# Introduction to Bayesian methods (continued) - Lecture 16 

## David Sontag <br> New York University

Slides adapted from Luke Zettlemoyer, Carlos Guestrin, Dan Klein, and Vibhav Gogate

## Outline of lectures

- Review of probability
(After midterm)
Maximum likelihood estimation
2 examples of Bayesian classifiers:
- Naïve Bayes
- Logistic regression


## Bayes' Rule

- Two ways to factor a joint distribution over two variables:

$$
P(x, y)=P(x \mid y) P(y)=P(y \mid x) P(x)
$$

- Dividing, we get:

$$
P(x \mid y)=\frac{P(y \mid x)}{P(y)} P(x)
$$

- Why is this at all helpful?

- Let's us build one conditional from its reverse
- Often one conditional is tricky but the other one is simple
- Foundation of many practical systems (e.g. ASR, MT)
- In the running for most important ML equation!


## Returning to thumbtack example...

- $\mathrm{P}($ Heads $)=\theta, \mathrm{P}($ Tails $)=1-\theta$

- Flips are i.i.d.: $D=\left\{x_{i} \mid i=1 \ldots n\right\}, P(D \mid \theta)=\Pi_{i} P\left(x_{i} \mid \theta\right)$
- Independent events
- Identically distributed according to Bernoulli distribution
- Sequence $D$ of $\alpha_{H}$ Heads and $\alpha_{T}$ Tails

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

Called the "likelihood" of the data under the model

## Maximum Likelihood Estimation

- Data: Observed set $D$ of $\alpha_{H}$ Heads and $\alpha_{T}$ Tails
- Hypothesis: Bernoulli distribution
- Learning: finding $\theta$ is an optimization problem
- What's the objective function?

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

- MLE: Choose $\theta$ to maximize probability of $D$

$$
\begin{aligned}
\hat{\theta} & =\arg \max _{\theta} P(\mathcal{D} \mid \theta) \\
& =\arg \max _{\theta} \ln P(\mathcal{D} \mid \theta)
\end{aligned}
$$

## Your first parameter learning algorithm

$\widehat{\theta}=\arg \max _{\theta} \ln P(\mathcal{D} \mid \theta)$
$=\arg \max _{\theta} \ln \theta^{\alpha_{H}}(1-\theta)^{\alpha_{T}}$

- Set derivative to zero, and solve!

$$
\begin{aligned}
& \frac{d}{d \theta} \ln P(\mathcal{D} \mid \theta)=\frac{d}{d \theta}\left[\ln \theta^{\alpha_{H}}(1-\theta)^{\alpha_{T}}\right] \\
& \quad=\frac{d}{d \theta}\left[\alpha_{H} \ln \theta+\alpha_{T} \ln (1-\theta)\right] \\
& \quad=\alpha_{H} \frac{d}{d \theta} \ln \theta+\alpha_{T} \frac{d}{d \theta} \ln (1-\theta) \\
& \quad=\frac{\alpha_{H}}{\theta}-\frac{\alpha_{T}}{1-\theta}=0 \quad \widehat{\theta}_{M L E}=\frac{\alpha_{H}}{\alpha_{H}+\alpha_{T}}
\end{aligned}
$$

## Data



$$
L(\theta ; \mathcal{D})=\ln P(\mathcal{D} \mid \theta)
$$



## What if I have prior beliefs?

- Billionaire says: Wait, I know that the thumbtack is "close" to 50-50. What can you do for me now?
- You say: I can learn it the Bayesian way...
- Rather than estimating a single $\theta$, we obtain a distribution over possible values of $\theta$

In the beginning


After observations



## Bayesian Learning

- Use Bayes' rule!
$\left.\left.P(\theta \mid \mathcal{D})=\frac{\text { Data Likelihood }}{\downarrow}{ }^{\downarrow} \right\rvert\, \theta\right) P(\theta)$
$P(\mathcal{D})$
- Or equivalently: $P(\theta \mid \mathcal{D}) \propto P(\mathcal{D} \mid \theta) P(\theta)$
- For uniform priors, this reduces to maximum likelihood estimation!

$$
P(\theta) \propto 1 \quad P(\theta \mid \mathcal{D}) \propto P(\mathcal{D} \mid \theta)
$$

## Bayesian Learning for Thumbtacks

$$
P(\theta \mid \mathcal{D}) \propto P(\mathcal{D} \mid \theta) P(\theta)
$$

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

- What should the prior be?
- Represent expert knowledge
- Simple posterior form
- For binary variables, commonly used prior is the Beta distribution:

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

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

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





- Since the Beta distribution is conjugate to the Bernoulli distribution, the posterior distribution has a particularly simple form:
$P(\theta \mid \mathcal{D}) \propto P(\mathcal{D} \mid \theta) P(\theta)$

$$
\begin{aligned}
& \propto \theta^{\alpha_{H}}(1-\theta)^{\alpha_{T}} \theta^{\beta_{H}-1}(1-\theta)^{\beta_{T}-1} \\
& \begin{aligned}
=\theta^{\alpha_{H}+\beta_{H}-1}(1-\theta)^{\alpha_{T}+\beta_{T}-1} \\
\quad=\operatorname{Beta}\left(\alpha_{H}+\beta_{H}, \alpha_{T}+\beta_{T}\right)
\end{aligned}
\end{aligned}
$$

## Using Bayesian inference for prediction

- We now have a distribution over parameters
- For any specific $f$, a function of interest, compute the expected value of $f$ :

$$
E[f(\theta)]=\int_{0}^{1} f(\theta) P(\theta \mid \mathcal{D}) d \theta
$$

- Integral is often hard to compute
- As more data is observed, posterior is more concentrated
- MAP (Maximum a posteriori approximation): use most likely parameter to approximate the expectation

$$
\begin{gathered}
\hat{\theta}=\arg \max _{\theta} P(\theta \mid \mathcal{D}) \\
E[f(\theta)] \approx f(\widehat{\theta})
\end{gathered}
$$

## Outline of lectures

- Review of probability
- Maximum likelihood estimation

2 examples of Bayesian classifiers:

- Naïve Bayes
- Logistic regression


## Bayesian Classification

- Problem statement:
- Given features $X_{1}, X_{2}, \ldots, X_{n}$
- Predict a label $Y$
[Next several slides adapted from:
Vibhav Gogate, Jonathan Huang, Luke Zettlemoyer, Carlos Guestrin, and Dan Weld]


## Example Application

- Digit Recognition

- $\mathrm{X}_{1}, \ldots, \mathrm{X}_{\mathrm{n}} \in\{0,1\}$ (Black vs. White pixels)
- $\mathrm{Y} \in\{\mathbf{0 , 1 , 2 , 3 , 4 , 5 , 6 , 7 , 8 , 9 \}}$


## The Bayes Classifier

- If we had the joint distribution on $\mathbf{X}_{1}, \ldots, \mathbf{X}_{\mathrm{n}}$ and $\mathbf{Y}$, could predict using:

$$
\arg \max _{Y} P\left(Y \mid X_{1}, \ldots, X_{n}\right)
$$

- (for example: what is the probability that the image represents a 5 given its pixels?)
- So ... How do we compute that?


## The Bayes Classifier

- Use Bayes Rule!

- Why did this help? Well, we think that we might be able to specify how features are "generated" by the class label


## The Bayes Classifier

- Let's expand this for our digit recognition task:

$$
\begin{aligned}
& P\left(Y=5 \mid X_{1}, \ldots, X_{n}\right)=\frac{P\left(X_{1}, \ldots, X_{n} \mid Y=5\right) P(Y=5)}{P\left(X_{1}, \ldots, X_{n} \mid Y=5\right) P(Y=5)+P\left(X_{1}, \ldots, X_{n} \mid Y=6\right) P(Y=6)} \\
& P\left(Y=6 \mid X_{1}, \ldots, X_{n}\right)=\frac{P\left(X_{1}, \ldots, X_{n} \mid Y=6\right) P(Y=6)}{P\left(X_{1}, \ldots, X_{n} \mid Y=5\right) P(Y=5)+P\left(X_{1}, \ldots, X_{n} \mid Y=6\right) P(Y=6)}
\end{aligned}
$$

- To classify, we'll simply compute these probabilities, one per class, and predict based on which one is largest


## Model Parameters

- How many parameters are required to specify the likelihood, $P\left(X_{1}, \ldots, X_{n} \mid Y\right)$ ?
- (Supposing that each image is $30 \times 30$ pixels)
- The problem with explicitly modeling $P\left(X_{1}, \ldots, X_{n} \mid Y\right)$ is that there are usually way too many parameters:
- We'll run out of space
- We'll run out of time
- And we'll need tons of training data (which is usually not available)


## Naïve Bayes

- Naïve Bayes assumption:
- Features are independent given class:

$$
\begin{aligned}
P\left(X_{1}, X_{2} \mid Y\right) & =P\left(X_{1} \mid X_{2}, Y\right) P\left(X_{2} \mid Y\right) \\
& =P\left(X_{1} \mid Y\right) P\left(X_{2} \mid Y\right)
\end{aligned}
$$

- More generally:

$$
P\left(X_{1} \ldots X_{n} \mid Y\right)=\prod_{i} P\left(X_{i} \mid Y\right)
$$

- How many parameters now?
- Suppose $\mathbf{X}$ is composed of $n$ binary features


## The Naïve Bayes Classifier

- Given:
- Prior P(Y)
$-n$ conditionally independent features $X_{1}, \ldots, X_{1}$, given the class $Y$
- For each feature $i$, we specify $P\left(X_{i} \mid Y\right)$

- Classification decision rule:

$$
\begin{aligned}
y^{*}=h_{N B}(\mathbf{x}) & =\arg \max _{y} P(y) P\left(x_{1}, \ldots, x_{n} \mid y\right) \\
& =\arg \max _{y} P(y) \prod_{i} P\left(x_{i} \mid y\right)
\end{aligned}
$$

If certain assumption holds, NB is optimal classifier! (they typically don't)

## A Digit Recognizer

- Input: pixel grids

- Output: a digit 0-9

Are the naïve Bayes assumptions realistic here?

## What has to be learned?

$P(Y)$

| 1 | 0.1 |
| :--- | :--- |
| 2 | 0.1 |
| 3 | 0.1 |
| 4 | 0.1 |
| 5 | 0.1 |
| 6 | 0.1 |
| 7 | 0.1 |
| 8 | 0.1 |
| 9 | 0.1 |
| 0 | 0.1 |



## MLE for the parameters of NB

- Given dataset
- Count $(A=a, B=b) \leftarrow$ number of examples where $A=a$ and B=b
- MLE for discrete NB, simply:
- Prior:

$$
P(Y=y)=\frac{\operatorname{Count}(Y=y)}{\sum_{y^{\prime}} \operatorname{Count}\left(Y=y^{\prime}\right)}
$$

- Observation distribution:

$$
P\left(X_{i}=x \mid Y=y\right)=\frac{\operatorname{Count}\left(X_{i}=x, Y=y\right)}{\sum_{x^{\prime}} \operatorname{Count}\left(X_{i}=x^{\prime}, Y=y\right)}
$$

## MLE for the parameters of NB

- Training amounts to, for each of the classes, averaging all of the examples together:



## MAP estimation for NB

- Given dataset
- Count $(A=a, B=b) \leftarrow$ number of examples where $A=a$ and $B=b$
- MAP estimation for discrete NB, simply:
- Prior:

$$
P(Y=y)=\frac{\operatorname{Count}(Y=y)}{\sum_{y^{\prime}} \operatorname{Count}\left(Y=y^{\prime}\right)}
$$

- Observation distribution:
$P\left(X_{i}=x \mid Y=y\right)=\frac{\operatorname{Count}\left(X_{i}=x, Y=y\right)+\mathbf{a}}{\sum_{x^{\prime}} \operatorname{Count}\left(X_{i}=x^{\prime}, Y=y\right)+|\mathbf{X} \mathbf{i}|^{*} \mathbf{a}}$
- Called "smoothing". Corresponds to Dirichlet prior!

