## Point Estimation

Zhang Hongxin zhx@cad.zju.edu.cn

State Key Lab of CAD\&CG, ZJU
2017-02-28

## What you need to know

－Point estimation：（点估计）
－Maximal Likelihood Estimation（MLE）
－Bayesian learning
－Maximize A Posterior（MAP）
－Gaussian estimation
－Regression（回归）
－Basis function＝features
－Optimizing sum squared error
－Relationship between regression and Gaussians
－Bias－Variance trade－off


点估计是非常重要的数据计算技术，是需要在入门阶段重点掌握的 Point estimation

## Your first consulting job

- An IT billionaire from Beijing asks you a question:
- B: I have thumbtack, if I flip it, what's the probability it will fall with the nail up?
- Y: Please flip it a few times ...

- B: Why???
- Y: Because...


## Binomial Distribution

- $P($ Heads $)=\theta, P($ Tails $)=1-\theta \quad D=\{T, H, H, T, T\}$

$$
P(D \mid \theta)=(1-\theta) \theta \theta(1-\theta)(1-\theta)
$$

- Flips are i.i.d. (Independent Identically distributed)
- Independent events
- Identically distributed according to Binomial distribution
- Sequence D of $\alpha_{H}$ Heads and $\alpha_{T}$ Tails

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

## Maximum Likelihood Estimation

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

$$
D=\{T, H, H, T, T\}
$$

- MLE: Choose $\theta$ that maximizes the probability of observed data:

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

## Maximum Likelihood Estimation (cont.)

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

- Set derivative to zero:

$$
\frac{d}{d \theta} \ln P(D \mid \theta)=0 \quad \hat{\theta}=\frac{\alpha_{H}}{\alpha_{H}+\alpha_{T}}=\frac{2}{2+3}
$$

## How many flips do I need?

$$
\hat{\theta}=\frac{\alpha_{H}}{\alpha_{H}+\alpha_{T}}
$$

- B: I flipped 2 heads and 3 tails.
- Y: 1- $\theta=3 / 5$, I can prove it!
- B: What if I flipped 20 heads and 30 tails?
- Y: Same answer, I can prove it!
- B: What's better?
- Y: Humm... The more the merrier???
- B: Is this why I am paying you the big bucks???


## Simple bound (based on Höffding's inequality)

- For $N=\alpha_{H}+\alpha_{T}$ and $\hat{\theta}=\frac{\alpha_{T}}{\alpha_{H}+\alpha_{T}}$ http://omega.albany.edu:8008/machine-learning-dir/notes-dir/vc1/vc-I.html
- Let $\theta^{*}$ be the true parameter, for any $\varepsilon>0$ :

$$
\begin{aligned}
& P\left(\left|\hat{\theta}-\theta^{*}\right| \geq \varepsilon\right) \leq 2 e^{-2 N \varepsilon^{2}} \leq \\
& N \geq \frac{1}{2 \varepsilon^{2}}[\ln 2-\ln \delta] \\
& N \geq 270 ;(\varepsilon=0.1, \delta=0.01)
\end{aligned}
$$

## PAC Learning

- PAC: Probably Approximate Correct
- B: I want to know the thumbtack parameter $\theta$, within $\varepsilon=0.1$, with probability at least $1-\delta=$ 0.99 . How many flips?
- Y: 270, ©


## Prior: <br> knowledge before experiments

- B: Wait, I know that the thumbtack is "close" to 50-50. What can you ...?
- Y: I can learn it the Bayesian way...
- Rather than estimating a single $\theta$, we obtain a distribution over possible values of $\theta$



## Bayesian Learning

- Bayes rule:

$$
\text { Posterior } \rightarrow P(\theta \mid D)=\frac{P(\theta) P(D \mid \theta)}{P(D) \longleftarrow \text { Data distribution }}
$$

- Or equivalently: evidence
(Normalization)

$$
P(\theta \mid D) \propto P(\theta) P(D \mid \theta)
$$

## Probability Theory


-Sum Rule

$$
\begin{aligned}
& p\left(X=x_{i}\right)=\frac{c_{i}}{N}=\frac{1}{N} \sum_{j=1}^{L} n_{i j} \\
& \quad=\sum_{j=1}^{L} p\left(X=x_{i}, Y=y_{j}\right)
\end{aligned}
$$

Product Rule

$$
\begin{aligned}
p\left(X=x_{i}, Y=y_{j}\right) & =\frac{n_{i j}}{N}=\frac{n_{i j}}{c_{i}} \cdot \frac{c_{i}}{N} \\
& =p\left(Y=y_{j} \mid X=x_{i}\right) p\left(X=x_{i}\right)
\end{aligned}
$$

## Probability concepts

－Random variables：$x$
－Probability（function）：$P(X \leq x), P(x)$
－Density（function）：$f(x)$ ，
－Independency：$P(x, y)=P(x) P(y)$
－Feature quantities：
－Mean，expectation $E(x)=\int x f(x) \mathrm{d} x$
－Covariance
－ $\operatorname{cov}(x, y)=0$ ，uncorrelatedness／irrelevant（统计无关）
－Higher order moments

## The Rules of Probability

- Sum Rule

$$
p(X)=\sum_{Y} p(X, Y)
$$

- Product Rule

$$
p(X, Y)=p(Y \mid X) p(X)
$$

## Bayes' Theorem

$$
\begin{aligned}
p(Y \mid X) & =\frac{p(X \mid Y) p(Y)}{p(X)} \\
p(X) & =\sum_{Y} p(X \mid Y) p(Y)
\end{aligned}
$$

posterior $\propto$ likelihood $\times$ prior

## Bayesian Learning in our case

－Likelihood function is simply Binomial：

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

－What about prior？
－Represent expert knowledge
－Simple posterior form
－Conjugate priors：（共轭先验）
－Closed－form representation of posterior
－For Binomial，conjugate prior is Beta distribution

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

- Prior: Beta distribution $\Gamma(x+1)=x \Gamma(x), \Gamma(1)=1$

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

- Likelihood: Binomial distribution

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

- Posterior:

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

## Using Bayesian posterior

- Posterior distribution:

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

- Bayesian inference:
- No longer single parameter:

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

- Integral, :


## Expectation（数学期望）

－Random variable：$\theta$
－Random function：$f(\theta)$
－Expectation：

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

## MAP:

## Maximum a posteriori approximation

$$
\begin{aligned}
& P(\theta \mid D) \sim \operatorname{Beta}\left(\alpha_{H}+\beta_{H}, \alpha_{T}+\beta_{T}\right) \\
& E[f(\theta)]=\int_{0}^{1} f(\theta) P(\theta \mid D) d \theta \longleftarrow
\end{aligned}
$$

- MAP: use most likely parameter

$$
\widehat{\theta}=\underset{\theta}{\arg \max } P(\theta \mid D) \quad E[f(\theta)] \approx f(\hat{\theta}) .
$$

## MAP for Beta distribution

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

- MAP: use most likely parameter

$$
\hat{\theta}=\underset{\theta}{\arg \max } P(\theta \mid D)=\frac{\alpha_{T}+\beta_{T}-1}{\alpha_{H}+\beta_{H}+\alpha_{T}+\beta_{T}-2}
$$

- Beta prior equivalent to extra thumbtack flips
- As $N=\alpha_{T}+\alpha_{H} \rightarrow \infty$, prior is "forgotten"
- But, for small sample size, prior is important!


## More ...

- B: Can we handle more complex cases?
- Y: Yes, :-D
- Prior: a mixture of beta distribution
- $P(\theta) \sim 0.4 \operatorname{Beta}(20,1)+0.4 \operatorname{Beta}(1,20)+0.2 \operatorname{Beta}(2,2)$


## Multinomial distribution

- B: Now if I give you a dice (骰子), then ...
- Y: I can solve this problem in a similar way.
- Likelihood:

$$
\begin{aligned}
& P\left(X=x^{k} \mid \boldsymbol{\theta}\right)=\theta_{k}, k=1,2, \ldots, r, \\
& \boldsymbol{\theta}=\left\{\theta_{1}, \ldots, \theta_{r}\right\}, \theta_{1}+\ldots+\theta_{r}=1 \\
& D=\left\{X_{1}=x_{1}, \ldots, X_{N}=x_{N}\right\} \Rightarrow\left\{N_{1}, \ldots, N_{r}\right\}
\end{aligned}
$$

$$
P(D \mid \boldsymbol{\theta})=\prod_{i=1}^{r} \theta_{i}^{N_{i}}
$$

## Multinomial distribution

- Conjugate prior (Dirichlet distribution):

$$
P(\boldsymbol{\theta})=\operatorname{Dir}\left(\boldsymbol{\theta} \mid \alpha_{1}, \ldots, \alpha_{r}\right)=\frac{\Gamma(\alpha)}{\prod_{k=1}^{r} \Gamma\left(\alpha_{k}\right)} \prod_{k=1}^{r} \theta_{k}^{\alpha_{k}-1}, \quad \alpha=\sum_{k=1}^{r} \alpha_{k}
$$

- Solution:

$$
P\left(X_{N+1}=x^{k} \mid D\right)=\int \theta_{k} \operatorname{Dir}\left(\boldsymbol{\theta} \mid \alpha_{1}+N_{1}, \ldots, \alpha_{r}+N_{r}\right) d \boldsymbol{\theta}=\frac{\alpha_{k}+N_{k}}{\alpha+N}
$$

- Important fact:

$$
P(D)=\frac{\Gamma(\alpha)}{\Gamma(\alpha+N)} \prod_{k=1}^{r} \frac{\Gamma\left(\alpha_{k}+N_{k}\right)}{\Gamma\left(\alpha_{k}\right)}
$$

## Beyond MAP: <br> Approximate Bayesian Computation

- Approximate Bayesian Computational methods [ARXIV]
- Approximate Bayesian Computation (ABC) in practice [PDF]

(1) Compute summary statistic $\mu$ from observational data
(3) Compute summary statistic $\mu_{i}$ for each simulation

$$
\rho\left(\mu_{i}, \mu\right) \stackrel{?}{\leq} \varepsilon
$$

(4) Based on a distance $\rho(\cdot \cdot \cdot$ and a tolerance $\varepsilon$, decide for each simulation whether its summary statistic is sufficiently close to that of the observed data.

$\mu_{1}$

$\mu_{2}$

$\mu_{3}$
(2) Given a certain model, perform $n$ simulations, each with a parameter drawn from the prior distribution

$\mu_{n}$

Posterior distribution of model parameter $\theta$
(5) Approximate the posterior distribution of $\theta$ from the distribution of parameter values $\theta_{i}$ associated with accepted simulations.


在数据处理中，高斯分布无处不在

## Gaussian distribution

## Gaussian distribution

## 均值

Continuous random variable：

$$
P(x \mid \mu, \delta) \sim \frac{1}{\delta \sqrt{2 \pi}} e^{-\frac{(x-\mu)^{2}}{2 \delta^{2}}}
$$

方差

Consider the difference between continuous and discrete variables？

## MLE for Gaussian

- Prob. of i.i.d. samples $D=\left\{x_{1}, x_{2}, \ldots, x_{N}\right\}$
likelihood $P(D \mid \mu, \sigma)=\left(\frac{1}{\sigma \sqrt{2 \pi}}\right)^{N} \prod_{i=1}^{N} e^{\frac{-\left(x_{i}-\mu\right)^{2}}{2 \sigma^{2}}}$
- The magic of log (to log-likelihood)

$$
\begin{aligned}
\ln P(D \mid \mu, \sigma) & =\ln \left(\frac{1}{\sigma \sqrt{2 \pi}}\right)^{N} \prod_{i=1}^{N} e^{\frac{-\left(x_{i}-\mu\right)^{2}}{2 \sigma^{2}}} \\
& =-N \ln (\sigma \sqrt{2 \pi})-\sum_{i=1}^{N} \frac{\left(x_{i}-\mu\right)^{2}}{2 \sigma^{2}}
\end{aligned}
$$

## MLE for mean of a Gaussian

$$
\begin{aligned}
\frac{\partial}{\partial \mu} \ln P(D \mid \mu, \sigma) & =\frac{\partial}{\partial \mu} \ln \left(\frac{1}{\sigma \sqrt{2 \pi}}\right)^{N} \prod_{i=1}^{N} e^{\frac{-\left(x_{i}-\mu\right)^{2}}{2 \sigma^{2}}} \\
& =\frac{\partial}{\partial \mu}-\sum_{i=1}^{N} \frac{\left(x_{i}-\mu\right)^{2}}{2 \sigma^{2}} \\
& =\sum_{i=1}^{N} \frac{\left(x_{i}-\mu\right)}{\sigma^{2}}=0 \\
\mu & =\frac{1}{N} \sum_{i} x_{i}
\end{aligned}
$$

## MLE for variance of a Gaussian

$$
\begin{aligned}
\frac{\partial}{\partial \sigma} \ln P(D \mid \mu, \sigma) & =\frac{\partial}{\partial \sigma} \ln \left(\frac{1}{\sigma \sqrt{2 \pi}}\right)^{N} \prod_{i=1}^{N} e^{\frac{-\left(x_{i}-\mu\right)^{2}}{2 \sigma^{2}}} \\
& =\frac{\partial}{\partial \sigma}[-N \ln \sigma \sqrt{2 \pi}]-\sum_{i=1}^{N} \frac{\partial}{\partial \sigma}\left[\frac{\left(x_{i}-\mu\right)^{2}}{2 \sigma^{2}}\right] \\
& =-\frac{N}{\sigma}+\sum_{i=1}^{N} \frac{(x-\mu)^{2}}{\sigma^{3}}=0 \\
\sigma^{2} & =\frac{1}{N} \sum_{i}\left(x_{i}-\mu\right)^{2}
\end{aligned}
$$

## Gaussian parameters learning

- MLE

$$
\begin{aligned}
\hat{\mu} & =\frac{1}{N} \sum_{i} x_{i} \\
\hat{\sigma}^{2} & =\frac{1}{N} \sum_{i}\left(x_{i}-\mu\right)^{2}
\end{aligned}
$$

- Bayesian learning: prior?
- Conjugate priors:
- Mean: Gaussian priors
- Variance: Wishart Distribution


回归的思想，对于理解数据本身的特点与规律，是一种利器

## Regression problems

## Prediction of continuous variable

- B: Wait, that's not what I meant!
- Y: Chill out, dude.
- B: I want to predict a continuous variable for continuous inputs: I want to predict salaries from GPA.
- Y: I can regress that...



## The regression problem

- Instances: < $\mathbf{x}_{i}, t_{i}>$
- Learn: mapping from $\mathbf{x}$ to $t(\mathbf{x})$.
- Hypothesis space: $t(\mathbf{x}) \approx \hat{f}(x)=\sum_{i=1}^{k} w_{i} h_{i}$
- Given, basis functions $H=\left\{h_{1}, \ldots, \bar{h}_{k}\right\}$
- Find coefficients $\mathbf{w}=\left\{w_{1}, \ldots, w_{k}\right\}$
- Problem formulation:

$$
\mathbf{w}^{*}=\underset{\mathbf{w}}{\arg \min } \sum_{j}\left[t\left(\mathbf{x}_{j}\right)-\sum_{i=1}^{k} w_{i} h_{i}(x)\right]^{2}
$$

## But, why sum squared error?

- Model:

$$
P(t \mid \mathbf{x}, \mathbf{w}, \sigma)=\frac{1}{\sigma \sqrt{2 \pi}} e^{-\frac{-\left[t-\sum_{i} w_{i} h_{i}(x)\right]^{2}}{2 \sigma^{2}}}
$$

- Learn w using MLE


## Maximizing log-likelihood

$$
\begin{aligned}
\ln P(D \mid \mathbf{w}, \sigma) & =\ln \prod_{j}\left(\frac{1}{\sigma \sqrt{2 \pi}} e^{\frac{-\left[t_{j}-\sum_{i} w_{i} h_{i}\left(x_{j}\right)\right]^{2}}{2 \sigma^{2}}}\right) \\
& \Rightarrow \min \sum_{j} \frac{-\left[t_{j}-\sum_{i} w_{i} h_{i}\left(x_{j}\right)\right]^{2}}{2 \sigma^{2}}
\end{aligned}
$$

## Bias-Variance Tradeoff

- Choice of hypothesis basis introduce learning bias:
- More complex basis:
- Less bias
- More variance (over-fitting)



Example

$$
\begin{array}{l|l|l|l|l}
\hline 0 & 1 & 2 & 3 & 4 \\
5 & 6 & 7 & 8 & 9
\end{array}
$$

## Polynomial Curve Fitting



$$
y(x, \mathbf{w})=w_{0}+w_{1} x+w_{2} x^{2}+\ldots+w_{M} x^{M}=\sum_{j=0}^{M} w_{j} x^{j}
$$

## Sum-of-Squares Error Function



$$
E(\mathbf{w})=\frac{1}{2} \sum_{n=1}^{N}\left\{y\left(x_{n}, \mathbf{w}\right)-t_{n}\right\}^{2}
$$

## $0^{\text {th }}$ Order Polynomial



## $1^{\text {st }}$ Order Polynomial



## $3^{\text {rd }}$ Order Polynomial



## 9th $^{\text {th }}$ Order Polynomial



## Over-fitting



Root-Mean-Square (RMS) Error: $E_{\text {RMS }}=\sqrt{2 E\left(\mathbf{w}^{\star}\right) / N}$

## Polynomial Coefficients

|  | $M=0$ | $M=1$ | $M=3$ | $M=9$ |
| ---: | ---: | ---: | ---: | ---: |
| $w_{0}^{\star}$ | 0.19 | 0.82 | 0.31 | 0.35 |
| $w_{1}^{\star}$ |  | -1.27 | 7.99 | 232.37 |
| $w_{2}^{\star}$ |  |  | -25.43 | -5321.83 |
| $w_{3}^{\star}$ |  |  | 17.37 | 48568.31 |
| $w_{4}^{\star}$ |  |  |  | -231639.30 |
| $w_{5}^{\star}$ |  |  |  | 640042.26 |
| $w_{6}^{\star}$ |  |  |  | -1061800.52 |
| $w_{7}^{\star}$ |  |  |  | 1042400.18 |
| $w_{8}^{\star}$ |  |  |  | -557682.99 |
| $w_{9}^{\star}$ |  |  |  | 125201.43 |

## Data Set Size: $N=15$

$9^{\text {th }}$ Order Polynomial


## Data Set Size: $N=100$

$9^{\text {th }}$ Order Polynomial


## Regularization

- Penalize large coefficient values

$$
\widetilde{E}(\mathbf{w})=\frac{1}{2} \sum_{n=1}^{N}\left\{y\left(x_{n}, \mathbf{w}\right)-t_{n}\right\}^{2}+\frac{\lambda}{2}\|\mathbf{w}\|^{2}
$$

## Regularization: $\ln \lambda=-18$



## Regularization: $\ln \lambda=0$



## Regularization: $E_{\text {RMS }}$ vs. $\ln \lambda$



## What you need to know

- Point estimation:
- Maximal Likelihood Estimation
- Bayesian learning
- Maximal a Posterior
- Gaussian estimation
- Regression
- Basis function = features
- Optimizing sum squared error
- Relationship between regression and Gaussians
- Bias-Variance trade-off


## Homework01

- Python programming
- 1-D regression
- Learn TensorFlow
- Finish the "Gaussian parameters learning"
- Please use google, ^_*


## The End

新浪溦溥：＠
微信公众号：


