# Machine Learning 10-601 

Tom M. Mitchell<br>Machine Learning Department<br>Carnegie Mellon University

January 21, 2015

Today:

- Bayes Rule
- Estimating parameters
- MLE
- MAP
some of these slides are derived from William Cohen, Andrew Moore, Aarti Singh, Eric Xing, Carlos Guestrin. - Thanks!

Readings:
Probability review

- Bishop Ch. 1 thru 1.2.3
- Bishop, Ch. 2 thru 2.2
- Andrew Moore' s online tutorial


## Announcements

- Class is using Piazza for questions/discussions about homeworks, etc.
- see class website for Piazza address
- http://www.cs.cmu.edu/~ninamf/courses/601sp15/
- Recitations thursdays 7-8pn, Wean 5409 ?
- videos for future recitations (class website)
- HW1 was accepted to Sunday 5pm for full credit
- HW2 out today on class website, due in 1 week
- HW3 will involve programming (in Octave )


# $P(A \mid B)=\frac{P(B \mid A){ }^{*} P(A)}{P(B)}$ Bayes' rule 

we call $P(A)$ the "prior"
and $P(A \mid B)$ the "posterior"
Bayes, Thomas (1763) An essay towards solving a problem in the doctrine of chances. Philosophical Transactions of the Royal Society of London, 53:370-418
...by no means merely a curious speculation in the doctrine of chances, but necessary to be solved in order to a sure foundation for all our reasonings concerning past facts, and what is likely to be hereafter.... necessary to be considered by any that would give a clear account of the strength of analogical or inductive reasoning...

Other Forms of Bayes Rule $\quad P(A \mid B)=\frac{P(B \mid A)^{* P(A)}}{P(B)}$

$$
P(A \mid B)=\frac{P(B \mid A) P(A)}{P(B \mid A) P(A)+P(B \mid \sim A) P(\sim A)}
$$

$$
P(A \mid B \wedge X)=\frac{P(B \mid A \wedge X) P(A \wedge X)}{P(B \wedge X)}
$$

## Applying Bayes Rule

$$
P(A \mid B)=\frac{P(B \mid A) P(A)}{P(B \mid A) P(A)+P(B \mid \sim A) P(\sim A)}
$$

$A=$ you have the flu, $B=$ you just coughed

$$
=0.17
$$

$\begin{aligned} & \text { Assume: } \\ & \begin{array}{l}P(A)=0.05 \\ P(B \mid A)=0.80\end{array}\end{aligned} \quad P(A \mid B)=\frac{.8 .05}{.8 .05+0.20 .95}$
$P(B \mid A)=0.80$
$P(A)=1-P(\neg A)$
what is $P($ flu $\mid$ cough $)=P(A \mid B)$ ?

# what does all this have to do with function approximation? 

## instead of $\mathrm{F}: \mathrm{X} \rightarrow \mathrm{Y}$, <br> learn <br> $P(Y \mid X)$

## The Joint Distribution

## Example: Boolean

 variables $A, B, C$Recipe for making a joint distribution of $M$ variables:

| $\mathbf{A}$ | $\mathbf{B}$ | $\mathbf{C}$ | Prob |
| :--- | :--- | :--- | :--- |
| $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{0 . 3 0}$ |
| $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{0 . 0 5}$ |
| $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{0 . 1 0}$ |
| $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{0 . 0 5}$ |
| $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{0}$ | $\mathbf{0 . 0 5}$ |
| $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{1}$ | $\mathbf{0 . 1 0}$ |
| $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{0}$ | $\mathbf{0 . 2 5}$ |
| $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{1}$ | $\mathbf{0 . 1 0}$ |

[A. Moore]

## The Joint Distribution

Example: Boolean variables $A, B, C$

Recipe for making a joint distribution of $M$ variables:

1. Make a truth table listing all combinations of values (M Boolean variables $\rightarrow 2^{\mathrm{M}}$ rows).

[A. Moore]

## The Joint Distribution

Example: Boolean variables $A, B, C$

Recipe for making a joint distribution of $M$ variables:

1. Make a truth table listing all combinations of values (M Boolean variables $\rightarrow 2^{\mathrm{M}}$ rows).
2. For each combination of values, say how probable it is.

| $A$ | $B$ | $C$ | Prob |
| :--- | :--- | :--- | :--- |
| 0 | 0 | 0 | 0.30 |
| 0 | 0 | 1 | 0.05 |
| 0 | 1 | 0 | 0.10 |
| 0 | 1 | 1 | 0.05 |
| 1 | 0 | 0 | 0.05 |
| 1 | 0 | 1 | 0.10 |
| 1 | 1 | 0 | 0.25 |
| 1 | 1 | 1 | 0.10 |


[A. Moore]

## The Joint Distribution

Example: Boolean variables $A, B, C$

Recipe for making a joint distribution of $M$ variables:

1. Make a truth table listing all combinations of values (M Boolean variables $\rightarrow 2^{\mathrm{M}}$ rows).
2. For each combination of values, say how probable it is.
3. If you subscribe to the axioms of probability, those probabilities must sum to 1 .

| $A$ | $B$ | $C$ | Prob |
| :--- | :--- | :--- | :--- |
| 0 | 0 | 0 | 0.30 |
| 0 | 0 | 1 | 0.05 |
| 0 | 1 | 0 | 0.10 |
| 0 | 1 | 1 | 0.05 |
| 1 | 0 | 0 | 0.05 |
| 1 | 0 | 1 | 0.10 |
| 1 | 1 | 0 | 0.25 |
| 1 | 1 | 1 | 0.10 |


[A. Moore]

## Using the Joint Distribution

| gender hours_worked wea |  |  |  |
| :---: | :---: | :---: | :---: |
| Female | v0:40.5- | poor | 0.253122 |
| Male |  | rich | 0.0245895 |
|  | V1:40.5+ | poor | 0.0421768 |
|  |  | rich | 0.0116293 |
|  | v0:40.5- | poor | 0.331313 |
|  |  | rich | 0.0971295 |
|  | V1:40.5+ | poor | 0.134106 |
|  |  | rich | 0.105933 |

One you have the JD you can ask for the
$P(E)=\quad \sum P($ row $)$
rows matching $E$ probability of any logical expression involving these variables
[A. Moore]

## Using the Joint

$\left.\begin{array}{|lllll|}\hline \text { gender } & \text { hours_worked } & \text { wealth } \\ \text { Female } & \text { v0:40.5- } & \text { poor } & 0.253122 & \\ & & \text { rich } & 0.0245895\end{array}\right]$
$P($ Poor Male $)=0.4654$

$$
P(E)=\sum_{\text {rows matching } E} P(\text { row })
$$

[A. Moore]

## Using the Joint

| gender | hours_worked | wealth |  |
| :--- | :--- | :--- | :--- |
| Female | v0:40.5- | poor | 0.253122 |
|  |  | rich | 0.0245895 |
|  | v1:40.5+ | poor | 0.0421768 |
|  | rich | 0.0116293 |  |
| Male | v0:40.5- | poor | 0.331313 |
|  |  | rich | 0.0971295 |
|  | v1:40.5+ | poor | 0.134106 |

$P($ Poor $)=0.7604$

$$
P(E)=\sum_{\text {rows matching } E} P(\text { row })
$$

[A. Moore]

## Inference with the Joint

| gender | hours_worked | wealth |  |
| :--- | :--- | :--- | :--- |
| Female | v0:40.5- | poor | 0.253122 |
|  |  | rich | 0.0245895 |
|  | v1:40.5+ | poor | 0.0421768 |
|  | rich | 0.0116293 |  |
|  |  | poor | 0.331313 |
|  | vich | 0.0971295 |  |
|  | vi:40.5+ | poor | 0.134106 |
|  | rich | 0.105933 |  |

$$
P\left(E_{1} \mid E_{2}\right)=\frac{P\left(E_{1} \wedge E_{2}\right)}{P\left(E_{2}\right)}=\frac{\sum_{\text {rows matching } E_{1} \text { and } E_{2}} P(\text { row })}{\sum_{\text {rows matching } E_{2}} P(\text { row })}
$$

$P($ Male $\mid$ Poor $)=0.4654 / 0.7604=0.612$
[A. Moore]

## Learning and the Joint Distribution

Suppose we want to learn the function $\mathrm{f}:<\mathrm{G}, \mathrm{H}>\rightarrow \mathrm{W}$
Equivalently, $\mathrm{P}(\mathrm{W} \mid \mathrm{G}, \mathrm{H})$
Solution: learn joint distribution from data, calculate $P(W \mid G, H)$
e.g., $\mathrm{P}(\mathrm{W}=$ rich $\mid \mathrm{G}=$ female, $\mathrm{H}=40.5-)=$

$$
\frac{P(w=r \wedge G=f \wedge H=40-)}{P( }=\frac{.024}{.27 \eta} \approx .09
$$

## sounds like the solution to <br> learning $\mathrm{F}: \mathrm{X} \rightarrow \mathrm{Y}$, or $P(Y \mid X)$.

Are we done?

## sounds like the solution to

 learning $F: X \rightarrow Y$,$$
\text { or } P(Y \mid X) . \quad 2^{10}=1024
$$

Main problem: learning $\mathrm{P}(\mathrm{Y} \mid \mathrm{X})$ can require more data than we have
consider learning Joint Dist. with 100 attributes \# of rows in this table? $2^{100} \geqslant 100^{10}=10^{30}$ \# of people on earth?
$10^{9}$
fraction of rows with 0 training examples? 0.9999

## What to do?

1. Be smart about how we estimate probabilities from sparse data

- maximum likelihood estimates
- maximum a posteriori estimates

2. Be smart about how to represent joint distributions

- Bayes networks, graphical models


## 1. Be smart about how we estimate probabilities

## Estimating Probability of Heads

- I show you the above coin $X$, and hire you to estimate the probability that it will turn up heads $(X=1)$ or tails $(X=0)$
- You flip it repeatedly, observing
- it turns up heads $\alpha_{1}$ times
- it turns up tails $\alpha_{0}$ times
- Your estimate for $P(X=1)$ is....?


Estimating $\theta=P(X=1)$


Test A:
100 flips: 51 Heads ( $X=1$ ), 49 Tails ( $X=0$ )

$$
\frac{\alpha_{1}}{\alpha_{1}+\alpha_{0}}=\frac{51}{100} \rightarrow \hat{P}(x=1)=0.51
$$

Test B:
3 flips: $\alpha_{\text {Heads }}(X=1), \stackrel{\alpha}{1}$ Tails $(X=0)$

$$
=\frac{2}{2+1}=0.666
$$

Estimating $\theta=P(X=1)$
Case C: (online learning)

- keep flipping, want single learning algorithm that gives reasonable estimate after each flip

$$
\begin{aligned}
& \alpha_{1}=\# \text { obs. heads }(x=1) \quad n=\alpha_{1}+\alpha_{0} \\
& \alpha_{0}=4 \text { obs } x=0 \\
& \beta_{1}=4 \text { hallucinated } x=1 \text { 's } \\
& \beta_{0}=4 \text { hallocinutad } x=0 \text { s } \\
& \frac{\alpha_{1}+10}{\left(\alpha_{1}+10\right)+\left(\alpha_{0}+10\right)} \rightarrow \frac{\left(\alpha_{1}+\beta_{1}\right)}{\left(\alpha_{1}+\beta_{1}\right)+\left(\alpha_{0}+\beta_{0}\right)}
\end{aligned}
$$

## Principles for Estimating Probabilities

Principle 1 (maximum likelihood):

- choose parameters $\theta$ that maximize $\mathbf{P ( d a t a | \theta )}$
- e.g.,

$$
\hat{\theta}^{M L E}=\frac{\alpha_{1}}{\alpha_{1}+\alpha_{0}}
$$

Principle 2 (maximum a posteriori prob.):

- choose parameters $\theta$ that maximize $\mathbf{P ( \theta | d a t a )}$
- e.g.

$$
\hat{\theta}^{M A P}=\frac{\alpha_{1}+\text { \#hallucinated_1s }}{\left(\alpha_{1}+\text { \#hallucinated_1s }\right)+\left(\alpha_{0}+\text { \#hallucinated_0s }\right)}
$$

## Maximum Likelihood Estimation

$P(X=1)=\underline{\theta} \quad P(X=0)=(1-\theta)$
$\begin{aligned} & \text { Data } \mathrm{D}:=\left\{\begin{array}{llll}1 & 0 & 0 & 1 \\ \uparrow & \uparrow & \uparrow & \uparrow\end{array}\right. \\ & P(\text { D| } \theta)\end{aligned}=\theta \cdot(1-\theta) \cdot(1-\theta) \cdot \theta \cdot \theta=\theta^{\alpha_{1}}(1-\theta)^{\alpha_{0}}$.

Flips produce data D with $\alpha_{1}$ heads, $\alpha_{0}$ tails

- flips are independent, identically distributed 1's and 0's (Bernoulli)
- $\alpha_{1}$ and $\alpha_{0}$ are counts that sum these outcomes (Binomial)

$$
P(D \mid \theta)=P\left(\alpha_{1}, \alpha_{0} \mid \theta\right)=\theta^{\alpha_{1}}(1-\theta)^{\alpha_{0}}
$$

## Maximum Likelihood Estimate for $\Theta$

$$
\begin{aligned}
\hat{\theta} & =\arg \max _{\theta} \ln P(\mathcal{D} \mid \theta) \\
& =\arg \max _{\theta} \ln \theta^{\alpha_{H}}(1-\theta)^{\alpha_{T}}
\end{aligned}
$$

- Set derivative to zero:

$$
\frac{d}{d \theta} \ln P(\mathcal{D} \mid \theta)=0
$$

[C. Guestrin]

$$
\begin{aligned}
& \hat{\theta}=\arg \max _{\theta} \ln P(D \mid \theta) \\
& =\arg \max _{\theta} \ln [\underbrace{\ln )}(1-\theta)^{\alpha_{0}}] \\
& \frac{\partial}{\partial \theta} \alpha_{1} \ln \theta+\alpha_{0} \ln (1-\theta) \\
& \alpha_{1} \frac{1}{\theta}+\alpha_{0} \frac{\partial \ln (1-\theta)}{\partial \theta} \\
& 0=\alpha_{1} \frac{1}{\theta}-\frac{\alpha_{0}}{1-\theta} \underbrace{\frac{\partial \ln (1-\theta)}{\partial(1-\theta)}}_{\frac{1}{1-\theta}} \cdot \underbrace{\frac{\partial(1-\theta)}{\partial \theta}}_{-1} \\
& \theta=\frac{\alpha_{1}}{\alpha_{1}+\alpha_{0}}
\end{aligned}
$$

## Summary: <br> Maximum Likelihood Estimate

- Each flip yields boolean value for $X$

(Bernoulli)

$$
X \sim \text { Bernoulli: } P(X)=\theta^{X}(1-\theta)^{(1-X)}
$$

- Data set $D$ of independent, identically distributed (iid) flips produces $\alpha_{1}$ ones, $\alpha_{0}$ zeros (Binomial)

$$
\begin{aligned}
& P(D \mid \theta)=P\left(\alpha_{1}, \alpha_{0} \mid \theta\right)=\theta^{\alpha_{1}}(1-\theta)^{\alpha_{0}} \\
& \hat{\theta}^{M L E}=\operatorname{argmax}_{\theta} P(D \mid \theta)=\frac{\alpha_{1}}{\alpha_{1}+\alpha_{0}}
\end{aligned}
$$

## Principles for Estimating Probabilities

Principle 1 (maximum likelihood):

- choose parameters $\theta$ that maximize $P($ data $\mid \theta)$

Principle 2 (maximum a posteriori prob.):

- choose parameters $\theta$ that maximize

$$
P(\theta \mid \text { data })=\frac{P(\text { data } \mid \theta) P(\theta)}{P(\text { data })}
$$

Beta prior distribution - $P(\theta)$

$$
P(\theta)=\frac{\theta^{\beta_{H}-1}(1-\theta)}{B\left(\beta_{H}, \beta_{T}\right)} \sim \operatorname{Beta}\left(\beta_{H}, \beta_{T}\right)
$$

- Likelihood function: $\quad P(\mathcal{D} \mid \theta)=\theta^{(\pi H}(1-\theta)^{\alpha_{T}}$
- Posterior: $P(\theta \mid \mathcal{D}) \propto P(\mathcal{D} \mid \theta) P(\theta)$

$$
\alpha \frac{\theta^{\left(\alpha^{2}+\beta_{H}-1\right.}}{}(1-\theta)^{\alpha_{T}+\beta_{T}-1}
$$

$$
\hat{\theta}^{\text {MAP }}=\frac{\left(\alpha_{H}+\beta_{H}-1\right)}{\left(\alpha_{H}+\beta_{1}-1\right)+\left(\alpha_{T}+\beta_{T}-1\right)}
$$

## Beta prior distribution - $\mathrm{P}(\theta)$

$$
P(\theta)=\frac{\theta^{\beta_{H}-1}(1-\theta)^{\beta_{T}-1}}{B\left(\beta_{H}, \beta_{T}\right)} \sim \operatorname{Beta}\left(\beta_{H}, \beta_{T}\right)
$$



[C. Guestrin]

Eg. 1 Coin flip problem
Likelihood is ~ Binomial

$$
P(\mathcal{D} \mid \theta)=\theta^{\alpha_{H}}(1-\theta)^{\alpha_{T}}
$$

If prior is Beta distribution,

$$
P(\theta)=\frac{\theta^{\beta_{H}-1}(1-\theta)^{\beta_{T}-1}}{B\left(\beta_{H}, \beta_{T}\right)} \sim \operatorname{Beta}\left(\beta_{H}, \beta_{T}\right)
$$

Then posterior is Beta distribution

$$
P(\theta \mid D) \sim \operatorname{Beta}\left(\alpha_{H}+\beta_{H}, \alpha_{H}+\beta_{H}\right)
$$

and MAP estimate is therefore

$$
\hat{\theta}^{M A P}=\frac{\alpha_{H}+\beta_{H}-1}{\left(\alpha_{H}+\beta_{H}-1\right)+\left(\alpha_{T}+\beta_{T}-1\right)}
$$

## Eg. 2 Dice roll problem (6 outcomes instead of 2)

Likelihood is $\sim$ Multinomial $\left(\theta=\left\{\theta_{1}, \theta_{2}, \ldots, \theta_{\mathrm{k}}\right\}\right)$

$$
P(\mathcal{D} \mid \theta)=\theta_{1}^{\alpha_{1}} \theta_{2}^{\alpha_{2}} \ldots \theta_{k}^{\alpha_{k}}
$$

If prior is Dirichlet distribution,

$$
P(\theta)=\frac{\theta_{1}^{\beta_{1}-1} \theta_{2}^{\beta_{2}-1} \ldots \theta_{k}^{\beta_{k}-1}}{B\left(\beta_{1}, \ldots, \beta_{k}\right)} \sim \operatorname{Dirichlet}\left(\beta_{1}, \ldots, \beta_{k}\right)
$$

Then posterior is Dirichlet distribution

$$
P(\theta \mid D) \sim \operatorname{Dirichlet}\left(\beta_{1}+\alpha_{1}, \ldots, \beta_{k}+\alpha_{k}\right)
$$

and MAP estimate is therefore

$$
\hat{\theta}_{i}^{M A P}=\frac{\alpha_{i}+\beta_{i}-1}{\sum_{j=1}^{k}\left(\alpha_{j}+\beta_{j}-1\right)}
$$

## Some terminology

- Likelihood function: $P($ data $\mid \theta)$
- Prior: P( $\theta$ )
- Posterior: P( $\theta$ | data)
- Conjugate prior: $P(\theta)$ is the conjugate prior for likelihood function $P($ data $\mid \theta)$ if the forms of $P(\theta)$ and $P(\theta \mid$ data) are the same.


## You should know

- Probability basics
- random variables, conditional probs, ...
- Bayes rule
- Joint probability distributions
- calculating probabilities from the joint distribution
- Estimating parameters from data
- maximum likelihood estimates
- maximum a posteriori estimates
- distributions - binomial, Beta, Dirichlet, ...
- conjugate priors


## Extra slides

## Independent Events

- Definition: two events $A$ and $B$ are independent if $P\left(A^{\wedge} B\right)=P(A)^{*} P(B)$
- Intuition: knowing A tells us nothing about the value of $B$ (and vice versa)


## Picture "A independent of B"

## Expected values

Given a discrete random variable $X$, the expected value of $X$, written $E[X]$ is

$$
E[X]=\sum_{x \in \mathcal{X}} x P(X=x)
$$

Example:

| $\mathbf{X}$ | $\mathbf{P}(\mathbf{X})$ |
| :---: | :---: |
| 0 | 0.3 |
| 1 | 0.2 |
| 2 | 0.5 |

## Expected values

Given discrete random variable $X$, the expected value of $X$, written $E[X]$ is

$$
E[X]=\sum_{x \in \mathcal{X}} x P(X=x)
$$

We also can talk about the expected value of functions of $X$

$$
E[f(X)]=\sum_{x \in \mathcal{X}} f(x) P(X=x)
$$

## Covariance

Given two discrete r.v.'s $X$ and $Y$, we define the covariance of $X$ and $Y$ as

$$
\operatorname{Cov}(X, Y)=E[(X-E(X))(Y-E(Y))]
$$

e.g., $X=$ gender, $Y=$ playsFootball
or $X=$ gender, $Y=l e f t H a n d e d$

Remember: $E[X]=\sum_{x \in \mathcal{X}} x P(X=x)$

