# NMAI059 Probability and statistics 1 Class 14 

Robert Šámal



## Overview

Permutation test

## Bootstrap

## Bayesian statistics

## Sampling random variables

## Situation

- We have two collections of pairwise independent r.v.'s (random samples): .... twonsruple test
- $X_{1}, \ldots, X_{n} \sim F_{X}$ a $Y_{1}, \ldots, Y_{m} \sim F_{Y}$
- We want to decide between $H_{0}: F_{X}=F_{Y}$ and $H_{1}: F_{X} \neq F_{Y}$.
- Examples: running time of an algorithm before/after modification, cholesterol level in people who do/don't eat Miraculous Superfood ${ }^{\top \mathrm{M}}$, frequency of short words in text by authors X and Y .
$\longrightarrow$ We do not assume anything about $F_{X}, F_{Y}$ (in particular they may not be normal).

Method $n=2, m=1 \quad x_{1}=1, x_{2}=9, \varphi_{1}=3$

$$
t_{0 b_{5}}=|5-3|=2
$$

- We choose an appropriate statistics, e.g.

$$
\left(\underline{\left.T\left(X_{1}, \ldots, X_{n}, Y_{1}, \ldots, Y_{m}\right)=\frac{\left|\bar{X}_{n}-\bar{Y}_{m}\right|}{\text { if tols isbed }}, X_{n}, Y_{1}, \ldots, Y_{m}\right)}\right.
$$

Assuming $H_{0}$, , all permutations of the data are the same": $X_{i}$ i $Y_{j}$ were generated from the same distribution.
We randomly permute the given $m+n$ numbers and for each permutation we calculate $T$ - we get numbers $T_{1}, T_{2}, \ldots, T_{(m+n)!}$ (each equally likely).

- $\overline{\overline{A s}} p$-value we take the probability that $T>t_{\text {obs }}$, or

$$
\frac{4}{6}=p_{t}=\frac{1}{(m+n)!} \sum_{j} I\left(T_{j}>t_{\mathrm{obs}}\right) \cdot=\frac{\# d_{j}: \pi_{j}>t_{\mathrm{od}} s}{(k+a)!}
$$

- This is the probability of Type I error. We reject $H_{0}$ whenever $p<\alpha$ (for our choice of $\alpha$, e.g. $\overline{\alpha=0.05}$ ).


## Improvement

- Enumerating all permutations can be too expensive. Instead, we take just an appropriate number $B$ of independently generated permutations and calculate just $B$ values $T_{1}, \ldots, \overline{T_{B}}$.
- As $p$-value we take the estimate of the probability that $T>t_{\text {obs }}$, or

$$
\frac{1}{B} \sum_{j=1}^{B} I\left(T_{j}>t_{\mathrm{obs}}\right)
$$

- For sufficiently large $m, n$ this gives similar results as tests based on CLT. So it is useful especially for medium sized samples.


## Overview



## Sampling random variables

## Empirical CDF - a reminder

- $X_{1}, \ldots, X_{n} \sim F$ i.i.d., $F$ is their CDF
- Definition: Empirical CDF is defined by

$$
\underline{\widehat{F}_{n}(x)}=\frac{\sum_{i=1}^{n} I\left(X_{i} \leq x\right)}{n},
$$

where $I\left(X_{i} \leq x\right)=1$ if $X_{i} \leq x$ and 0 otherwise.


We hope $F=\hat{F}$
weite Êce to sample
$T$ wenk dah
To semple foo $\hat{\hat{E}}$, we daose unfeatiocalon are folvady massued dite.

## Boostrap - basic idea

- from the measured data $X_{1}=x_{1}, \ldots, X_{n}=x_{n} \sim F$ we create $\widehat{F}_{n}$
- other data can be sampled from $\widehat{F}_{n}$
- to do this we select a uniformly random $i \in\{1, \ldots, n\}$ and outputing $x_{i}$


## Bootstrap - basic usage

- $T_{n}=g\left(X_{1}, \ldots, X_{n}\right)$ some statistics (function of the data)
- we want to estimate $\operatorname{var}\left(T_{n}\right)$
- sample $X_{1}^{*}, \ldots, X_{n}^{*} \sim \widehat{F}_{n}$ (see last slide)
- calculate $T_{n}^{*}=g\left(X_{1}^{*}, \ldots, g_{n}^{*}\right)$
- repeat $B$ times to get $T_{n, 1}^{*}, \ldots, T_{n, B}^{*}$
- the variance estimate:

$$
\frac{1}{B} \sum_{b=1}^{B}\left(T_{n, b}^{*}-\frac{1}{B} \sum_{k=1}^{B} T_{n, k}^{*}\right)^{2}
$$


$\operatorname{var}\left(T_{n}\right)$
when ye lion F
$\stackrel{2}{\approx} \operatorname{var}_{\text {when }}\left(T_{0}^{x}\right)$

## Overview

## Permutation test

## Bootstrap

Bayesian statistics

Sampling random variables

## Two approaches to statistics

## Frequentists'/classical approach

- Probability is a long-term frequency (out of 6000 rolls of the dice, a six was rolled 1026 times). It is an objective property of the real world.
- Parameters are fixed, unknown constants. We can't make meaningful probabilistic statements about them.
- We design statistical procedures to have desirable long-run properties. E.g. 95 \% of our interval estimates will cover the unknown parameter.


## Bayesian approach

- Probability describes how much we believe in a phenomenon, how much we are willing to bet. (Prob. that T. Bayes had a cup of tea on December 18, 1760 is $90 \%$.) (Prob. that COVID-19 virus did leak from a lab is ?50? \%.)
- We can make probabilistic statements about parameters (even though they are fixed constants).
- We compute the distribution of $\vartheta$ and form point and interval estimates from it, etc.


## Bayesian method - basic description

- The unknown parameter is treated as a random variable $\Theta$
- We choose prior distribution, the pmf $p_{\Theta}(\vartheta)$ or the pdf $f_{\Theta}(\vartheta)$ independent of the data.
- We choose a statistical model $f_{X \mid \Theta}(x \mid \vartheta)$ that describes what we measure (and with what probability), depending on the value of the parameter.
- After we observe $X=x$, we compute the posterior using distribution $f_{\Theta \mid X}(\overline{\vartheta \mid x})$ or $P_{\Theta \mid X}(v / x)$
- and then derive what we need e.g. find $a, b$ so that

$$
\begin{aligned}
& =\int_{a}^{b} f_{\Theta \mid X}(\vartheta \mid x) d \vartheta \geq 1-\alpha \\
& \text { th. fest if } P(\theta=0) X=x)<\alpha \text { we reject Ho }
\end{aligned}
$$

- $\vartheta=\theta$ lower-case theta, $\Theta$ is upper-case theta


Theorem (Bayes theorem for continuous r.v.'s)
$X, \Theta$ are continuous r.v.s with pdf's $f_{X}, f_{\Theta}$ and joint pdf $f_{X, \Theta}$

$$
f_{\Theta \mid X}(\vartheta \mid x)=\frac{\overbrace{f_{X \mid \Theta}(x \mid \vartheta) f_{\Theta}}(\vartheta)}{\int_{\vartheta^{\prime} \in \mathbb{R}} f_{X \mid \Theta}\left(x \mid \vartheta^{\prime}\right) f_{\Theta}\left(\vartheta^{\prime}\right) d \vartheta^{\prime}} .
$$

(terms with $f_{\Theta}\left(\vartheta^{\prime}\right)=0$ with $f_{\Theta}\left(\vartheta^{\prime}\right)=0$ are considered 0 ).

- Two more variants omitted.


## Bayesian point estimates - MAP and LMS

MAP - Maximum A-Posteriori We choose $\hat{\vartheta}$ to maximize
${ }^{-} p_{\Theta \mid X}(\vartheta \mid x)$ in the discrete case

- $f_{\Theta \mid X}(\vartheta \mid x)$ in the continuous case
- Similar to the ML method in the classical approach if we choose a "flat prior" - $\Theta$ is supposed to be uniform/discrete uniform.


LMS - Least Mean Square Also the conditional mean method.

- We choose $\hat{\vartheta}=\mathbb{E}(\Theta \mid X=x)$.
- Unbiased point estimate, takes the smallest possible Li.ĩ * value

Example 1 caveat: Are $x_{1}, x_{2}-$ independent?

Bayesian spam classifier: $\mathcal{f}\left(k_{1}, \lambda_{2} \ldots\right)$

- create a list of suspicious words (money, win, pharmacy, ...)
- Riv. $X_{i}$ describes whether the email contains the suspicious word $w_{i}$.
- R.v. $\Theta$ describes whether the email is spam $\Theta=1$ or not $\overline{\Theta=0}$.
- From the previous emails, we get estimates of $p_{X \mid \Theta}$ and $p_{\Theta}^{\prime}$
- We use Bayes' theorem to calculate $p_{\Theta \mid X}$

$$
P_{0}(0) \text {. foot. of } u \text { on-spous }
$$

Example 2
Romeo and Juliet are to meet at noon sharp. But Juliet is late by the time described by the random variable, $X \sim U(0, \vartheta)$. We model the parameter $\vartheta$ by the random variable $\Theta \sim U(0,1)$.
What do we infer about $\vartheta$ from the measured value of $X=x$ ?


## Example 3

Observing random variables $X=\left(X_{1}, \ldots, X_{n}\right)$, assume $X_{i} \sim N\left(\vartheta, \sigma_{i}^{2}\right)$ and $\vartheta$ is the value of the random variable $\Theta \sim N\left(x_{0}, \sigma_{0}\right)$. What can we conclude about $\vartheta$ from the measured values $X=x=\left(x_{1}, \ldots, x_{n}\right)$ ?

## Example 4

We flip a coin, the probability of getting heads is $\vartheta$. Out of $n$ flips, the coin comes up heads in $X=k$ cases. If our a priori distribution was $U(0,1)$, what would be the distribution of the posterior distribution?

## Overview

## Permutation test

## Bootstrap

## Bayesian statistics

Sampling random variables

## Basic method - inverse transformation method

Theorem
Let $F$ be a function "of CDF-type": nondecreasing
right-continuous function with $\lim _{x \rightarrow-\infty} F(x)=0$ a
$\lim _{x \rightarrow+\infty} F(x)=1$.
Let $Q$ be the corresponding quantile function. Let $U \sim U(0,1)$ and $X=Q(U)$. Then $X$ has CDF $F$.

- It works well if we can quin施部y $Q$, for example for exponential or geometric distributions.
- The gamma distribution is the sum of several exponential distributions - so we generate it that way.


## Variant of the basic method for discrete variables

- We want a r.v. $X$ that takes values $x_{1}, x_{2}, \ldots$ with probabilities $p_{1}, p_{2}, \ldots\left(\sum_{i} p_{i}=1\right)$.
- We generate $U \sim U(0,1)$.
- Find $i$ such that $p_{1}+\cdots+p_{i-1}<U<p_{1}+\cdots+p_{i}$.
- We set $X:=x_{i}$.

- Works nicely when we have a formula for $p_{1}+\cdots+p_{i}$ (e.g. geometric distribution).
- The binomial distribution is better simulated as the sum of $n$ independent Bernoulli variables.
- There are special tricks for other ones (Poisson).


## Rejection sampling

- We want to generate a rev. with density $f$.
- We can generate a r.v. with density $g$ (which is similar"), namely
- $\frac{f(y)}{g(y)} \leq c$ for some constant $c$. $\quad$ Hz
- The method:


1. Generate $Y$ with density $g$, and $U \sim U(0,1)$. ( $Y$, c. U. $\cdot g(\varphi))$ pas)
2. If $U \leq \frac{f(Y)}{c g(Y)}$, then $X:=Y$.
3. Otherwise, reject the value of $Y, U$ and repeat from point $T$.

- Rationale: generating a random value of $X$ with density $f$ is the same as generating a random point under the graph of the function $f$ whose horizontal $(x)$ coordinate is $X$ (and whose vertical coordinate is uniformly random between 0



## Follow-up classes

$\Longrightarrow$ Probability and Statistics 2 - NMAI073
$\rightarrow$ Introduction to Approximation and Randomized Algorithms - NDMI084

- Introduction to Machine Learning in Python|R NPFL129|NPFL054
- and many master-level lectures

