Statistics

Bayesian Hypothesis Testing
Consider $H_{0}: \theta=\theta_{0}$
vs. $H_{1}: \theta \neq \theta_{0}$

Technique: Put a prior on $\theta$ and $H_{0}$, then compute : $\mathbb{P}\left(H_{0} \mid \vec{X}\right)$
${ }^{2}$ this 15 our observe data.

Ex: Put the prior $\mathbb{P}\left(H_{0}\right)=\frac{1}{2}, \mathbb{P}\left(H_{1}\right)=1 / 2$.
Thin compute $\mathbb{P}\left(H_{0} \mid \vec{X}\right)$

$$
\begin{aligned}
\mathbb{P}\left(H_{0} \mid \vec{x}\right) & =\frac{f\left(\vec{x}, H_{0}\right)}{f(\vec{x})} \\
& =\frac{f\left(\vec{x} \mid H_{0}\right) \cdot \mathbb{P}\left(H_{0}\right)}{f\left(\vec{x} \mid H_{0}\right) \mathbb{P}\left(H_{0}\right)+f\left(\vec{x} \mid H_{1}\right) \mathbb{P}\left(H_{1}\right)} \\
& =\frac{f\left(\vec{x} \mid \theta_{0}\right)}{f\left(\vec{x} \mid \theta_{0}\right)+f\left(\vec{x} \mid H_{1}\right)} \\
& =\frac{f\left(\vec{x} \mid \theta_{0}\right)}{f\left(\vec{x} \mid \theta_{0}\right)+\int f(\vec{x} \mid \theta) f(\theta) d \theta} \\
& =\frac{\mathcal{L}\left(\theta_{0}\right)}{\left.\mathcal{L} \mid \theta_{0}\right)+\int \mathcal{L}(\theta) f(\theta) d \theta}
\end{aligned}
$$

in our case,

$$
\mathbb{P}\left(H_{0}\right)=\mathbb{P}\left(H_{1}\right)
$$

Notes = prior $f$ can have a lager influence on $H_{0} \mid \vec{x}$

- Improper props an not allowed
- $\mathbb{P}\left(+l_{0} \mid \vec{x}\right)$ is the probability that $H_{0}$ is true gavin $\vec{x} \rightarrow$ this doss not tell us when to rujat the null hypothesis. When do we reject? When do we retain? We need mon detailed analysis.

Read 11.9 in All of stats fir mon steengths/waknesses.
Regression (standard Inion regnssion)
Goal: Fit noisy data using a curve.


Denote: $\quad Y \sim$ random variable, response variable
$X \sim$ covariate, predictor, feature
Regression: $\quad r(x)=\mathbb{E}(Y \mid \vec{X}=\vec{x})=\int y f(y \mid \vec{x}) d y$
Goal: Gavin data $\left(X_{1}, Y_{1}\right) \ldots\left(X_{n}, Y_{n}\right) \sim F_{X, Y}$, estimate $r(x)$.

Basic Linear Regression
Mode: $r(x)=\beta_{0}+\beta_{1} x$ "Simple liver. regression
Ohseru some data: $X_{i}, Y_{i}$
Assumption: $\operatorname{Var}(Y \mid X=x)=\sigma^{2}$

$$
\begin{aligned}
& L Y_{i}=\beta_{0}+\beta, X_{i}+\epsilon_{i} \\
& \hat{E}\left(\epsilon_{i} \mid X_{i}\right)=0 \\
& \operatorname{Var}\left(\epsilon_{i} \mid X_{i}\right)=\sigma^{2} .
\end{aligned}
$$

Given this data for the statistical model, find estimator for the unknown coefficients $\beta_{0}, \beta_{1} \rightarrow \hat{\beta}_{0}, \hat{\beta}_{1}$.

Fitted line : $\hat{r}(x)=\hat{\beta}_{0}+\hat{\beta}_{1} x$
Predicted values: $\quad \hat{Y}_{i}=\hat{r}\left(x_{i}\right)$
Residuals: $\quad \hat{\epsilon}_{i}=Y_{i}-\hat{Y}_{i}=Y_{i}-\left(\hat{\beta}_{0}+\hat{\beta}_{1} x_{i}\right)$
Residual som: $\quad$ PS $=\sum \hat{\epsilon}_{i}^{2} \leftarrow$ only one such of squares metric to determine how all $\hat{r}$ fits the data.


Definition:
Least squares estimates: $\hat{\beta}_{0}, \hat{\beta}$, minimize $R S S=\sum_{i=1}^{n} \hat{\epsilon}_{i}^{2}$.
$\hat{\beta}_{0}, \hat{\beta}_{1}$ can la found using. calculus), livens algebra, statistics, etc.'

Thu:

$$
\begin{array}{ll}
\left.\hat{\beta}_{1}=\frac{\sum_{i}\left(X_{i}-\bar{X}\right)\left(Y_{i}-\bar{Y}\right)}{\sum\left(X_{i}-\bar{X}\right)^{2}}\right\} & \approx \frac{\operatorname{cov}(X, Y)}{\operatorname{var}(X)} \\
\begin{aligned}
\hat{\beta}_{0}=\bar{Y}-\hat{\beta}_{1} \bar{X} & \frac{\sigma_{X} \sigma_{Y} \rho_{X Y}}{\sigma_{X}^{2}} \\
& =\frac{\sigma_{Y}}{\sigma_{X}} P_{X Y}
\end{aligned}
\end{array}
$$

Unbiased estimate of $\sigma^{2}$ :

$$
\hat{\sigma}^{2}=\frac{1}{n-2} \sum \hat{\epsilon}_{i}^{2} .
$$

To find $\hat{\beta}_{0}, \hat{\beta}_{\text {, }}$ using calculus:
solve $\frac{\partial}{\partial \beta_{0}}$ TS $=0$

$$
\frac{\partial}{\partial \beta_{1}} \text { PS }=0
$$

To solve using linear algebra:

compute the orthogonal projection of $Y_{i}$ onto $\operatorname{span}\{\overrightarrow{1}, \vec{x}\}$

Least Squares and Maximum Likelihood Estimator
Add assumption that $\epsilon_{i} \mid x_{i} \sim N\left(0, \sigma^{2}\right)$

$$
\begin{aligned}
& \Rightarrow Y_{i} \mid X_{i} \sim N\left(\mu_{i}, \sigma^{2}\right) \\
& \quad \\
& \rightarrow \mu_{i}=\beta_{0}+\beta_{1} X_{i}
\end{aligned}
$$

Write down the likelihood function:
L. does not depend on any parameters.
$\mathcal{L}_{2}$ is known a) the conditional likelihood and contain all the parameter.

$$
\begin{aligned}
& I_{2}\left(\beta_{0}, \beta_{1}, \sigma^{2}\right) \alpha \frac{1}{\sigma^{n}} e^{-\frac{1}{2 \sigma^{2}}} \sum_{i}\left(Y_{i}-\beta_{0}-\beta_{1} x_{i}\right)^{2} \\
\Rightarrow & \log \mathcal{L}_{2}=l_{2}=-n \log \sigma-\frac{1}{2 \sigma^{8}} \underbrace{\sum\left(Y_{i}-\beta_{0}-\beta_{1} X_{i}\right)^{2}}
\end{aligned}
$$

Thu If $\epsilon_{i} \mid x_{i} \sim N\left(0, \sigma^{2}\right)$, then the MLE for $\beta_{0}, \beta_{1}$ is the same as the least squars estimate, and $\quad \hat{\sigma}_{\text {ALE }}^{2}=\frac{\perp}{n} \sum \hat{\epsilon}_{i}^{2}$
biased estimator.
Properties of then least square estimintos
Thu: Define $\hat{\beta}=\binom{\hat{\beta}_{0}}{\hat{\beta}_{1}}$
then $\mathbb{E}(\hat{\beta} \mid \vec{X})=\binom{\beta_{0}}{\beta_{1}}$

$$
\operatorname{Var}(\hat{\beta} \mid \vec{x})=\frac{\sigma^{2}}{n S_{x}^{2}}\left(\begin{array}{cc}
\frac{1}{n} \sum x_{t}^{2} & -\bar{x} \\
-\bar{x} & 1
\end{array}\right)
$$

when $\quad s_{x}^{2}=\frac{1}{n} \sum\left(x_{i}-\bar{x}\right)^{2}$

The estimated standard eros an $\begin{gathered}\text { syst of } \\ \text { diagonals of } \\ \text { matron })\end{gathered}$

$$
\begin{aligned}
& \hat{\operatorname{se}\left(\hat{\beta}_{0}\right)}=\frac{\hat{\sigma}}{\sqrt{n} s_{x}} \sqrt{\frac{1}{n} \sum x_{i}^{2}} \\
& \hat{\operatorname{se}}\left(\hat{\beta}_{1}\right)=\frac{\hat{\sigma}}{\sqrt{n} s_{x}}
\end{aligned}
$$

Thu These estimation an

- consistent
- asymptotically normal
- and thenfure me un apply the wald test, e, g. $H_{0}: \beta_{1}=0$ vs. $H_{1}: \beta_{1} \neq 0$,

Prediction
Setup: Have the data $X_{1}, Y, \ldots X_{n}, Y_{n}$ and estimate $\hat{r}(x)=\hat{\beta}_{0}+\hat{\beta}_{1} x$ using least squares. Observe (or pict ) a new covariate $X=x_{*}$, and we want to predict $Y_{*}$.
Estimate $\hat{Y}_{x}=\hat{\beta}_{0}+\hat{\beta}_{1} x_{*}$

$$
\begin{aligned}
\operatorname{Var}\left(\hat{Y}_{*}\right) & =\operatorname{Var}\left(\hat{\beta}_{0}+\hat{\beta}_{1} x_{*}\right) \\
& =\operatorname{Var}\left(\hat{\beta}_{0}\right)+x_{*}^{2} \operatorname{Var}\left(\hat{\beta}_{1}\right)+2 x_{*} \operatorname{cov}\left(\hat{\beta}_{0}, \hat{\beta}_{1}\right)
\end{aligned}
$$

and $\hat{\operatorname{se}}\left(\hat{Y}_{X}\right)=\sqrt{\operatorname{Var}\left(Y_{x}\right)}$ using $\hat{\sigma}^{2}$.
How about a $(1-\alpha)$ confidence interval for $Y_{x}$ ?
Nainly, $\hat{Y}_{x} \pm z_{\alpha / 2} \hat{\operatorname{se}}\left(\hat{Y}_{x}\right)$ is a $1-\alpha$ confidime interval, lout this is incornet. (Sue Them 13.11, do exeniñ 10)

The ida behind the mistake:
the above confidence interval is only corot if we never observed the independent noise $\epsilon_{i}$, ie. in the real world re observe $Y_{*}=\beta_{0}+\beta_{1} X_{*}+\epsilon$.

Multiple Regression


$$
\vec{\beta}=\left(\begin{array}{c}
\beta_{1} \\
\vdots \\
\beta_{k}
\end{array}\right) \quad, \quad \vec{\epsilon}=\left(\begin{array}{c}
\epsilon_{1} \\
\vdots \\
\epsilon_{n}
\end{array}\right)
$$

Modi : $\quad \stackrel{\rightharpoonup}{Y}=X \vec{\beta}+\stackrel{\rightharpoonup}{\epsilon}$
See The 13,13 for the least square solution, same as from liver algebra (las).

Logistic Regression
Change the model: Imagine that $Y_{i} \in\{0,1\}$, lie. $\mathbb{P}\left(Y_{i} \mid X_{i}\right)=p_{i}$

We wart to modal $p_{i}$, not $Y_{i}$.

Choose a particular parametric form:

$$
\left.\begin{array}{rl}
p_{i} & =p_{i}\left(\beta_{0}, \ldots, \beta_{k}\right)=\mathbb{P}\left(Y_{i}=1 \mid X_{i}\right) \\
& =\frac{e^{\beta_{0}+\sum_{1}^{h} \beta_{j} x_{i j}}}{1+e^{\beta_{0}+\sum_{i}^{k} \beta_{j} x_{i j}}}
\end{array}\right\}
$$

The logistic function: $\frac{e^{x}}{1+e^{x}}$

$$
\log _{i t}(p)=\log \left(\frac{p}{1-p}\right)
$$


$\operatorname{Sina} \quad Y_{i} \mid X_{i} \sim \operatorname{Bernsu} l_{i}\left(p_{i}\right)$
The conditional likelihood is:

$$
\mathcal{L}(\vec{\beta})=\prod_{i} p_{i}(\vec{\beta})^{Y_{i}}\left(1-p_{i}(\vec{\beta})\right)^{1-Y_{i}}
$$

1 most he maximized numerically.

Multivariate Models
$\underset{\text { Vector }}{\operatorname{Random}} \vec{x}=\left(\begin{array}{c}X_{1} \\ \vdots \\ X_{k}\end{array}\right)$, mean $\mathbb{E}(\vec{x})=\left(\begin{array}{c}\mathbb{E}\left(X_{1}\right) \\ \mathbb{E}\left(X_{L}\right) \\ \vdots \\ \mathbb{E}\left(X_{k}\right)\end{array}\right)=\vec{\mu}$

Covariance matrix

$$
C=\underbrace{\left(\begin{array}{cccc}
\operatorname{cov}\left(x_{1}, x_{1}\right) & \operatorname{cov}\left(x_{1}, x_{2}\right) & \cdots & \operatorname{cov}\left(x_{1}, x_{k}\right) \\
\operatorname{cov}\left(x_{2}, x_{1}\right) & & \vdots \\
\operatorname{cov}\left(x_{3}, x_{1}\right) & & \vdots \\
\vdots & & \\
\operatorname{cov}\left(x_{k}, x_{1}\right) & \operatorname{cov}\left(x_{k}, x_{k}\right)
\end{array}\right)}
$$

$k \times k$ matrix.
$C^{-1}$ is known as the precision matrix,
Ferthermon: $-C$ is symmetric positive definite

- eigenvalue of $C$ ar all positive.

Thu Let $\vec{a} \in \mathbb{R}^{k}$ be a constant vector,
then (i) $\mathbb{E}\left(\vec{a}^{\top} \vec{x}\right)=\vec{a}^{\top} \vec{\mu}$
(2) $\operatorname{Var}\left(\vec{a}^{\top} \vec{x}\right)=\vec{a}^{\top} C \vec{a}$

Let $A \in \mathbb{R}^{n x k}$ constant matrix, then
(1) $\mathbb{E}(A \vec{x})=A \vec{\mu}$
(2) $\left.\operatorname{Var}(A \vec{x})=A \subset A^{\top}\right\}$ another covarisines matrix

Next consider we han samples

$$
\begin{aligned}
& X_{11}, X_{12}, \ldots, X_{i n} \\
& \vdots \\
& X_{k 1}, \ldots . . . X_{k n}
\end{aligned}
$$

The sample mean is then

$$
\bar{x}=\left(\begin{array}{c}
\bar{x}_{1}=\frac{1}{n} \sum_{j} x_{1 j} \\
\vdots \\
\bar{X}_{k m}=\frac{1}{n} \sum_{j} x_{k n}
\end{array}\right)
$$

The sample variance matrix:

$$
\begin{aligned}
& S=\left(\begin{array}{ccc}
s_{11} & s_{12} & \cdots \\
& & \ddots \\
& & \\
s_{u n}
\end{array}\right) \\
& \text { when } \quad s_{i j}=\frac{1}{n-1} \sum_{l=1}^{n}\left(X_{i l}-\bar{X}_{i}\right)\left(X_{j,}-\bar{X}_{j}\right)\left\{\begin{array}{l}
\text { unbiaized } \\
\text { estimate } \\
\text { cove }\left(x_{i}, X_{j}\right)
\end{array}\right. \\
& \mathbb{E}(\bar{X})=\vec{\mu} \quad 1 \quad \mathbb{E}(S)=C .
\end{aligned}
$$

And since the correlation of two variable is


$$
\sqrt{\operatorname{Var}\left(x_{i}\right) \operatorname{Var}\left(x_{j}\right)}
$$

estimator for the correlation is:

$$
\hat{p}_{i j}=\frac{s_{i j}}{s_{i i} s_{j j}}
$$

Example Multivariate Normal
$\vec{X} \in \mathbb{R}^{k} \sim N(\vec{\mu}, C)$ if its density is

$$
\begin{aligned}
& \quad f\left(x_{1} \ldots x_{k} ; \vec{\mu}, c\right)=\frac{1}{(2 \pi)^{k / 2}} \frac{1}{\sqrt{\operatorname{det} c}} e^{-\frac{1}{2}(\vec{x}-\vec{\mu})^{\top} c^{-1}(\vec{x}-\vec{\mu})} \\
& \quad \vec{\mu} \in \mathbb{R}^{k} \\
& C \in \mathbb{R}^{k \times k} \\
& \Rightarrow \\
& \\
& \mathbb{E}(\vec{x})=\vec{\mu} \\
& \operatorname{Var}(\vec{x})=C .
\end{aligned}
$$

Thu Let $\vec{Z} \sim N(\overrightarrow{0}, I)$, and $C$ be a sped matrix.
(1) Let $c^{1 / 2}$ be such that $c^{1 / 2} \cdot c^{1 / 2}=C$., then

$$
\vec{X}=\vec{\mu}+c^{1 / 2} \vec{z} \sim N(\vec{\mu}, c)
$$

(2) $\quad^{-1 / 2}(\vec{x}-\vec{\mu}) \sim N(\overrightarrow{0}, I)$.
(3) $\vec{a}^{\top} \vec{x} \sim N\left(\vec{a}^{\top} \stackrel{\vec{a}^{\top}}{ } C \vec{a}\right)$
(4) $\quad V=(\vec{x}-\mu)^{\top} C^{-1}(\vec{X}-\mu)$, then $V \sim X_{w}^{2}$

$$
=\vec{z}^{\top} \dot{z}
$$

Thu For multivariate dat $\left(\begin{array}{c}x_{11} \\ x_{21} \\ \vdots \\ x_{k 1}\end{array}\right) \ldots\left(\begin{array}{c}x_{i n} \\ \vdots \\ x_{k n}\end{array}\right)$, the log-likalihood

$$
l(\vec{\mu}, C) \propto \frac{-n}{2}(\vec{x}-\vec{\mu})^{\top} C^{-1}(\vec{x}-\mu)-\frac{n}{2} \operatorname{tr}\left(C^{-1} S\right)-\frac{n}{2} \log d t C
$$

Recall that $\operatorname{tr}(A)=\sum A_{i i}$,
$S=$ as befon, the sample covariminu matrix.
The MLE's arc $\hat{\mu}=\bar{x}, \quad \hat{c}=\frac{n-1}{n} S$

Gaussian Processes
Simplest interpretation: the extension of random vector to random functions

Best single source refernce:
Rasmussen \& Williams, Gaussian Processes fur Machine Learning

We say that $f \sim G P(m, k)$ is a Gaussian process with mean $m$ and covariance function $k$.

$$
\begin{align*}
& \Rightarrow \quad \mathbb{E}(f(x))=m(x) \\
& \operatorname{Cov}\left(f(x), f\left(x^{\prime}\right)\right)=k\left(x, x^{\prime}\right) \\
& \Rightarrow \quad \mathbb{F}\left(\begin{array}{c}
f\left(x_{1}\right) \\
\vdots \\
f\left(x_{n}\right)
\end{array}\right)=\left(\begin{array}{c}
m\left(x_{1}\right) \\
m\left(x_{2}\right) \\
\vdots \\
m\left(x_{n}\right)
\end{array}\right) \\
& \operatorname{Cov}\left(\begin{array}{c}
f\left(x_{1}\right) \\
\vdots \\
f\left(x_{n}\right)
\end{array}\right)=1<\text { with entries } K_{L}\left(x_{i}, x_{j}\right)  \tag{12}\\
& \Leftrightarrow \quad k\left(x, x^{\prime}\right)=\mathbb{E}\left((f(x)-m(x))\left(f\left(x^{\prime}\right)-m\left(x^{\prime}\right)\right)\right)
\end{align*}
$$

Example covariance functions:

$$
k\left(x, x^{\prime}\right)=A e^{-\left(x-x^{\prime}\right) / b}
$$

$\left.k\left(x, x^{\prime}\right)=B e^{-\left|x-x^{\prime}\right| / c}\right\}$ one function from the Maters family of covarinie kernels.

Graphically

$$
k\left(x, x^{\prime}\right)=e^{-\left(x-x^{\prime}\right)^{2} / .0001}
$$



$h\left(x, x^{\prime}\right)=\delta\left(x-x^{\prime}\right)=1$ if $x=x^{\prime}$ 0 otherwise



Brownian $\quad B(t)=\int_{0}^{t} w(\tau) d \tau$
$\hat{L}_{G P} \operatorname{Cotion}(0, \delta)$

