speech recognition. Proceedings of the IEEE 77 (2): 257286.
- Roweis, S. and Ghahramani, Z. (1999) A Unifying Review of Linear Gaussian Models. Neural Computation 11(2):305–345.
- Buesing, L., Macke, J.H. and Sahani, M. (2012) Spectral learning of linear dynamics from generalised-linear observations with application to neural population data. NIPS 25.
- Lafferty, J. and McCallum, A. and Pereira, F. (2001) Conditional random fields: Probabilistic models for segmenting and labeling sequence data. Proceedings of the 18th International Conference on Machine Learning.