# Learning in Bayes Nets 

Mausam

## (Based on slides by Stuart Russell, Marie desJardins, Subbarao Kambhampati, Dan Weld)

## Parameter Estimation

- Learn all the CPTs in a Bayesian Net
- Data $\rightarrow$ Model $\rightarrow$ Queries
- Key idea: counting!


## Burglars and Earthquakes



## Counting



| E | B | A | \# |
| :--- | :--- | :--- | :--- |
| 0 | 0 | 0 | 1000 |
| 0 | 0 | 1 | 10 |
| 0 | 1 | 0 | 20 |
| 0 | 1 | 1 | 100 |
| 1 | 0 | 0 | 200 |
| 1 | 0 | 1 | 50 |
| 1 | 1 |  | 1 |
| 1 | e,b <br> $\mathrm{e}, \overline{\mathrm{b}}$ <br> $\overline{\mathrm{e}, \mathrm{b}}$ <br> $\overline{\mathrm{e}, \overline{\mathrm{b}}}$ |  |  |

## Counting




## Counting




## Counting




## Counting




## Counting



| E | B |  | A | \# |
| :---: | :---: | :---: | :---: | :---: |
| 0 | 0 |  | 0 | 1000 |
| 0 | 0 |  | 1 | 10 |
| 0 | 1 |  | 0 | 20 |
| 0 | 1 |  | 1 | 100 |
| 1 | 0 |  | 0 | 200 |
| 1 | 0 |  | 1 | 50 |
| 1 | 1 |  | 0 | 0 |
| 1 | 1 |  | 1 | 5 |
|  | $\begin{aligned} & \hline e, b \\ & e, b \\ & \bar{e}, b \\ & \bar{e}, \bar{b} \end{aligned}$ | $\begin{array}{\|l} \mid \operatorname{Pr}(\mathrm{A} \mid \mathrm{E}, \mathrm{~B}) \\ \hline 1 \\ 0.2 \\ 0.83 \\ \sim 0.01 \end{array}$ |  | Bad idea to have prob as 0 or 1 <br> - stumps Gibbs sampling <br> - low prob states become impossible |

## Solution: Smoothing

- Why?
- To deal with events observed zero times.
- "event": a particular ngram
- How?
- To shave a little bit of probability mass from the higher counts, and pile it instead on the zero counts
- Laplace Smoothing/Add-one smoothing
- assume each event was observed at least once.
- add 1 to all frequency counts
- Add $m$ instead of 1 ( $m$ could be $>$ or $<1$ )


## Counting w/ Smoothing




## ML vs. MAP Learning

- ML: maximum likelihood (what we just did)
- find parameters that maximize the prob of seeing the data $D$
$-\operatorname{argmax}_{\theta} P(D \mid \theta)$
- easy to compute (for example, just counting)
- assumes uniform prior
- Prior: your belief before seeing any data
- Uniform prior: all parameters equally likely
- MAP: maximum a posteriori estimate
- maximize prob of parameters after seeing data $D$
$-\operatorname{argmax}_{\theta} P(\theta \mid D)=\operatorname{argmax}_{\theta} P(D \mid \theta) P(\theta)$
- allows user to input additional domain knowledge
- better parameters when data is sparse...
- reduces to ML when infinite data


## Example

Suppose there are five kinds of bags of candies:
$10 \%$ are $h_{1}: 100 \%$ cherry candies
$20 \%$ are $h_{2}: 75 \%$ cherry candies $+25 \%$ lime candies
$40 \%$ are $h_{3}: 50 \%$ cherry candies $+50 \%$ lime candies $20 \%$ are $h_{4}: 25 \%$ cherry candies $+75 \%$ lime candies $10 \%$ are $h_{5}: 100 \%$ lime candies


Then we observe candies drawn from some bag:
What kind of bag is it? What flavour will the next candy be?
Learning
Inference

## Full Bayesian learning

View learning as Bayesian updating of a probability distribution over the hypothesis space
$H$ is the hypothesis variable, values $h_{1}, h_{2}, \ldots$, prior $\mathbf{P}(H)$
$j$ th observation $d_{j}$ gives the outcome of random variable $D_{j}$ training data $\mathrm{d}=d_{1}, \ldots, d_{N}$

Given the data so far, each hypothesis has a posterior probability:


$$
P\left(h_{i} \mid \mathbf{d}\right)=\alpha P\left(\mathbf{d} \mid h_{i}\right) P\left(h_{i}\right)
$$

where $P\left(\mathbf{d} \mid h_{i}\right)$ is called the likelihood
Predictions use a likelihood-weighted average over the hypotheses:

$$
\mathbf{P}(X \mid \mathbf{d})=\sum_{i} \mathbf{P}\left(X \mid \mathbf{d}, h_{i}\right) P\left(h_{i} \mid \mathbf{d}\right)=\sum_{i} \mathbf{P}\left(X \mid h_{i}\right) P\left(h_{i} \mid \mathbf{d}\right)
$$

No need to pick one best-guess hypothesis!

## Posterior probability of hypotheses



True hypothesis eventually dominates...
probability of indefinitely producing uncharacteristic data $\rightarrow 0$


## ML vs. MAP Learning

- ML: maximum likelihood (what we just did)
- find parameters that maximize the prob of seeing the data $D$
$-\operatorname{argmax}_{\theta} P(D \mid \theta)$
- easy to compute (for example, just counting)
- assumes uniform prior
- Prior: your belief before seeing any data
- Uniform prior: all parameters equally likely
- MAP: maximum a posteriori estimate
- maximize prob of parameters after seeing data $D$
$-\operatorname{argmax}_{\theta} P(\theta \mid D)=\operatorname{argmax}_{\theta} P(D \mid \theta) P(\theta)$
- allows user to input additional domain knowledge
- better parameters when data is sparse...
- reduces to ML when infinite data


## Learning the Structure

- Problem: learn the structure of Bayes nets
- Search thru the space...
- of possible network structures!
- Heuristic search/local search
- For each structure, learn parameters
- Pick the one that fits observed data best
- Caveat - won't we end up fully connected????

When scoring, add a penalty
$\propto$ model complexity

## Local Search



## How to learn when some data missing?

- Expectation Maximization (EM)


## Example



| Examples: | 0 | 1 | 1 |
| :--- | :--- | :--- | :--- |
|  | 1 | 0 | 0 |
|  | 1 | 1 | 1 |
|  | 1 | $?$ | 0 |

$\begin{array}{lrrl}\text { Initialization: } & P(B \mid A) & = & P(C \mid B) \\ P(A)= & P(B \mid \neg A) & = & P(C \mid \neg B)\end{array}$

## Chicken \& Egg Problem

- If we knew the missing value
- It would be easy to learn CPT
- If we knew the CPT
- Then it'd be easy to infer the (probability of) missing value
- But we do not know either!


## Example


$\begin{array}{lrlrl}\text { Initialization: } & P(B \mid A) & =0 & P(C \mid B) & =0 \\ P(A)=0.75 & P(B \mid \neg A) & =0 & P(C \mid \neg B) & =0\end{array}$
E-step: $P(?=1)=P(B \mid A, \neg C)=\frac{P(A, B, \neg C)}{P(A, \neg C)}=\ldots=0$

M-step:
$P(A)=$
E-step: $P(?=1)=$

## Example


$\begin{array}{lrlrl}\text { Initialization: } & P(B \mid A) & =0 & P(C \mid B) & =0 \\ P(A)=0.75 & P(B \mid \neg A) & =0 & P(C \mid \neg B) & =0\end{array}$
E-step: $P(?=1)=P(B \mid A, \neg C)=\frac{P(A, B, \neg C)}{P(A, \neg C)}=\ldots=0$
M-step:

$$
P(B \mid A)=0.33
$$

$$
P(C \mid B)=1
$$

$P(A)=0.75 \quad P(B \mid \neg A)=1$
$P(C \mid \neg B)=0$
E-step: $P(?=1)=$

## Expectation Maximization

- Guess probabilities for nodes with missing values (e.g., based on other observations)
- Compute the probability distribution over the missing values, given our guess
- Update the probabilities based on the guessed values
- Repeat until convergence
- Guaranteed to converge to local optimum


## Learning Summary

- Known structure, fully observable: only need to do parameter estimation
- Unknown structure, fully observable: do heuristic/local search through structure space, then parameter estimation
- Known structure, missing values: use expectation maximization (EM) to estimate parameters
- Known structure, hidden variables: apply adaptive probabilistic network (APN) techniques
- Unknown structure, hidden variables: too hard to solve!


## Other Graphical Models

- Directed
- Bayesian Networks
- Undirected
- Markov Network (Markov Random Field)
- BN $\rightarrow$ MN (moralization: marry all co-parents)

- Mixed
- Chain Graph



## Other Graphical Models

Naïve Bayes


Conditional


Sequence

HMMs


General Graphs

Conditional

Generative directed models


Conditional

Logistic
Regression


Linear-chain CRFs


General CRFs


