

L'brahur: Worsermann: all of statistics

Shoudard setup in parametric shatistics

We assume that do to is quivaled by a particular
family of distributions, for example

$$\mathcal{F} = \left(N(\mu, \sigma^{e}) \right) \mu \in \mathbb{R}, \sigma^{2} > 0 \right\}$$
.
The family \mathcal{F} is called the stabistical model.
Hore quivally, $\mathcal{F} = \left\{ f \ominus \mid \Theta \in \Theta \right\}$
 $\int_{\mathcal{F}} \int_{\mathcal{F}} \int_{\mathcal{F}}$

Guvenhour

Pavamehr space
$$\Theta$$
 ("capital Hieta")
Fare (unlinown) pavamehr Θ ("bour case Heba")
 P_{Θ} , E_{Θ} ... refere to the probability, expectation
under the distribution f_{Θ}
Estimates hypically get a "hot": $\hat{\Theta}$, $\hat{\mu}$, ...

Point estimation

$$\frac{\Phi_{4}}{\Theta_{n}} = \frac{1}{2} \left\{ \begin{array}{c} \frac{\Phi_{4}}{\Theta_{n}} \\ \frac{\Phi_{4}}{\Theta_{$$

Variance and standard error

$$\mathcal{D}_{ef}$$
 the voriance of an estimator it defined as
 Var_{Θ} ($\hat{\Theta}_{n}$). The corresponding standard deviation
is called the standard error se. Typically, pe
is called the standard error se. Typically, pe
is unknown, but it can be estimated : \hat{se} .

Example
$$[o_{1}1]^{2}$$

 $X_{1},...,X_{n} \sim \text{Perudull}(\rho), \text{ porametry } \rho \in [o_{1}1],$
 $\hat{p}_{n} := \frac{1}{n} \sum_{i=1}^{n} X_{i}$ an orthinate of p .
 $E_{p}(\hat{p}_{n}) = E_{p}(\frac{1}{n} \sum_{i=1}^{n} X_{i}) = \frac{1}{n} \sum_{i=1}^{n} E_{p}(X_{i}) = p.$
 Mur, \hat{p}_{n} is unbiand because
 $E_{p}(\hat{p}_{n}) - p = p - p = 0.$

The standard error of his estimate is

$$se = \sqrt{\operatorname{Var}_{p}} \left(\hat{p}_{n} \right) = \sqrt{\frac{\Lambda}{n}} \operatorname{Var}_{p} \left(X_{n} \right) = \sqrt{\frac{p(\Lambda - p)}{n}}$$

We can for example estimate if by
$$\int_{se}^{\Lambda} = \sqrt{\frac{\hat{p}_{n}(\Lambda - \hat{p}_{n})}{n}}$$

Example: weight of Laby





Mean squared estor

$$\frac{\partial f}{\partial t} \quad \text{The mean squared error (HSE) of an estimate of the quantity
He quantity
$$\frac{\mu SE(\hat{\theta}, \theta)}{\theta} = E_{\theta} \left(\left(\hat{\theta}_{n} - \theta \right)^{2} \right) \\ \frac{\mu SE(\hat{\theta}, \theta)}{\theta} = E_{\theta} \left(\left(\hat{\theta}_{n} - \theta \right)^{2} \right) \\ \frac{\mu SE(\hat{\theta}, \theta)}{\theta} = E_{\theta} \left(\left(\hat{\theta}_{n} - \theta \right)^{2} \right) \\ \frac{\mu SE(\hat{\theta}, \theta)}{\theta} = E_{\theta} \left(\left(\hat{\theta}_{n} - \theta \right)^{2} \right) \\ \frac{\mu SE(\hat{\theta}, \theta)}{\theta} = E_{\theta} \left(\left(\hat{\theta}_{n} - \theta \right)^{2} \right) \\ \frac{\mu SE(\hat{\theta}, \theta)}{\theta} = E_{\theta} \left(\left(\hat{\theta}_{n} - \theta \right)^{2} \right) \\ \frac{\mu SE(\hat{\theta}, \theta)}{\theta} = E_{\theta} \left(\left(\hat{\theta}_{n} - \theta \right)^{2} \right) \\ \frac{\mu SE(\hat{\theta}, \theta)}{\theta} = E_{\theta} \left(\left(\hat{\theta}_{n} - \theta \right)^{2} \right) \\ \frac{\mu SE(\hat{\theta}, \theta)}{\theta} = E_{\theta} \left(\left(\hat{\theta}_{n} - \theta \right)^{2} \right) \\ \frac{\mu SE(\hat{\theta}, \theta)}{\theta} = E_{\theta} \left(\left(\hat{\theta}_{n} - \theta \right)^{2} \right) \\ \frac{\mu SE(\hat{\theta}, \theta)}{\theta} = E_{\theta} \left(\left(\hat{\theta}_{n} - \theta \right)^{2} \right) \\ \frac{\mu SE(\hat{\theta}, \theta)}{\theta} = E_{\theta} \left(\left(\hat{\theta}_{n} - \theta \right)^{2} \right) \\ \frac{\mu SE(\hat{\theta}, \theta)}{\theta} = E_{\theta} \left(\left(\hat{\theta}_{n} - \theta \right)^{2} \right) \\ \frac{\mu SE(\hat{\theta}, \theta)}{\theta} = E_{\theta} \left(\left(\hat{\theta}_{n} - \theta \right)^{2} \right) \\ \frac{\mu SE(\hat{\theta}, \theta)}{\theta} = E_{\theta} \left(\left(\hat{\theta}_{n} - \theta \right)^{2} \right)$$$$

Mearen: bias-vaniance - de composition

$$MSE(\hat{\theta}_{n}, \theta) = bias^{2}(\hat{\theta}_{n}) + Var_{\theta}(\hat{\theta}_{n})$$
how good is our
cotimete

$$= \underbrace{E}_{\Theta} \left(\left(\widehat{\theta}_{u} - \widehat{E} \, \widehat{\theta}_{u} \right) \right) + \underbrace{K} \left(\left(\underbrace{E}_{\theta} \, \widehat{\theta}_{u} - \Theta \right)^{2} \right) \\ deterministic} \\ = \left(\underbrace{E}_{\theta} \, \widehat{\theta}_{u} \right)^{2} \\ = \left(\underbrace{E}_{\theta} \, \widehat{\theta}_{u} \right)^{2}$$

 \sim

M

Example

$$\begin{aligned} \mathcal{F} &= \left\{ \begin{array}{l} \mathsf{N}(p_{1} \sigma^{2}) \mid p \in \mathbb{R}, \ 6 > 0 \right\} \\ \text{Sample} &: \mathsf{K}_{1} \cdots \mathsf{r} \mathsf{K}_{n} \sim \mathsf{N}(p_{1} \sigma^{2}) \quad \text{with unknowly } p_{2} \sigma^{2} \quad \text{ind} \\ \hat{\mu} &= \begin{array}{l} \frac{n}{n} \sum_{i=n}^{n} \mathsf{K}_{i} \quad \text{is an unbrand which of } \mu. \\ \hat{\mu} &= \begin{array}{l} \frac{n}{n} \sum_{i=n}^{n} \mathsf{K}_{i} \quad \text{is an unbrand which of } \mu. \\ \hat{\sigma}_{n}^{2} &:= \begin{array}{l} \frac{n}{n} \sum_{i=n}^{n} \left(\mathsf{K}_{i} - \hat{\mu}\right)^{2} \quad \text{first when}^{n} \\ \hat{\sigma}_{2}^{2} &:= \begin{array}{l} \frac{n}{n} \sum_{i=n}^{n} \left(\mathsf{K}_{i} - \hat{\mu}\right)^{2} \quad \text{first when}^{n} \\ n = n = \begin{array}{l} \frac{n}{n-n} \sum_{i=1}^{n} \left(\mathsf{K}_{i} - \hat{\mu}\right)^{2} \quad \text{first when}^{n} \\ \end{array} \end{aligned}$$

$$E(\hat{\sigma}_{1}^{2}) = \frac{n-\lambda}{n} \sigma^{2} \quad \text{so } \text{ He biar is } \frac{1}{n} \sigma^{2}$$

$$E(\hat{\sigma}_{2}^{2}) = \sigma^{2} \quad \text{unbiand!}$$

$$Var(\hat{\sigma}_{1}^{2}) = \frac{2(u-\lambda)\sigma^{4}}{n^{2}}$$

$$Var(\hat{\sigma}_{2}^{2}) = \frac{2\sigma^{4}}{n-\lambda}$$

$$HSE(\hat{\sigma}_{1}^{2}) = bis^{2} + var = \dots = \left(\frac{2u-\lambda}{u^{2}}\right) \sigma^{4}$$

$$HSE(\hat{\sigma}_{2}^{2}) = \dots = \frac{2}{u-\lambda} \sigma^{4}$$

$$u = \frac{2}{u-\lambda} \sigma^{4}$$

=)
$$\mu SE(\hat{e}_1) < \mu SE(\hat{e}_2)$$

Consistant estimator

$$\frac{Def}{Def} \qquad A \quad psiut estimator \quad \widehat{\Theta}_n \quad of \quad \Theta \quad ir \quad coustituent} \\ (rtrougly courrictient) \quad if \\ \widehat{\Theta}_n \quad - > \quad \Theta \quad in \quad probability \quad (a.r.) \\ ar \quad u \rightarrow \infty \\ \\ Meanin \qquad If an istimuch so this firs \quad bias \rightarrow 0 \quad and \quad re \quad - > 0 \\ ar \quad u \rightarrow \infty, \quad the \quad the istimuch \quad ir \quad coustistent. \end{cases}$$

Gufiduce rets

$$\frac{\partial ef}{\partial e R} = A (1-\alpha) - Confidence inhoral for a parameter
\Theta \in R is an interval $C_{11} = (\alpha_{11}, \alpha_{11})$ where
 $\alpha_{11} = \alpha(K_{11}, \dots, K_{11})$, $b_{11} = b(K_{11}, \dots, K_{11})$ or functions
of the source K_{11}, \dots, K_{11} such that

$$P_{\Theta} (\Theta \in C_{11}) \ge 1 - \alpha$$
 for all $\Theta \in \Theta$.
Where
(underval)
prometer
 $(1-\alpha)$ is called the coverage of the confidence interval.$$

lurhahion



Example

Coin flips, with
$$P(X = \Lambda) = \rho$$
, $P(X = 0) = \Lambda - \rho$,
 $\rho \in [\partial_{1}\Lambda]$ unknown. Want to estimation it.
 ~ 1 Observe $X_{1}, \dots, X_{n} \sim f\rho$

٠

$$\mathcal{E}_n^2 := \frac{\log(2/\alpha)}{2n}$$

Proposition:
$$Cn := \left(\hat{p}_n - \mathcal{E}_n, \hat{p}_n + \mathcal{E}_n \right)$$
 is a CI with cover \mathcal{A} .

Proof (example)

Proof: By Hoeffding mequality, for any two han $P(|p_n - p| > t) \leq 2 \exp(-2ut^2)$ Set $\alpha := 2 \exp(-2 \pi t^2)$ $r_{n-E} p \hat{p} n \hat{p}_{n+E}$ and rolve for t: $\log\left(\frac{\alpha}{2}\right) = -2ut^2 = t^2 = -\frac{\log(\frac{\alpha}{2})}{2n} = \frac{\log(2/\alpha)}{2n}$ Chor Enst.



Lihelihood

More forwally: Parametric family
$$\mathcal{F} = \{f_{\theta} \mid \theta \in \Theta\},\$$

observe idd points $X_{A_{1}}, \dots, X_{A_{n}} \sim f_{\theta} \in \mathcal{F}.$
The likelihood of the data given a parameter Θ_{0} is
 $P_{\Theta_{0}}(X_{A_{1}},\dots,X_{n}) = P(X_{A_{1}},\dots,X_{n} \mid \Theta_{0})$
 $= \prod_{i=A}^{n} P(X_{i} \mid \Theta_{0})$ ustation!

Maximum likelikord

To ashimate the true parameter θ , we now select θ such that this likelihood is maximized:

$$\hat{\Theta} := \operatorname{argmax} P(X_{1,\dots,} X_{1} | \Theta) = \operatorname{argmax} \prod_{i=1}^{n} P(X_{i} | \Theta)$$

 $\Theta \in \Theta$

θ

$$\hat{\Theta} = \operatorname{argmin} \left\{ \operatorname{lsg} \left(\operatorname{tr} P(X_{c} \mid \theta) \right) = \operatorname{argmax} \left[\operatorname{lsg} \left(\operatorname{lsg} P(X_{c} \mid \theta) \right) \right] = \operatorname{argmax} \left[\operatorname{lsg} P(X_{c} \mid \theta) \right] = \operatorname{argmin} \left[\operatorname{lsg} P(X_{c} \mid \theta) \right] = \operatorname$$

20

Example for an analytic polution

Model:
$$X \sim Poirrow(A)$$
, their means that

$$P(X=x) = \frac{\lambda^{x} e^{-\lambda}}{x!}, \quad i' + her E(X) = \lambda$$

$$Var(X) = \lambda.$$
Observe $X_{1,..., X_{u}} \sim Poirrow(A)$

$$W and to construct the HL - estimator. for $\lambda.$

$$Goupule the likelihood:$$

$$Y(\lambda) = P(X_{1,..., X_{u}} \mid \lambda) = \prod_{i < A} \frac{\lambda^{x_{i}} e^{-\lambda}}{x_{i}!}$$$$

Example (continued)

$$(o_{3}(\dots) = \sum_{i=n}^{n} (\log \left(\frac{x^{i} e^{-\lambda}}{x_{c}!}\right)$$

$$= \sum_{i=n}^{n} (X_{i} \log \lambda - \lambda - \log (x_{i}!))$$

$$= \sum_{i=n}^{n} (X_{i} \log \lambda - \lambda - \log (x_{i}!))$$

Now would to optimize for
$$\lambda$$
. Take the absorbed (work λ):

$$f'(\lambda) = \sum_{i=n}^{n} \left(\frac{x_i}{\lambda} - 1\right) = \frac{1}{\lambda} \left(\sum_{i=1}^{n} x_i\right) - n \stackrel{!}{=} 0$$

$$\Rightarrow \lambda = \frac{1}{n} \sum_{i=n}^{n} k_i$$
So $\hat{\lambda} := \frac{1}{n} \sum_{i=1}^{n} x_i$ is the the cohinete of λ .

MLE properties

$$\frac{\Theta_{\text{HLE}}}{Se} = \Theta_{\text{indivitr.}}$$

 $N(0, 1)$ and

$$\frac{\Theta_{\rm MCE}}{Se} = \Theta_{\rm inductr.}$$

 $N(O_{\rm I} \Lambda)$

(s) This can be used to construct car fidure introde:

$$C_{n} := \left(\begin{array}{c} \hat{\theta}_{HLE} - \frac{3}{2}x_{12} & \hat{r}e \\ -E \end{array}\right) \quad \hat{\theta}_{HLE} + \frac{3}{2}x_{12} & \hat{r}e \\ \frac{1}{2}e^{-E} & \frac{1}{2}e^{-E} \\ \frac{$$



Sufficiency

Sufficiency

Which properties would we used to arost sufficiency?

when we observe two soundes X₁,..., X_n and X₁',..., X_n',
and T(X₁,..., X_n) = T(X_n',..., X'), then we would infer the same Θ.
When we have T(K₁,..., X_n), then we would used some way to estimate Θ just backdon Θ.

Formal definition is kohning, shipped.

(den hifig Sility

Sometimes families of distribution can be described in depent
ways with differt sets of parameter.
Dest A parameter & for a family
$$\overline{F} = \{f_{\Theta} \mid \Theta \in \Theta\}$$
 is
intentificable if distrived values of Θ correspond to district
pdfr in \overline{F} :
 $G \neq G' = \{f_{\Theta} \neq f_{\Theta}\}$

(Identifiability is a projecty of the model iralf, not of the data)

Kypomis teoling
Hohivahoy



Gewool idea





But how do we know what "for yout" is?

Example Want to lest whether a coin it fair.
Unit hypothesis: Ho: coin it fair
Altoretive hypothesis: Ho: coin it unfair
Somple many coin flips and estimate
$$\hat{p}_n = \int_{1}^{\infty} \sum_{i=1}^{\infty} t_i$$
.
We want to kject Ho if \hat{p}_n is "for away" from O.J.
Question: "for away"?
Losh at the division of \hat{p} under the unit hypothesis:
 $f_n = riject$
retain $\hat{p}_n = 0.5$ $\hat{p}_n = riject$
 $retain \hat{p}_n = 0.5$ $\hat{p}_n = riject$

More formal setup

Statistical model
$$\mathcal{F} = \{f_{\Theta} \mid \Theta \in \Theta\}$$
. Assume that
 $\Theta_{O} \subset \Theta_{1} \quad \Theta_{1} \subset \Theta_{1} \quad \Theta_{O} \cap \Theta_{1} = \mathcal{J} \quad \Theta_{O} \cup \Theta_{1} = \widehat{\mathcal{J}}$
Want to test
 $\underbrace{H_{O}: \Theta \in \Theta_{O}}_{\text{null hey}}$ against $\underbrace{H_{1}: \Theta \in \Theta_{1}}_{\text{null hey}}$.
Sample data from the unknown f_{Θ} , compatin a test obtatistic
 $T(X_{1}, ..., X_{n})$. Now we construct a rejection region R_{n}
such that $T(X_{1}, ..., X_{n}) \in R_{n} = 3$ reject the
 $T(X_{1}, ..., X_{n}) \notin R_{n} = 3$ reject the

Typical hypotheses are of the form
•
$$H_0: \Theta = \theta_0$$
 vs $H_1: \Theta \neq \Theta_0$
• $H_0: \Theta \leq \Theta_0$ vs $H_1: \Theta \geq \Theta_0$

Level of a test

Def live say that a test is of level
$$\alpha$$
 if
 $\sup \beta(\theta) \leq \alpha$
 $\theta \in \Theta_{\theta}$

Pour of a kst
We always resilve the pour of a list against a fixed albruchive
parameter
$$\Theta_A \in \Theta_A$$
.
The pour of hot against albruchive Θ_A is pinn as
 $1 - \beta(\Theta_A)$

Standard approach for traking

Uniformly most poweful kst Let I be a set of tests of level & for testing Yef. $H_0: \Theta \in \Theta_0$ vs $H_1: \Theta \notin \Theta_0$. A fest in I will pour function B(O) is uniformly most poareful (UMP) if $\beta(\Theta) \ge \beta'(\Theta)$ for all $\Theta \in \Theta^{C}$ and for all p' that are power functions for other lests in J.

Kunsk: la practice it is often impossible to find an UMP test.

Example: runne or winh? Assume we want to jest whether it currently is summer as Say, we hypically believe it is summer melers wink. evidence ir apairit it : Hg = Jrunne g Hy = {winter }

Now we construct reveal potr band on the away temperatur of the current day (1 measurement)

Example: runne or winh?



Example: runne or winte?



Binomial chample
We throw 5 Dains with
$$P({}^{*}\Lambda^{*}) = \theta \in E_{0}$$
 fJ.
Wout to hat $H_{0}: \Theta \stackrel{<}{=} \frac{\Lambda}{2}$ vs $H_{1}: \Theta > \frac{\Lambda}{2}$
Test 1: reject the if we observe 5 times 1: $\Lambda \wedge \Lambda \wedge \Lambda$
Power function: $[I(\Theta) = P_{\Theta}(reject) = \Theta^{5}$
Test 2: reject Ho if we observe at least 3 times 1: $0 \ge \Lambda n \wedge r$
 $\theta \land \theta \land \eta \land \sigma \sim \dots$
Power fet : $[I(\Theta) = P_{\Theta}(T_{1}H, 5 \text{ times } 1) =$
 $= \left(\frac{5}{3}\right) \stackrel{?}{\Theta}(\Lambda - \Theta)^{2} \in \left(\frac{5}{4}\right) \Theta^{4}(\Lambda - \Theta) + \Theta^{5}$







Binomial example Pours functions of both tests



Binomial crample



z - Test (Wald hut)

If the distribution of the last statistic is asymptotically normal: $\frac{\delta}{\delta} - \theta \longrightarrow \mathcal{W}(0, \Lambda)$ $(t_0: \Theta = \Theta_0, H_1: \Theta \neq \Theta_0)$ Test: reject when $\left|\frac{\hat{\theta}-\theta_0}{\hat{\rho}}\right| > \frac{2\alpha/2}{\omega}$ where

z - Test (Wald hit)



z - Test (Wald hat)

Pour of test increases with Effect size 16-201:



Theorem Suppose we let
$$H_0: \theta = \theta_0$$
 against $H_1: \theta = \theta_1$.
Consider
 $T = \frac{\mathcal{I}(\theta_n)}{\mathcal{I}(\theta_0)} = \frac{\prod_{i=1}^n f(x_i \mid \theta_i)}{\prod_{i \leq n} f(x_i \mid \theta_0)} \int_{i=1}^n \frac{1}{n t_i} \frac{1}{n$

$$\frac{\text{Marc pureal likelihood-ratio-kst:}}{\text{Porometr space } \Theta_{1} \quad \Theta_{0} \subset \Theta_{1} \quad \Theta_{1} = \Theta_{0}^{C}. \text{ Then we consider the test statistic } \\ \frac{\mathcal{N}}{T} = \frac{\sup_{\Theta \in \Theta_{0}}}{\operatorname{\Theta \in \Theta_{1}}} \quad \text{or even simpler } T = \frac{\sup_{\Theta \in \Theta_{0}} \mathcal{X}(O)}{\sup_{\Theta \in \Theta_{1}} \mathcal{X}(\Theta)}$$

and we detuning a parameter & such that the rejection repions
is of the form
$$R = \{T \leq J \}$$
.



Courider a test at level
$$d_{1}$$
 and densk its
rejection region as R_{α} .
Recall: $k = P(T_{\gamma p}e - E - error)$.
The smalles d_{1} the word ifficult does it get to reject the
(are of the even han that $k < \hat{\alpha} = \sum R_{\alpha} \subset R_{\tilde{\alpha}}$)



Def the p-value is defined as

$$p = \inf \{ \mathcal{A} \mid T(x_{1,\dots,x_{n}}) \in \mathbb{R}_{d} \}$$

 $i_{e} \in Me smallest & for which the level-e-but would
reject the null hypothesis.
Cutuition: Anally p-values are below a more evidence for
rejecting the null$



for a large lest will find a statistically significant difference. As small p

Tratifically similicant Guardine: $p \leq 0.05 \rightarrow "rignificant"$ $p \leq 0.01 \rightarrow "highly significant"$

"We can have small y-values, yet the "effect rize" can be hing.

Theorem Suppose we let
$$H_0: \theta = \theta_0$$
 against $H_1: \theta = \theta_1$.
Consider
 $T = \frac{\mathcal{I}(\theta_1)}{\mathcal{I}(\theta_0)} = \frac{\prod_{i=1}^n f(x_i \mid \theta_1)}{\prod_{i \leq n} f(x_i \mid \theta_2)} \int_{i=1}^n \frac{1}{n t_i} \frac{1}{n$
Marc ground likelihood-ratio-kst:
Narc ground likelihood-ratio-kst:
Norametr grace
$$\Theta_1 \quad \Theta_0 \subset \Theta_1 \quad \Theta_1 = \Theta_0^C$$
. Then we
consider the test statistic
 $\sum_{\substack{n=1\\n \in \Theta_1}} \sum_{\substack{n \in \Theta_1\\ \theta \in \Theta_0}} \sum_{\substack{n \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_0}} \sum_{\substack{n \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_0}} \sum_{\substack{n \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_0}} \sum_{\substack{n \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_0}} \sum_{\substack{n \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_0}} \sum_{\substack{n \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_0}} \sum_{\substack{n \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_0}} \sum_{\substack{n \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_1}} \sum_{\substack{n \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_1}} \sum_{\substack{n \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_1}} \sum_{\substack{n \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_1}} \sum_{\substack{n \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_1}} \sum_{\substack{n \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_1\\ \theta \in \Theta_1}} \sum_{\substack{n \in \Theta_1\\ \theta \in \Theta_1$

and we detuning a parameter & such that the rejection region
is of the form
$$R = \{T \leq A\}$$
.



Courider a test at level
$$d_{1}$$
 and densk its
rejection region as R_{α} .
Recall: $k = P(T_{\gamma p}e - E - error)$.
The smaller d_{1} the word ifficult does it get to reject the
(are of the even han that $k < \hat{\alpha} = \sum R_{\alpha} \subset R_{\tilde{\alpha}}$)





for a large lest will find a statistically significant difference. As small p

Multiple tisting



Asrume we van, for each june, a lest of level & P(Fest i mohen type-E-error) = 5%. Now we have in tests.

$$P\left(at \text{ least one of the hate welles a $\frac{1}{7}/(-12 - 2000)\right) =$

$$= P\left(\frac{1}{14} \text{ makes error at } \frac{1}{52} \text{ error or } \frac{1}{52} \text{ or } \frac{1}{50} \text{ melles error}\right)$$

$$= 1 - P\left(\text{no error in tr ound us error in } \frac{1}{52} \text{ and } \frac{1}{50}\right) = \frac{1}{120} + \frac{1}{120} P\left(\frac{1}{100} \text{ error in } \frac{1}{120}\right) = \frac{1}{120} + \frac{1}{120} \left(\frac{1}{500}\right)^{100} + \frac{1}{120} + \frac{1}{120}$$$$

Boukroui correction

Astrume we run in tests, and we want to a derive
The FWER of
$$(e.g. d = \partial.05)$$
. Then we run
the individual lests with less $\frac{d}{m} = : \frac{d}{ringle}$. Then:

FINER =
$$P(a + least one type - I - error) =$$

= $P(t_1 = rror error) = P(t_1 = rror) = m \cdot e_{single} = e_{single} =$

•

Boufroni, direursion

$$\frac{Def}{E} = \left(\begin{array}{c} \frac{H}{F} \text{ fabre njeching} \\ F \text{ all rejections} \end{array}\right) = : FDR$$

$$He \text{ fabre discovery rate.}$$

· Fix FDR & in advance.

· Row the mindividual tests and evaluate their p-values.

Find the largest index is
such that
$$P(i_0) \leq L_{i_0}$$
.
(below the red line)

•



Theorem : If the Renjamini - thelibery procedure is applied
(and the testr or independent), then repord less of
how many will hypotheses are true and toponol less of
the distribution of y-values when the will is false,
we obtain FDR
$$\leq \alpha$$
.

Intuition





General Remarks

- . BH fends to han mar your than Touferoui
- · BH could FDR, ust FWGR (swell type [-error)!
- · Blt worker best in spore regime where only few tests reject the null
- · Blt gives gaarantees on FDR, but in queval alses ust minimize it.

Non-parametric tests

Standard (parametic scenario): . Statistical model $F = \{f_{\Theta} \mid \Theta \in \Theta\}$ distibution of the samples · Observe data, compute a list d'atities, for example the mean X · Need to know the distribution of the fest statistics T under pre null distribution: distibution of Tn mole he wall ligg. reject reject

We counido her colf
We counido her colf

$$F_0 = colf$$
 of the pinn distribution
 $F_u = colf$ of the platon
 $D_u := \sup_{x \in \mathbb{R}} \left[F_u(x) - F_v(x) \right]$
By the alliventes- Cautelli hereren we hnow that under
Her will hypothesis, $F_u \rightarrow F_v$ uniformity, a.s.



two saught hat

Two sample kert:
$$X_{n_1} \cdots X_n \sim F_n$$
 a first sample
distributed accoroling to F_{n_1}
 $Y_{n_1} \cdots Y_m \sim F_2$ a second sample distributed acc. to F_2
Question: $F_n = F_2^{n_1}$

 $H_0: F_1 = F_2 \qquad H_1: F_1 \neq F_2$

Wilcoxon - Hauy - Whitney test (based ou rankp) · Pool Hie saugle : X1,..., X4, Y1, ..., Ym GR Test: · Sort the pooled sample in increasing order and retrier pre rouch of all points ~ raule (x;) rauh (4;)





· Compute the rach sums for both groups:



. Post the sople

· Coupule the diffuce
$$T = mean (red) - moan (blue)$$



· Check whether the observed Tobserved on the true date is 5 t.

Bootrhrep tests

Kohvahpy

Mohivahioui
$$X_{1},...,X_{n} \sim F$$
, no hnowledge on F
want to estimate a parameter $\theta = t(F)$. You
purvate an estimate $\hat{\theta}$ based on $X_{1}...,X_{n}$, would
be known how releadle $\hat{\theta}$ is.



Algorithm in preudo code
under of original sample paints
In put:
$$x_{1},..., x_{n}$$
 number of bootohrap replications
For $b = A_{1},..., B$
 \cdot Sample $x_{1}^{*},..., x_{n}^{*}$ uniformly with replace ment
from $X_{1},..., x_{n}$
 \cdot Ortimate the parameter $\hat{\Theta}_{b}^{*}$
Gitimate the standard error $\hat{\Gamma}_{c}$ of the original estimate $\hat{\Theta}_{c}^{i}$
by the standard dev. of the bootohry replicates:
 $\hat{\Gamma}_{c} = \left(\begin{array}{c} A \\ B-A \\ B-A \end{array}\right)^{i} \left(\begin{array}{c} \hat{\Theta}_{b}^{*} - \left(\begin{array}{c} A \\ B \\ B \end{array}\right)^{2} \\ mean of replicates\end{array}\right)^{2}$

Dow it dways work?

Confirming mult for boatstrep
Theorem (Consistency of the shimate of the standard arror)
Assume that
$$x_{1,...,1} x_{n-1} \in T$$
, iid, and
 $E(\|X_{1}\|^{2}) \leq \infty$.
Let $\hat{G}_{n} = g(X_{1,...,1} x_{n})$ be the parameter that we estimate.
Assume that g is continuously differentiable in a
neighborhood of $p = E x_{1}$, with a non-zero gradient.
Then the boatstrap estimate of the standard error is
strongly consistent.
Want to estimate Θ . The UL astimate of Θ is simply the largest number are observe:

Confidure sets les Lootohop



$$CE = [a, b]$$

$$It has coverage 1 - \alpha be cank (appleinstely, be cause a, 7 finite)$$

$$p_{0}(\hat{G} \in CE) = 1 - \alpha \qquad be cause a, 7 finite)$$



Frequentist vs. Payesian statistics

Frequentist of histors:

- . Prosobility = limiting frequency
- · parametus & are courtants, we cannot assign probabilities to know
- · statistics believes well when agreeded sylhin

Bayesian Adhidic probability = depree of belief porameters do have probabilities have a prior belief about the world, update it band ou observed data.

Jaquian tatisfies: me model

Bayesian opprovels: pris distribution





Statistics derived pour joshoios

• You can continue confidence in houses:
find a, b such that
$$P(\Theta \in [\alpha, b]) = 95\%$$
.

Discussion

le'bratur:

Kigh-dimensional statistics is different!

Norms of high-dim vectors are ouccubated
Couridu
$$X_{1,\dots,X_{d}}$$
 independent, $E(x_{i})=0$, $Var(X_{i})=1$, $X = \begin{pmatrix} x_{i} \\ \vdots \\ x_{d} \end{pmatrix}$

Expected norm of X:

$$E(\|X\|^2) = E(\sum_{i=1}^{d} X_i^2) = \sum_{i=1}^{d} EX_i^2 = d$$

Concentration of the worm:

By Boustin inequality are can prove:

$$P(||X|| - |d| | > t) \leq 2 \exp(-ct^2)$$

Norms of high-dim vectors are ouccubated Junhahion:







2.5

3.0

10.0 -Hedneuck

2.5

0.0 + 0.0

0.5

1.0

1.5

Points drawn from N(0, I/d), n=100, d=100

2.0

Volume of balls is concentrated along "surface" Concentration can be seen from the point of view of sampling (previous stide), het also pour the point of view of geometry:



Uslume of unit balls gets hing
the Euclidean volume of the unit ball in Rd
goes hoor
$$d \rightarrow \infty$$
.
 $1^{15} \int_{0}^{10} \int_{0$

Figur cadit: Norius Wilf, Midiu 2024

Random vectors on the ophere are nearly orthogonal

$$X = \begin{pmatrix} x_n \\ \dot{x}_d \end{pmatrix}, \quad Y = \begin{pmatrix} Y_n \\ \dot{y}_d \end{pmatrix} \text{ drawn uniformly from sphere}$$

$$E\left(\left| \langle x, y \rangle\right|\right) = 1/d \quad (\approx 0.14 \text{ dir large!})$$



~> pet gete prollematic, ree later

Landous vectors on the sphere are nearly orthogonal Simulations: points drawn from d-dim normal distr. $N(\partial_1, \frac{1}{d}, Jd)$ in increasing dimensions (d=2,3, 10, 100, 1000)







Let's book at distances in high-dim spaces If we are lucky ... pairwise distances ar still meaning full ~ Johnson-Lindustranss

Can work "or if we were in a low dimention";

Messeur of Johnson- Lindenstromps:
Consider a pointe in IR^d, fix EDD. One can find a linear map
$$f: \mathbb{R}^d \to \mathbb{R}^k$$
 with $k \approx \log n/c^2$ such that
the distances do not change by "more than E:

$$(1-\varepsilon) || x_i - x_j || \le \| f(x_i) - f(x_j) \| \le (1+\varepsilon) \| x_i - x_j \|$$

But of his is one not lucky: randous projections Look like (n'uple) Gauppions X ~ Unif (Sd-1 Vd'), Project X on a fixed vector V: Then <x, v > N(D, A) in distribution in prochice, this offen also holds for non-uniform data.

We dou't see any class or clush muchur!



22 between-class-distances ?? Within- class - distances

Histor of Gaurnious with different variances:

$$x_{i}x' \sim \mathcal{V}(\mu_{i} \mathfrak{s}^{*}_{i} \mathfrak{s}_{i}), \quad \forall_{i} \forall' \sim \mathcal{N}(\mu_{2}, \mathfrak{s}^{*}_{2} \mathfrak{s}_{d})$$

Within-class distance:
 $E(\mathcal{U}x - x'\mathcal{V}) = 2d\mathfrak{s}^{2}_{1} \quad (c(arr 1))$
 $2d\mathfrak{s}^{2}_{1} \quad (c(arr 2))$

Between-class-distances:

$$E(\pi x - 4\pi^2) = |\mu_{\mu_1} - \mu_2||^2 + d(6^2_1 + 6^2_2)$$

17 61 22 62, Hum

 $d \cdot (e_1^2 + e_2^2) < d \cdot 2e_2^2$

clushing is difficult.

between

willin - class -2

ML conrequence:

- · different clarses aveleg heavily and are hard to distinguish bard on individual distances
- . Hight hope for <u>amuniphous</u> that tell us that actually, our space does not have ruch a high dimension: . sparsity (~ book by Waterwight)
 - · manifold vorumption

Lef's take a look at matrices...

Estimating avaniance matrices in high dim
Assume on how in data points in d dimension ~ nod "measurements"
Guariance matrix has d'entres... obviews troussie if u « d.
The default columnte for the avaniance matrix, the sample cov. matrix
$$\hat{\Sigma} = \frac{1}{4} \sum_{i=1}^{2} x_i x_i^{i}$$
 might not converp to the true
mudulying cov. matrix.
If u is low and d is high, (In large), the entries of the
sample covariance matrix are unreliable and do not convex
to the true values. If u << d, there is no vary to fix that!!!

Estimating avaniance matrices in high dim Alternatively, if al >> u, one needs to note structural assumptions, for wangle:

· Spilerd Covariance model where the "signal" just lives in few data dimensions

· Ploch - matrix orrunghous

•• •

$$Cov. \in \left(\begin{array}{c} 0\\ 0\\ 0\end{array}\right)$$

The true countrance matrix is the identity.

Eigenvalues in high dimension

If dzcn, Kun eigenvalues concentrate:

Concentration Herrenn (Verhyni'n book Th. 4.7.1)
X pilogausprion vector in
$$\mathbb{R}^{d}$$
, denote by Σ the true covariance
wathix, and by Σ_n the unpirical covariance matrix based on
a somple pointer. Then
 $\mathbb{E}(\mathbb{I}\sum_{n} - \Sigma\mathbb{I}) \stackrel{<}{=} court \cdot (\sqrt{\frac{d}{n}} + \frac{d}{n})\mathbb{I}\Sigma\mathbb{I}$
subgaurrion:
 $p(rxr, rt) \stackrel{<}{=} 2 eq(-f^2, court)$
 $\mathbb{E}(\frac{1}{2} - operation usen (x distance between lager signature)$

Eigenvalues in high dimensions

17 d/n-5 court., then the distribution of eigenvalues is legowser: Courside vandou vectors whole entries are iid with variance 1.

He theorem of Hardwelles-Partur clearacterises the distribution of eigenvalues of the sample row. Matrix as $\frac{4}{3} \rightarrow \lambda$: (It has a closed - prus expression that I did not would to put on the rhide).

General behavior: the large λ , the ware the cignivalues ""prood", see next dide:



Eigunvectors (P(A) in high dim

PCA an high-dim data is problematic, to say the least.
Weighter in a DNN van belæve furmy ar well i K you by bo surolyze kreie Arabishiep

(ree for example paper Ly K. Heliovery)

Outlook: double descent

Finally, let's look at "model rive" (in hour of the number of parameter) as well.

Relation between model tize in HL and onrfitting:



Long debake ... see the Statistical HL lecture vest tour ...