## Probabilistic Graphical Models

## Spectral Learning for Graphical Models

Eric Xing<br>Lecture 24, April 14, 2014



Acknowledgement: slides drafted by Ankur Parikh

## Latent Variable Models



Sequence models



Ho. et al. 2012
Mixed membership models

## Latent Variable PCFG ${ }_{\text {Inatsuraki ie tal., 2005, }}$

Petrov et al. 2006]

## PCFG



## Latent Variable PCFG



## Learning Parameters (EM)



$$
\mathbb{P}\left[X_{1}, \ldots, X_{5}, H_{1}, \ldots, H_{5}\right]=\mathbb{P}\left[H_{1}\right] \prod_{i=2}^{5} \mathbb{P}\left[H_{i} \mid H_{i-1}\right] \prod_{i=1}^{5} \mathbb{P}\left[X_{i} \mid H_{i}\right]
$$

Since latent variables are not observed in the data, we have to use Expectation Maximization (EM) to learn parameters

- Slow
- Local Minima


## Spectral Learning

- Different paradigm of learning in latent variable models based on linear algebra
- Theoretically,
- Provably consistent
- Can offer deeper insight into the identifiability
- Practically,
- Local minima free
- As if now, performs comparably to EM with 10-100x speed-up
- Can also model non-Gaussian continuous data using kernels (usually performs much better than EM in this case)


## Related References

- Relevant works
- Hsu et al. 2009 - Spectral HMMs (also Bailly 2009)
- Siddiqi et al. 2009 - Features in Spectral Learning
- Parikh et al. 2011/2012 -Tensors to Generalize to Trees/Low Treewidth Graphs
- Cohen et al. 2012 I 2013 - Spectral Learning of latent PCFGs
- Will present it from "matrix factorization" view:
- Balle et al. 2012 - Connection between Spectral Learning / Hankel Matrix Factorization
- Song et al. 2013 - Spectral Learning as Hierarchical Tensor Decomposition


## Focusing on Prediction

- In many applications that use latent variable models, the end task is not to recover the latent states, but rather to use the model for prediction among observed variables.
- Dynamical Systems - Predict future given past



## Focusing on Prediction

- We will only be concerned with quantities related to the observed variables:

$$
\mathbb{P}\left[X_{1}, X_{2}, X_{3}, X_{4}, X_{5}\right]
$$

- We do not care about the latent variables explicitly.

- Do we still need EM to learn the parameters?


## But if we don't care about the latent variables....

- Why don't we just integrate them out?
- Because integrating them out results in a clique $)^{\circ}$



## Marginal Does Not Factorize


$\mathbb{P}\left[X_{1}, X_{2}, X_{3}, X_{4}, X_{5}\right]=\sum_{H_{1}, \ldots, H_{5}} \mathbb{P}\left[H_{1}\right] \mathbb{P}\left[H_{1}\right] \prod_{i=2}^{5} \mathbb{P}\left[H_{i} \mid H_{i-1}\right] \prod_{i=1}^{5} \mathbb{P}\left[X_{i} \mid H_{i}\right]$
Does not factorize due to the outer sum (Can somewhat distribute the sum, but doesn't solve problem)

## But isn't an HMM different from a clique?

- It depends on the number of latent states.
- Consider the following model.



## If H has only one state.....

- Then the observed variables are independent!



## What if H has many states?

- Let us say the observed variables each have $m$ states.
- Then if H has $m^{3}$ states then the latent model can be exactly equivalent to a clique (depending on how parameters are set).

- But what about all the other cases?


## The Question

- Under existing methods, latent models all require EM to learn regardless of the number of hidden states.
- However, is there a formulation of latent variable models where the difficulty of learning is a function of the number of latent states?
- This is the question that the spectral view will answer.


## Sum Rule (Matrix Form)

- Sum Rule

$$
\mathbb{P}[X]=\sum_{Y} \mathbb{P}[X \mid Y] \mathbb{P}[Y]
$$

- Equivalent view using Matrix Algebra

$$
\begin{aligned}
& \mathcal{P}[X]= \\
& \mathcal{P}[X \mid Y] \quad \times \mathcal{P}[Y] \\
& \binom{\mathbb{P}[X=0]}{\mathbb{P}[X=1]}=\left(\begin{array}{l}
\mathbb{P}[X=0 \mid Y=0] \\
\mathbb{P}[X=1 \mid Y=0] \quad \mathbb{P}[X=0 \mid Y=1] \\
\mathbb{P}[| | Y=1]
\end{array}\right) \times\binom{\mathbb{P}[Y=0]}{\mathbb{P}[Y=1]}
\end{aligned}
$$

## Important Notation

- Calligraphic P to denotes that the probability is being treated as a matrix/vector/tensor
- Probabilities

$$
\mathbb{P}[X, Y]=\mathbb{P}[X \mid Y] \mathbb{P}[Y]
$$

- Probability Vectors/Matrices/Tensors

$$
\mathcal{P}[X]=\mathcal{P}[X \mid Y] \mathcal{P}[Y]
$$

## Chain Rule (Matrix Form)

- Chain Rule

$$
\mathbb{P}[X, Y]=\mathbb{P}[X \mid Y] \mathbb{P}[Y]=\mathbb{P}[Y \mid X] \mathbb{P}[Y]
$$

- Equivalent view using Matrix Algebra

$$
\begin{aligned}
& \mathcal{P}[X, Y]= \\
& \mathcal{P}[X \mid Y] \\
& \left(\begin{array}{lll}
\mathbb{P}[X=0, Y=0] & \mathbb{P}[X=0, Y=1] \\
\mathbb{P}[X=1, Y=0] & \mathbb{P}[X=1, Y=1]
\end{array}\right) \quad= \\
& \left(\begin{array}{lll}
\mathbb{P}[X=0 \mid Y=0] & \mathbb{P}[X=0 \mid Y=1] \\
\mathbb{P}[X=1 \mid Y=0] & \mathbb{P}[X=1 \mid Y=1]
\end{array}\right) \quad X\left(\begin{array}{cc}
\mathbb{P}[Y=0] & 0 \\
0 & \mathbb{P}[Y=1]
\end{array}\right)
\end{aligned}
$$

Means on diagonal

- Note how diagonal is used to keep $\boldsymbol{Y}$ from being marginalized out.


## Graphical Models: The Linear Algebra View

$$
\mathcal{P}[A, B]
$$

## $A$ and $B$ have $m$ states each.



- In general, nothing we can say about the nature of this matrix.


## Independence: The Linear Algebra View

- What if we know $A$ and $B$ are independent?


## $\mathcal{P}[A, B]$



$$
\begin{aligned}
\longleftarrow & (\mathbb{P}[A=1, B=1], \ldots, \mathbb{P}[A=1, B=m]) \\
= & (\mathbb{P}[A=1](\mathbb{P}[B=1], \ldots, \mathbb{P}[B=m]))
\end{aligned}
$$

- Joint probability matrix is rank one, since all rows are multiples of one another!!


## Independence and Rank



## $\mathcal{P}[A, B]$ has rank $\boldsymbol{m}$ (at most)


has rank 1

- What about rank in between 1 and $m$ ?


## Low Rank Structure

- $\boldsymbol{A}$ and $\boldsymbol{B}$ are not marginally independent (They are only conditionally independent given $\boldsymbol{X}$ ).

- Assume $\boldsymbol{X}$ has $\boldsymbol{k}$ states (while $\boldsymbol{A}$ and $\boldsymbol{B}$ have $\boldsymbol{m}$ states).
- Then, $\operatorname{rank}(\mathcal{P}[A, B]) \leqslant k$
- Why?


## Low Rank Structure



## The Spectral View

- Latent variable models encode low rank dependencies among variables (both marginal and conditional)
- Use tools from linear algebra to exploit this structure.
- Rank
- Eigenvalues
- SVD
- Tensors


## A More Interesting Example



$$
\left\{X_{3}, X_{4}\right\}
$$

$\mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]$
has rank $k$

## Low Rank Matrices "Factorize"

## $M=\boldsymbol{L} \boldsymbol{R} \quad$ If m has rank k <br> mby $n \quad m$ by k $k$ by $n$

We already know one factorization!!!

$$
\begin{aligned}
& \mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]=\boldsymbol{P}\left[X_{\{1,2\}} \mid H_{2}\right] \boldsymbol{P}\left[\oslash H_{2}\right] \boldsymbol{P}\left[X_{\{3,4\}} \mid H_{2}\right]^{\top} \\
& \text { Factor of 4 variables } \quad \text { Factor of } 3 \text { variables }
\end{aligned}
$$

Factor of 1 variable

## Alternate Factorizations

- The key insight is that this factorization is not unique.
- Consider Matrix Factorization. Can add any invertible transformation:

$$
\begin{gathered}
\boldsymbol{M}=\boldsymbol{L} \boldsymbol{R} \\
\boldsymbol{M}=\boldsymbol{L} \boldsymbol{S} \boldsymbol{S}^{-1} \boldsymbol{R}
\end{gathered}
$$

- The magic of spectral learning is that there exists an alternative factorization that only depends on observed variables!


## An Alternate Factorization

- Let us say we only want to factorize this matrix of 4 variables

$$
\boldsymbol{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]
$$

such that it is product of matrices that contain at most three observed variables e.g.

$$
\begin{aligned}
& \mathcal{P}\left[X_{\{1,2\}}, X_{3}\right] \\
& \mathcal{P}\left[X_{2}, X_{\{3,4\}}\right]
\end{aligned}
$$

## An Alternate Factorization

- Note that

$$
\begin{aligned}
\mathcal{P}\left[X_{\{1,2\}}, X_{3}\right] & =\mathcal{P}\left[X_{\{1,2\}} \mid H_{2}\right] \mathcal{P}\left[\oslash H_{2}\right] \mathcal{P}\left[X_{3} \mid H_{2}\right]^{\top} \\
\mathcal{P}\left[X_{2}, X_{\{3,4\}}\right] & =\mathcal{P}\left[X_{2} \mid H_{2}\right] \mathcal{P}\left[\oslash H_{2}\right] \mathcal{P}\left[X_{\{3,4\}} \mid H_{2}\right]^{\top}
\end{aligned}
$$

- Product of green terms (in some order) is

$$
\mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]
$$

- Product of red terms (in some order) is $\mathcal{P}\left[X_{2}, X_{3}\right]$


## An Alternate Factorization

$$
\mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]=\mathcal{P}\left[X_{\{1,2\}}, X_{3}\right] \mathcal{P}\left[X_{2}, X_{3}\right]^{-1} \mathcal{P}\left[X_{2}, X_{\{3,4\}}\right]
$$

factor of 4 variables factor of 3 variables factor of 3 variables
Advantage: Factors are only functions of observed variables! Can be directly computed from data without EM!!!!

Caveat: some factors are no longer probability tables (do not have to be non-negative)

We will call this factorization the observable factorization.

## Graphical Relationship

$$
\mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]=\mathcal{P}\left[X_{\{1,2\}}, X_{3}\right] \mathcal{P}\left[X_{2}, X_{3}\right]^{-1} \mathcal{P}\left[X_{2}, X_{\{3,4\}}\right]
$$



## Another Factorization

$\mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]=\mathcal{P}\left[X_{\{1,2\}}, X_{4}\right] \mathcal{P}\left[X_{1}, X_{4}\right]^{-1} \mathcal{P}\left[X_{1}, X_{\{3,4\}}\right]$


- Seems we would do better empirically if you could "combine" both factorizations. Will come back to this later.


## Relationship to Original Factorization

- What is the relationship between the original factorization and the new factorization?

$$
\frac{\mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]}{\boldsymbol{M}}=\frac{\mathcal{P}\left[X_{\{1,2\}} \mid H_{2}\right] \mathcal{P}\left[\oslash H_{2}\right]}{\boldsymbol{L}} \frac{\mathcal{P}\left[X_{\{3,4\}} \mid H_{2}\right]^{\top}}{\boldsymbol{R}}
$$

$$
\begin{gathered}
M=\boldsymbol{L} \boldsymbol{R} \\
M=\boldsymbol{L} \boldsymbol{S} \boldsymbol{S}^{-1} \boldsymbol{R}
\end{gathered}
$$

Can I choose S to get the observable factorization?

## Relationship to Original Factorization

- Let

$$
\boldsymbol{S}:=\mathcal{P}\left[X_{3} \mid H_{2}\right]
$$

$\mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]=\frac{\mathcal{P}\left[X_{\{1,2\}}, X_{3}\right]}{=\boldsymbol{L} \boldsymbol{P}} \frac{\mathcal{S}}{} \frac{\left.X_{2}, X_{3}\right]^{-1} \mathcal{P}\left[X_{2}, X_{\{3,4\}}\right]}{=\boldsymbol{S}^{-1} \boldsymbol{R}}$

## Our Alternate Factorization

$$
\mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]=\mathcal{P}\left[X_{\{1,2\}}, X_{3}\right] \mathcal{P}\left[X_{2}, X_{3}\right]^{-1} \mathcal{P}\left[X_{2}, X_{\{3,4\}}\right]
$$

factor of 4 variables factor of 3 variables
factor of 3 variables

- It may not seem very amazing at the moment (we have only reduced the size of the factor by 1)
- What is cool is that every latent tree of $\boldsymbol{V}$ variables has such a factorization where:
- All factors are of size 3
- All factors are only functions of observed variables


## Training / Testing with Spectral Learning

- We have that

$$
\mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]=\mathcal{P}\left[X_{\{1,2\}}, X_{3}\right] \mathcal{P}\left[X_{2}, X_{3}\right]^{-1} \mathcal{P}\left[X_{2}, X_{\{3,4\}}\right]
$$

- In training, we compute estimates:

$$
\mathcal{P}_{M L E}\left[X_{\{1,2\}}, X_{3}\right] \quad \mathcal{P}_{M L E}\left[X_{2}, X_{3}\right]^{-1} \quad \mathcal{P}_{M L E}\left[X_{2}, X_{\{3,4\}}\right]
$$

- In test time, we can compute probability estimates (let lowercase letters denote fixed evidence values):
$\widehat{\mathbb{P}}_{\text {spec }}\left[x_{1}, x_{2}, x_{3}, x_{4}\right]=\mathcal{P}_{M L E}\left[x_{\{1,2\}}, X_{3}\right] \mathcal{P}_{M L E}\left[X_{2}, X_{3}\right]^{-1} \mathcal{P}_{M L E}\left[X_{2}, x_{\{3,4\}}\right]^{\top}$


## Generalizing To More Variables

- Consider HMM with 5 observations. Using similar arguments as before we will get that:

$$
\mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4,5\}}\right]=\mathcal{P}\left[X_{\{1,2\}}, X_{3}\right] \mathcal{P}\left[X_{2}, X_{3}\right]^{-1} \mathcal{P}\left[X_{2}, X_{\{3,4,5\}}\right]
$$


reshape and decompose recursively
$\mathcal{P}\left[X_{\{2,3\}}, X_{\{4,5\}}\right]=\mathcal{P}\left[X_{\{2,3\}}, X_{4}\right] \mathcal{P}\left[X_{3}, X_{4}\right]^{-1} \mathcal{P}\left[X_{3}, X_{\{4,5\}}\right]$

## Consistency

- A trivial consistent estimator is to simply attempt to estimate the "big" probability table from the data without making any conditional independence assumptions
$\mathcal{P}_{M L E}\left[X_{1}, X_{2} ; X_{3}, X_{4}\right] \rightarrow \mathcal{P}\left[X_{1}, X_{2} ; X_{3}, X_{4}\right]$
as number of samples increases
- While this is consistent, it is not very statistically efficient


## Consistency

- A better estimate is to get compute likelihood estimates of the factorization:

$$
\begin{gathered}
\mathcal{P}_{M L E}\left[X_{\{1,2\}} \mid H_{2}\right] \mathcal{P}_{M L E}\left[\oslash H_{2}\right] \mathcal{P}_{M L E}\left[X_{\{3,4\}} \mid H_{2}\right]^{\top} \\
\rightarrow \mathcal{P}\left[X_{1}, X_{2} ; X_{3}, X_{4}\right]
\end{gathered}
$$

- But this requires running EM, which will get stuck in local optima and is not guaranteed to obtain the MLE of the factorized model


## Consistency

- In spectral learning, we estimate the alternate factorization from the data

$$
\begin{aligned}
& \mathcal{P}_{M L E}\left[X_{\{1,2\}}, X_{3}\right] \mathcal{P}_{M L E}\left[X_{2}, X_{3}\right]^{-1} \mathcal{P}_{M L E}\left[X_{2}, X_{\{3,4\}}\right] \\
& \rightarrow \mathcal{P}\left[X_{1}, X_{2} ; X_{3}, X_{4}\right]
\end{aligned}
$$

- This is consistent and computationally tractable (at some loss of statistical efficiency due to the dependence on the inverse)


## Where's the Catch?

- Before we said that if the number of latent states was very large then the model was equivalent to a clique.
- Where does that scenario enter in our factorization?

$$
\mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]=\mathcal{P}\left[X_{\{1,2\}}, X_{3}\right] \mathcal{P}\left[X_{2}, X_{3}\right]^{-1} \mathcal{P}\left[X_{2}, X_{\{3,4\}}\right]
$$

## When Does the Inverse Exist

$$
\mathcal{P}\left[X_{2}, X_{3}\right]=\mathcal{P}\left[X_{2} \mid H_{2}\right] \mathcal{P}\left[\oslash H_{2}\right] \mathcal{P}\left[X_{3} \mid H_{2}\right]^{\top}
$$

- All the matrices on the right hand side must have full rank. (This is in general a requirement of spectral learning, although it can be somewhat relaxed)


## When $\mathrm{m}>\mathrm{k}$

- The inverse cannot exist, but this situation is easily fixable (project onto lower dimensional space)

$$
\begin{aligned}
& \mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]= \\
& \quad \mathcal{P}\left[X_{\{1,2\}}, X_{3}\right] \boldsymbol{V}\left(\boldsymbol{U}^{\top} \boldsymbol{P}\left[X_{2}, X_{3}\right] \boldsymbol{V}\right)^{-1} \boldsymbol{U}^{\top} \boldsymbol{P}\left[X_{2}, X_{\{3,4\}}\right]
\end{aligned}
$$

- Where $\boldsymbol{U}, \boldsymbol{V}$ are the top left/right $\mathbf{k}$ singular vectors of $\boldsymbol{\mathcal { P }}\left[X_{2}, X_{3}\right]$


## When k > m

- The inverse does exist. But it no longer satisfies the following property, which we used to derive the factorization

$$
\mathcal{P}\left[X_{2}, X_{3}\right]^{-1}=\left(\mathcal{P}\left[X_{3} \mid H_{2}\right]^{\top}\right)^{-1} \mathcal{P}\left[\oslash H_{2}\right]^{-1} \mathcal{P}\left[X_{2} \mid H_{2}\right]^{-1}
$$

- This is much more difficult to fix, and intuitively corresponds to how the problem becomes intractable if $\boldsymbol{k} \gg \boldsymbol{m}$.


## What does $\mathrm{k}>\mathrm{m}$ mean?

- Intuitively, large $\boldsymbol{k}$, small $\boldsymbol{m}$ means long range dependencies
- Consider following generative process:
(1) With probability 0.5 , let $S=X$, and with probability 0.5 let $S=Y$.
(2) Print $\boldsymbol{A} \boldsymbol{n}$ times.
(3) Print S
(4) Go back to step (2)

With $n=1$ we either generate:
AXAXAXA...... or AYAYAYA.....

With $n=2$ we either generate:
AAXAAXAA..... or AAYAAYAA.......

## How many hidden states does HMM need?

- HMM needs $2 n$ states.
- Needs to remember count as well as whether we picked $\mathbf{S}=\boldsymbol{X}$ or $\boldsymbol{S}=\boldsymbol{Y}$
- However, number of observed states $\boldsymbol{m}$ does not change, so our previous spectral algorithm will break for $\boldsymbol{n} \boldsymbol{>} \mathbf{2}$.
- How to deal with this in spectral framework?


## Making Spectral Learning Work In Practice

- We are only using marginals of pairs/triples of variables to construct the full marginal among the observed variables.
- Only works when $\boldsymbol{k}<\boldsymbol{m}$.

- However, in real problems we need to capture longer range dependencies.


## Recall our factorization

$$
\mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]=\mathcal{P}\left[X_{\{1,2\}}, X_{3}\right] \mathcal{P}\left[X_{2}, X_{3}\right]^{-1} \mathcal{P}\left[X_{2}, X_{\{3,4\}}\right]
$$



## Key Idea: Use Long-Range Features



Construct feature vector of left side

$$
\phi_{L}
$$

Construct feature vector of right side

## Spectral Learning With Features

$$
\mathcal{P}\left[X_{2}, X_{3}\right]=\mathbb{E}\left[\delta_{\boldsymbol{\delta}}^{\otimes} \otimes \delta_{3}\right]:=\mathbb{E}\left[\delta_{2} \delta_{3}^{\top}\right]
$$

Use more complex feature instead:

$$
\mathbb{E}\left[\phi_{L} \otimes \phi_{R}\right]
$$

$$
\begin{aligned}
& \mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]=\mathbb{E}\left[\boldsymbol{\delta}_{1 \otimes 2}, \boldsymbol{\delta}_{3 \otimes 4}\right] \\
& \quad=\mathbb{E}\left[\boldsymbol{\delta}_{1 \otimes 2}, \boldsymbol{\phi}_{R}\right] \boldsymbol{V}\left(\boldsymbol{U}^{\top} \mathbb{E}\left[\boldsymbol{\phi}_{L} \otimes \boldsymbol{\phi}_{R}\right] \boldsymbol{V}\right)^{-1} \boldsymbol{U}^{\top} \mathcal{P}\left[\boldsymbol{\phi}_{L}, X_{\{3,4\}}\right]
\end{aligned}
$$

## Experimentally,

- Has been shown by many authors that (with some work) spectral methods achieve comparable results to EM but are 10-50x faster
- Parikh et al. 2011 / 2012
- Balle et al. 2012
- Cohen et al. 2012 / 2013
- The following are some synthetic and real data results demonstrating the comparison between EM and spectral methods.


## Synthetic Data [Parikh et al. 2012]

- Synthetic $3^{\text {rd }}$ order HMM Example (Spectral/EM/Online EM):


Runtime vs. Sample Size

Error vs. Sample Size


## Empirical Results for Latent PCFGs [cohen et al. 2013]

|  | section 22 |  | section 23 |  |
| :--- | :---: | :---: | :---: | :---: |
|  | EM | spectral | EM | spectral |
| $m=8$ | 86.87 | 85.60 | - | - |
| $m=16$ | 88.32 | 87.77 | - | - |
| $m=24$ | 88.35 | 88.53 | - | - |
| $m=32$ | 88.56 | 88.82 | 87.76 | 88.05 |

## Timing Results on Latent PCFGs[Cohen et al. 2013]

|  | single | EM |  |  |  |  |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | EM iter. | best model | total | feature | spectral algorithm <br> transfer + scaling | SVD | $a \rightarrow b c$ | $a \rightarrow x$ |
| $m=8$ | 6 m | 3 h | 3 h 32 m |  |  | 36 m | 1 h 34 m | 10 m |
| $m=16$ | 52 m | 26 h 6 m | 5 h 19 m | 22 m | 49 m | 34 m | 3 h 13 m | 19 m |
| $m=24$ | 3 h 7 m | 93 h 36 m | 7 h 15 m |  |  | 36 m | 4 h 54 m | 28 m |
| $m=32$ | 9 h 21 m | 187 h 12 m | 9 h 52 m |  |  | 35 m | 7 h 16 m | 41 m |

## Dealing with Nonparametric, Continuous Variables

- It is difficult to run EM if the conditional/marginal distributions are continuous and do not easily fit into a parametric family.



- However, we will see that Hilbert Space Embeddings can easily be combined with spectral methods for learning nonparametric latent models.


## Connection to Hilbert Space Embeddings

- Recall that we could substitute features for variables

$$
\mathcal{P}\left[X_{2}, X_{3}\right]=\mathbb{E}\left[\boldsymbol{\delta}_{2} \otimes \boldsymbol{\delta}_{3}\right]:=\mathbb{E}\left[\boldsymbol{\delta}_{2} \boldsymbol{\delta}_{3}^{\top}\right]
$$

Use more complex feature instead:

$$
\mathbb{E}\left[\phi_{L} \otimes \phi_{R}\right]
$$

## Can Also Use Infinite Dimensional Features

- Replace

$$
\mathcal{P}\left[X_{2}, X_{3}\right]=\mathbb{E}\left[\boldsymbol{\delta}_{2} \otimes \boldsymbol{\delta}_{3}\right]:=\mathbb{E}\left[\boldsymbol{\delta}_{2} \boldsymbol{\delta}_{3}^{\top}\right]
$$

- with
$\mathcal{C}\left[X_{2}, X_{3}\right]=\mathbb{E}\left[\phi_{X_{2}} \otimes \phi_{X_{3}}\right] \quad \begin{gathered}\text { covariance } \\ \text { operator }\end{gathered}$
- (and similarly for other quantities)


## Connection to Hilbert Space Embeddings

Discrete case:

$$
\begin{aligned}
& \mathcal{P}\left[X_{\{1,2\}}, X_{\{3,4\}}\right]= \\
& \quad \mathcal{P}\left[X_{\{1,2\}}, X_{3}\right] \boldsymbol{V}\left(\boldsymbol{U}^{\top} \boldsymbol{\mathcal { P }}\left[X_{2}, X_{3}\right] \boldsymbol{V}\right)^{-1} \boldsymbol{U}^{\top} \mathcal{P}\left[X_{2}, X_{\{3,4\}}\right]
\end{aligned}
$$

Continuous case:

$$
\begin{aligned}
& \mathcal{C}\left[X_{\{1,2\}} ; X_{\{3,4\}}\right]= \\
& \quad \mathcal{C}\left[X_{\{1,2\}} ; X_{3}\right] \boldsymbol{V}\left(\boldsymbol{U}^{\top} \mathcal{C}\left[X_{2}, X_{3}\right] \boldsymbol{V}\right)^{-1} \boldsymbol{U}^{\top} \mathcal{C}\left[X_{2} ; X_{\{3,4\}}\right]
\end{aligned}
$$

## Summary - EM \& Spectral (Part I)

## EM

- Aims to Find MLE so more "statistically" efficient
- Can get stuck in local-optima
- Lack of theoretical guarantees
- Slow
- Easy to derive for new models


## Spectral

- Does not aim to find MLE so less statistically efficient.
- Local-optima-free
- Provably consistent
- Very fast
- Challenging to derive for new models (Unknown whether it can generalize to arbitrary loopy models)


## Summary - EM \& Spectral (Part II)

## EM

- No issues with negative numbers
- Allows for easy modelling with conditional distributions
- Difficult to incorporate long-range features (since it increases treewidth).
- Generalizes poorly to nonGaussian continuous variables.


## Spectral

- Problems with negative numbers. Requires explicit normalization to compute likelihood.
- Allows for easy modelling with marginal distributions
- Easy to incorporate long-range features.
- Easy to generalize to nonGaussian continuous variables via Hilbert Space Embeddings

