## Thus far, fully supervised learning

- We have assumed fully supervised learning:

- Many real problems have missing data:

$$
\left\langle A=t_{1} F=f, S=?, H=t, N=?\right\rangle
$$

## The general learning problem with missing data

- Marginal likelihood $-\mathbf{x}$ is observed, $\mathbf{z}$ is missing:

$$
\begin{aligned}
\ell(\mathcal{D}: \theta) & =\log \prod_{j=1}^{m} P\left(x^{\text {Dobsenca }} \mid \theta\right) \\
& =\sum_{j=1}^{m} \log P\left(x^{(j)} \mid \theta\right) \\
& =\sum_{j=1}^{m} \log \sum_{z} P\left(z, x^{(j)} \mid \theta\right)
\end{aligned}
$$

## E-step

$\mathbf{x}$ is observed, $\mathbf{z}$ is missing

- Compute probability of missing data given current choice of $\theta$
$\square \mathrm{Q}\left(\mathbf{z} \mid \mathbf{x}^{(j)}\right)$ for each $\mathbf{x}^{(j)}$
- e.g., probability computed during classification step
- corresponds to "classification step" in $K$-means

$$
\underbrace{\hat{i n f e n c e ~ i n ~ a ~ B N ~}^{P}}_{\hat{Q^{(t+1)}\left(z \mid x^{(j)}\right)=P\left(z \mid x^{(j)}, \theta^{(t)}\right)}}
$$

## Jensen's inequality

$$
\ell(\mathcal{D}: \theta)=\sum_{j=1}^{m} \log \sum_{z} P\left(z, x^{(j)} \mid \theta\right)
$$

- Theorem: $\log \sum_{\mathbf{z}} \mathrm{P}(\mathbf{z}) \mathrm{f}(\mathbf{z}) \geq \sum_{\mathbf{z}} \mathrm{P}(\mathbf{z}) \log f(\mathbf{z})$

$$
\log 2=\log (1+1) \geqslant \log |+|\log |=0
$$

## 

- Use: $\log \sum_{\mathbf{z}} \mathrm{P}(\mathbf{z}) \mathrm{f}(\mathbf{z}) \geq \sum_{\mathbf{z}} \mathrm{P}(\mathbf{z}) \log \mathrm{f}(\mathbf{z})$
$\ell\left(\mathcal{D}: \theta^{(t)}\right)=\sum_{j=1}^{m} \log \sum_{z} \underbrace{Q^{(t+1)}\left(z \mid x^{(j)}\right.}_{" \prime P(z)^{\prime \prime}}) \underbrace{\frac{P\left(z, x^{(j)} \mid \theta^{(t)}\right)}{Q^{(t+1)}\left(z \mid x^{(j)}\right)}}_{" f(z)^{\prime \prime}}$
qQ(Di $\left.\left.\theta^{(t)}\right) \geqslant \sum_{j=1}^{m} \sum_{z} Q^{(t+1)}\left(z \mid x^{(j)}\right)\right) \log \frac{p\left(z, x^{(s)} \mid \theta^{(t)}\right)}{Q^{(t+1)}\left(z(z)^{(j)}\right)}$

wrighted log-ikelibood of
fully observed dat
fully observed data m $\hat{H}_{a^{(t+1)}}(Z \mid X)$.


## The M-step maximizes lower bound on weighted data <br> Lower bound from Jensen's: <br> $\ell\left(\mathcal{D}: \theta^{(t)}\right) \geq \sum_{j=1}^{m} \sum_{z} Q^{(t+1)}\left(z \mid x^{(j)}\right) \log P\left(z, x^{(j)} \mid \theta^{(t)}\right)+{ }^{\hat{H}} H\left(Q^{(t+1)}\right)$ <br> optimize $\theta$ <br> 

- Corresponds to weighted dataset:

$$
\begin{aligned}
& <\mathbf{x}^{(1)}, \mathbf{z}=1>\text { with weight } Q^{(t+1)}\left(\mathbf{z}=1 \mid \mathbf{x}^{(1)}\right) \quad .8 \\
& <\mathbf{x}^{(1)}, \mathbf{z}=2>\text { with weight } Q^{(t+1)}\left(\mathbf{z}=2 \mid \mathbf{x}^{(1)}\right) \\
& <\mathbf{x}^{(1)}, \mathbf{z}=3>\text { with weight } Q^{(t+1)}\left(\mathbf{z}=3 \mid \mathbf{x}^{(1)}\right) .05 \\
& <\mathbf{x}^{(2)}, \mathbf{z}=1>\text { with weight } Q^{(t+1)}\left(\mathbf{z}=1 \mid \mathbf{x}^{(2)}\right) \\
& <\mathbf{x}^{(2)}, \mathbf{z}=2>\text { with weight } Q^{(++1)}\left(\mathbf{z}=2 \mid \mathbf{x}^{(2)}\right) .8 \\
& <\mathbf{x}^{(2)}, \mathbf{z}=3>\text { with weight } Q^{(t+1)}\left(\mathbf{z}=3 \mid \mathbf{x}^{(2)}\right), 1
\end{aligned}
$$

## The M-step <br>  <br> $\ell\left(\mathcal{D}: \theta^{(t)}\right) \geq \sum_{j=1}^{m} \sum_{z} Q^{(t+1)}\left(z \mid x^{(j)}\right) \log P\left(z, x^{(j)} \mid \theta^{(t)}\right)+H\left(Q^{(t+1)}\right)$

- Maximization step:

$$
\begin{aligned}
& \theta^{(t+1)} \leftarrow \arg \max _{\theta} \sum_{j=1}^{m} \sum_{z} Q^{(t+1)}\left(z \mid x^{(j)}\right) \log P\left(z, x^{(j)} \mid \theta\right) \\
& E_{Q^{(t+1)}}\left[C_{\text {ont }}(x=t, z=f)\right]=\sum_{j=1}^{m} 11\left(x^{(j)}=t\right) Q^{(t+1)}\left(z=f\left(x^{(j)}\right)\right.
\end{aligned}
$$

- Use expected counts instead of counts:
$\square$ If learning requires $\operatorname{Count}(\mathbf{x}, \mathbf{z})$
Use $\mathrm{E}_{\mathrm{Q}(t+1)}[$ Count $(\mathbf{x}, \mathbf{z})]$


## Convergence of EM

- Define potential function $\mathrm{F}(\theta, \mathrm{Q})$ :

$$
\left.\ell\left(\mathcal{D}: \theta^{(t)}\right) \geq F(\theta, Q)=\sum_{j=1}^{m} \sum_{z} \overline{Q(z \mid} x^{(j)}\right) \log \frac{P\left(z, x^{(j)} \mid \theta\right)}{Q\left(z \mid x^{(j)}\right)}
$$



- EM corresponds to coordinate ascent on F

Thus, maximizes lower bound on marginal log likelihood
As seen in Machine Learning class last semester

$$
\begin{aligned}
& \text { whiting } Q \rightarrow \text { max over } \theta \\
& \text { fixing } \theta \rightarrow \max \text { over } Q
\end{aligned}
$$

## Data likelihood for ENs

- Given structure, log likelihood of fully observed data:
$\log P\left(\mathcal{D} \mid \theta_{\mathcal{G}}, \mathcal{G}\right)$
$=\sum_{j=1}^{m}$

Idocompoes


## E-step for BNs



- E-step computes probability of hidden vars $\mathbf{h}$ given $\mathbf{o}$ $\underbrace{Q^{(t+1)}\left(\mathbf{h} \mid \mathbf{o}^{(i)}\right)} \leftarrow P\left(\mathbf{h} \mid \mathbf{o}^{(j)}, \theta^{(t)}\right)$ compute with:
- Corresponds to inference in BN

$$
\begin{aligned}
& \text { Naively, must represent joint over H } 10^{(j)} \\
& \text { if there are k hidden vars, then joint ; hugh! }
\end{aligned}
$$

## The M-step for ENs

- Maximization step:

$$
\theta^{(t+1)} \leftarrow \arg \max _{\theta} \sum_{j=1}^{m} \sum_{\mathbf{h}} Q^{(t+1)}\left(\mathbf{h} \mid \mathbf{o}^{(j)}\right) \log P\left(\mathbf{h}, \mathbf{o}^{(j)} \mid \theta\right)
$$

- Use expected counts instead of counts:
$\square$ If learning requires Count( $\mathbf{h}, \mathbf{0}$ )
$\square$ Use $\mathrm{E}_{\mathrm{Q}(t+1)}[\operatorname{Count}(\mathbf{h}, \mathbf{o})]$


## M-step for each CPT

## M-step decomposes per CPT

$\square$ Standard MLE:
$P(\underbrace{}_{\left.X_{i}=x_{i} \mid \mathbf{P a}_{X_{i}}=\mathrm{z}\right)}=\frac{\operatorname{Count}\left(X_{i}=x_{i}, \mathbf{P a}_{X_{i}}=\mathbf{z}\right)}{\operatorname{Count}\left(\mathbf{P a}_{X_{i}}=\mathrm{z}\right)}$

$\square$ M-step uses expected counts:

$P\left(X_{i}=x_{i} \mid \mathbf{P a}_{X_{i}}=\mathbf{z}\right)=\frac{\operatorname{ExCount}\left(X_{i}=x_{i}, \mathrm{~Pa}_{X_{i}}=\mathbf{z}\right)}{\operatorname{ExCount}\left(\operatorname{Pa}_{X_{i}}=\mathbf{z}\right)}$
erg., $P\left(H^{\prime \prime} \mid \delta=f\right) \stackrel{\text { me }}{=} C$ count $(1 f=t, S=f)$
$\operatorname{count}\left(\int=f\right)$
ExComt $(H=t, \delta=f)=\sum_{j=1}^{m} \pi\left(H^{(j)}=t\right)$


## Computing expected counts

$$
P\left(X_{i}=x_{i} \mid \mathrm{Pa}_{X_{i}}=\mathrm{z}\right)=\frac{\operatorname{ExCount}\left(X_{i}=x_{i}, \mathrm{~Pa}_{X_{i}}=z\right)}{\operatorname{ExCount}\left(\mathrm{Pa}_{X_{i}}=\mathrm{z}\right)}
$$

- M-step requires expected counts:
$\square$ Observe $\mathbf{O = 0}$
$\square$ For a set of vars $\mathbf{A}$, must compute $\operatorname{ExCount(A=a)}$
$\square$ Some of $\mathbf{A}$ in example $j$ will be observed
- denote by $\mathbf{A}_{0}=\mathbf{a}_{0}{ }^{(j)} \quad A_{0} \leqslant 0$

Some of $\mathbf{A}$ will be hidden

- denote by $\mathbf{A}_{H} \quad A_{H} \cap O=\varnothing$
- Use inference (E-step computes expected counts):



## Data need not be hidden in the same way



- When data is fully observed
$\square$ A data point is $\langle F=t, A=f, S=t\rangle$
ugly notation
- When data is partially observed $A=f, S=f$ )


A data point is

$$
\angle F=f, A=f, S=?\rangle
$$

( $F=F, A=f \quad S=$ ?

- But unobserved variables can be different for different data points
$\square$ egg.,
$\angle F=t, A=t, s=?\rangle$
$\langle F=?, A=?, S=t\rangle$
$C F=t, A=?, S=?$
- Same framework, just change definition of expected counts
$\square$ Observed vars in point $j$, $O_{j} t$ changes for each $j$
$\square$ Consider set of vars $\mathbf{A}$

$$
A O_{j}=A \wedge O_{j} \quad A_{A j}=A / A_{D_{j}}
$$

$\square$ ExCount $^{(t+1)}(\mathbf{A}=\mathbf{a})$
$=\sum_{j=1}^{m}$
$I\left(A_{0 j}=a_{0 j}\right) Q^{(t+1)}\left(A_{i H_{j}}=a_{H j} \mid O_{j}^{(i)}\right)$

17

## Poster printing

## facilities wo sully dons

- Poster session:
$\square$ Friday Dec 1st, 3-6pm in the NSH Atrium.
$\square$ There will be a popular vote for best poster. Invite your friends!
$\square$ please be ready to set up your poster at 2:45pm sharp.
- We will provide posterboards, easels and pins.
$\square$ The posterboards are $30 \times 40$ inches
$\square$ We don't have a specific poster format for you to use.
- You can either bring a big poster or a print a set of regular sized pages and pin them together.
- Unfortunately, we don't have a budget to pay for printing. If you are an SCS student, SCS has a poster printer you can use which prints on a 36 " wide roll of paper.
( If you are a student outside SCS, you will need to check with your department to see if there are printing facilities for big posters (I don't know what is offered outside SCS), or print a set of regular sized pages.
- We are looking forward to a great poster session!


## EM for BN \& identifiability: a superficial discussion

What happens if a leaf is never observed?


## Learning structure with missing data <br> [K\&F 18.4]

- Known BN structure: Use expected counts, learning algorithm doesn't change
- Unknown BN structure:
$\square$ Can use expected counts and score model as when we talked about structure learning


But, very slow...

- e.g., greedy algorithm would need to redo inference for every edge we test...
- (Much Faster) Structure-EM: Iterate:
$\square$ compute expected counts
$\square$ do a some structure search (e.g., many greedy steps)
repeat
- Theorem: Converges to local optima of marginal loglikelihood
$\square$ details in the book


## What you need to know about learning ENs with missing data

- EM for Bayes Nets
- E-step: inference computes expected counts
$\square$ Only need expected counts over $\mathrm{X}_{\mathrm{i}}$ and $\mathrm{Pa}_{\mathrm{xi}}$
- M-step: expected counts used to estimate parameters
- Which variables are hidden can change per datapoint
$\square$ Also, use labeled and unlabeled data $\rightarrow$ some data points are complete, some include hidden variables
- Structure-EM:
$\square$ iterate between computing expected counts \& many structure search steps


## MAs \& CRFs with missing data

- Ns with missing data
$\square$ Models $\mathbf{P}(\mathbf{X})$, part of $\mathbf{X}$ hidden
$\square$ Use EM to optimize
$\square$ Same ideas as BN
- CRFs with missing data
$\square$ Models $\mathrm{P}(\mathbf{Y} \mid \mathbf{X})$
$\square$ What's hidden?
- Part of $\mathbf{Y}: Y_{0}, Y_{H}$

- Part of $\mathbf{X}$ :



## Kalman Filters Gaussian BNs

Graphical Models - 10708
Carlos Guestrin
Carnegie Mellon University
November 24 ${ }^{\text {th }}, 2008$

## Adventures of our BN hero

- Compact representation for 1. Naïve Bayes probability distributions
- Fast inference
- Fast learning
- Approximate inference

2 and 3.
Hidden Markov models (HMMs)

- But... Who are the most popular kids?

Kalman Filters

gaussion "CPTS"

## The Kalman Filter

An HMM with Gaussian distributions

- Has been around for at least 60 years
- Possibly the most used graphical model ever
- It's what
$\square$ does your cruise control
$\square$ tracks missiles
$\square$ controls robots
- And it's so simple...
$\square$ Possibly explaining why it's so used
- Many interesting models build on it...
$\square$ An example of a Gaussian BN (more on this later)


## Example of KF - SLAT <br> Simultaneous Localization and Tracking

[Funiak, Guestrin, Paskin, Sukthankar '06]

- Place some cameras around an environment, don't know where they are
- Could measure all locations, but requires lots of grad. student (Stano) time
- Intuition:
$\square$ A person walks around
$\square$ If camera 1 sees person, then camera 2 sees person, learn about relative positions of cameras




## Multivariate Gaussian

$$
p\left(X_{1}, \ldots, X_{n}\right)=\frac{1}{(2 \pi)^{n / 2}|\Sigma|^{1 / 2}} \exp \left\{-\frac{1}{2}(\mathbf{x}-\mu)^{T} \Sigma^{-1}(\mathbf{x}-\mu)\right\}
$$

Mean vector:


Covariance matrix:
 $\delta_{32}=\sigma_{23}$




## Conditioning a Gaussian

$\begin{array}{rr}\text { - } & \\ \text { - Joint Gaussian: } \\ \quad \mathrm{p}(X, Y) \sim N(\mu ; \Sigma)\end{array}$

- Conditional linear Gaussian:
$\square \mathrm{p}(\mathrm{Y\mid X}) \sim N\left(\mu_{Y \mid X} ; \sigma_{Y \mid X}^{2}\right) \longleftarrow$ gaussian

$$
\mu_{Y \mid X}=\frac{\mu_{Y}}{\pi_{\text {pili }}}+\frac{\sigma_{Y X}}{\sigma_{X}^{2}}\left(x-\mu_{x}\right)
$$

