## Statistical methods in Data Science and AI

DAT405, 2019/2020, Module 3-4. Marina Axelson-Fisk


## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Module 3.1: Bayesian statistics

## CHALMERS

## Probability theory and statistics

- a quick refresher



## Sample space, events and random experiments

- A random experiment is a process that produces random outcomes.
- The sample space is the set of all possible outcomes in an experiment.
- An event is the outcome, or a subset of possible outcomes, of an experiment.



## CHALMERS

## Example: roll a die

- Sample space: $S=\{1,2, \ldots, 6\}=6$ outcomes
- Events:
- "At least 3" $=\{\mathbf{3}, 4,5,6\}$
- "Six" $=\{6\}$
- "Odd" $=\{1,3,5\}$
- Probabilities
$P($ at least 3$)=4 / 6$

$$
\begin{aligned}
& P(\text { six })=1 / 6 \\
& P(\text { odd })=3 / 6
\end{aligned}
$$



## Venn diagrams of set operations

Union: $A \cup B$


S

Intersection: $A \cap B$


S

Mutually exclusive: $A \cap B=\phi$


## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Conditional probability

- The conditional probability of an event $A$ given the knowledge that event $B$ occurred

$$
P(A \mid B)=\frac{P(A \cap B)}{P(B)}=\frac{P(A, B)}{P(B)}
$$

- Note also

$$
P(A, B)=P(A \mid B) P(B)=P(B \mid A) P(A)
$$



## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Mutually exclusive and exhaustive events

Events $E_{1}, E_{2}, \ldots, E_{n}$ are

- mutually exclusive if they cannot occur simultaneously

$$
E_{i} \cap E_{j}=\phi, i \neq j
$$

- exhaustive if they cover the sample space

$$
E_{1} \cup E_{2} \cup \cdots \cup E_{n}=\bigcup_{i=1}^{n} E_{i}=S
$$

Sample space $S$


## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Total law of probability

- For mutually exclusive and exhaustive events $E_{1}, E_{2}, \ldots, E_{n}$ we get for any other event $B$

$$
P(B)=\sum_{i=1}^{n} P\left(B \mid E_{i}\right)
$$



## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Bayes' rule

- Bayes' rule

$$
P(A \mid B)=\frac{P(B \mid A) P(A)}{P(B)}
$$



- For mutually exclusive and exhaustive events
$E_{1}, E_{2}, \ldots, E_{n}$ we get

$$
P(A \mid B)=\frac{P(B \mid A) P(A)}{P(B)}=\frac{P(B \mid A) P(A)}{\sum_{i=1}^{n} P\left(B \mid E_{i}\right)}
$$

## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Example

- Assume that 0.0015 individuals in our population has a certain disease $D$.
- When testing for the disease
- an ill person always tests positive
- a healthy person tests positive with probability 0.0002

- Given that you tested positive, what is the probability that you have the disease?


## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Example (cont.)

Bayes' rule: $\quad \boldsymbol{P}($ ill $\mid+)=\frac{\boldsymbol{P}(+\mid \mathrm{ill}) \boldsymbol{P}(\mathrm{ill})}{\boldsymbol{P}(+)}$

- We have
- $P($ ill $)=0.0015$ and $P($ healthy $)=1-0.0015=0.9985$
- $P(+\mid$ ill $)=1, P(+\mid$ healthy $)=0.002$
- The denominator

- $\boldsymbol{P}(+)=\boldsymbol{P}(+\mid \mathrm{ill}) \boldsymbol{P}($ ill $)+\boldsymbol{P}(+\mid$ healthy $) \boldsymbol{P}($ healthy $)$

$$
\boldsymbol{P}(\text { ill } \mid+)=\frac{\boldsymbol{P}(+\mid \mathrm{ill}) \boldsymbol{P}(\text { ill })}{\boldsymbol{P}(+)}=\frac{1 \cdot 0.0015}{1 \cdot 0.0015+0.002 \cdot 0.9985}=\mathbf{0 . 4 3}
$$

## Random variables and probability distributions

- A random variable is a function of the outcomes in a random experiment.
- Assumes values according to a probability distribution.
- Discrete r.v.: finite or countable number of values,
- Continuous r.v: takes all real values in

$$
\begin{gathered}
P(X=a)=0 \\
P(a \leq X \leq b)=\int_{a}^{b} f(x) d x
\end{gathered}
$$

## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Probability distributions

- Typically depend on one or more parameters
- Common discrete distributions
- Uniform: $\boldsymbol{U}(\boldsymbol{a}, \boldsymbol{b})$
- Binomial: Bin(n,p)
- Geometric: Geo(p)
- Hypergeometric: HGeo(N,K,n)
- Poisson: Poi( $\lambda$ )
- Negative binomial: $N B(r, p)$



## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Probability distributions

- Common continuous distributions
- Uniform: $\boldsymbol{U}[a, b]$
- Normal (Gaussian): $N\left(\mu, \sigma^{2}\right)$
- Student's t: $\boldsymbol{t}_{n-1}$
- Exponential: $\operatorname{Exp}(\lambda)$
- Chi-square: $\chi_{n-1}^{2}$
- Beta: Beta $(\alpha, \beta)$



## CHALMERS

## Statistical inference

Estimation and analysis of these parameters in random samples to draw conclusions of the underlying population.

Two main paradigms:

- Frequentism
- Bayesianism


## Classical or frequentist probability theory:

- Probabilities are frequencies of random repeatable experiments
- Probabilities quantify variability.
- Parameters are (unknown) constants.


## Bayesian probability theory:

- Probabilities correspond to reasonable expectation of an event.
- Probabilities quantify uncertainty.
- Unknown parameters are treated as random variables.


## CHALMERS

UNIVERSITY OF TECHNOLOGY

## DID THE SUN JUST EXPLODE? <br> (TSS NGIT, SO WERE NOT SURE.)



FREQUENTIST STATISTCIAN:


BAYESIAN STATISTCAN:


## Bayes' rule interpretation



We have prior information $P(A)$ of event $A$, and then update the posterior probability $P(A \mid B)$ as more information/data $B$ is achieved.

## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Example



Random number $p \in(0,1)$
Random numbers $q_{1}, q_{2}, q_{3}, \ldots$

- If $q_{i}<p$ Alice wins
- If $q_{i}>p$ Bob wins

First to 6 wins the game.

Only the scores are visible!

## Example

What is the probability that Alice wins?


## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Example



## - Alice • Bob • 5 3

## For known $p$ :

- $P(\mathrm{Bob})=(1-p)^{3}$
- $P($ Alice $)=1-P($ Bob $)$
$p=0.5 \Rightarrow P($ Alice $)=7 / 8$


## Example



Frequentists approach (ML):
$\widehat{\boldsymbol{p}}=5 / 8 \Rightarrow P($ Alice $) \approx 0.95$

Fair odds: 19:1

## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Example

- Bayesian approach
- Consider $\boldsymbol{p}$ a random variable.
- Let $D=\left\{n_{A}=5, n_{B}=3\right\}$ denote our observed data
- The expected probability that Bob wins is given by

$$
E_{B}=\int_{0}^{1}(1-p)^{3} P(p \mid D) d p
$$



- Bayes' rule
likelihood

$$
P(p \mid D)=\frac{P(D \mid p) P(p)}{P(D)}=\frac{P(D \mid p) P(p)}{\int_{0}^{1} P\left(D \mid p^{\prime}\right) P\left(p^{\prime}\right) d p^{\prime}}
$$

## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Example

- The likelihood $P(D \mid p)$ :
- Let $X=$ the number of times Alice wins out of 8
- Probability of winning $=p$

$$
\boldsymbol{X} \sim \operatorname{Bin}\left(\boldsymbol{n}_{\boldsymbol{A}}, \boldsymbol{p}\right)
$$

- The likelihood of observing our data, given $p$ becomes


$$
P(X=5)=\binom{8}{5} p^{5}(1-p)^{3}
$$

## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Example

- The prior $\boldsymbol{P}(\boldsymbol{p})$ :
- Assume $p \sim U(0,1) \Rightarrow P(p)=$ constant

$$
\begin{aligned}
& E_{B}=\int_{0}^{1}(1-p)^{3} P(D \mid p) d p=\frac{\int_{0}^{1} p^{5}(1-p)^{6} d p}{\int_{0}^{1} p^{5}(1-p)^{3} d p}=1 / 11 \\
& E_{A}=1-1 / 11=10 / 11
\end{aligned}
$$



Beta-integral: $\int_{0}^{1} p^{m-1}(1-p)^{n-1} d p=\frac{\Gamma(m) \Gamma(n)}{\Gamma(n+m)}, \Gamma(n)=(n-1)$ !

## CHALMERS

UNIVERSITY OF TECHNOLOGY
Example
Alice


## - Alice • Bob •

## Frequentist approach

- $P($ Alice $) \approx 0.95$
- Fair odds: 19:1

Bayesian approach

- $P($ Alice $) \approx 0.91$
- Fair odds: 10:1

Simulation confirms Bayesian computation!

## Bayesianism versus frequentism

What is the probability of an event?

- Frequentists: the relative frequency of the event in a large number of trials.
- Bayesians: a reasonable expectation, quantifying personal beliefs and prior knowledge, and including the degrees of certainty in these beliefs.



## CHALMERS

## Bayesianism versus frequentism

## Frequentists:

- A distribution parameter $\theta$ is an (unknown) constant.
- $P(\theta=a)=$ ? becomes meaningsless.
- The density of a random variable $X$ : $f_{\theta}(X)$


Bayesians:

- An unknown parameter $\theta$ is treated as a random variable.
- The density of a random variable $X$ is a conditional probability: $\boldsymbol{f}(\boldsymbol{X} \mid \boldsymbol{\theta})$



## CHALMERS

UNIVERSITY OF TECHNOLOGY

## The likelihood function

The likelihood function introduces a third view


- The density of $X$ as a function of $\theta: L_{x}(\theta)$
- Same thing, different names

$$
L_{\theta}(x)=f_{\theta}(x)=f(x \mid \theta)
$$



- But with Bayesian statistics we can use Bayes' theorem on $\theta$

$$
f(\theta \mid x)=\frac{f(x, \theta)}{f(x)}=\frac{f(x \mid \theta) f(\theta)}{\int f\left(x \mid \theta^{\prime}\right) f\left(\theta^{\prime}\right) d \theta^{\prime}}
$$

## Bayesianism versus frequentism

Frequentists:

- $X$ is random, but $\theta$ is not.

Bayesians:

- $\theta$ is random, but after having seen data, $x$ is not



## Frequentism versus Bayesianism

| Frequentism | Bayesianism |
| :--- | :--- |
| + Objective | + More natural |
| + Trade of between errors | + Logically rigourous |
| + Design controls bias | + Can explore different priors |
| + Long prosperous history | + Data can be added |
| - p-value depends on design | - Prior is subjective |
| - Ad-hoc notions of "data more | - Assigning probabilities to |
| extreme" | hypotheses |
| - Fully specified designs ahead |  |

## The effect of different priors

Bayesian feature: different priors will give different posteriors


## CHALMERS

## Example

- Alice has moved to a new city
- She takes the bus to work
- Out of 5 attempts:
- 2 got her to the right place
- 3 forced her to walk another 20 min

What is the proportion of "good" buses for her to take?


## CHALMERS

## Example

Let $\theta=$ the fraction of "good" buses.

- Prior $\boldsymbol{f}(\boldsymbol{\theta})$ : Uniform $(\mathbf{0}, \mathbf{1})$

Let $X=$ the number of good buses of $n$

- Likelihood $f(x \mid \theta): \operatorname{Bin}(n, \theta)$

Observed data:


- $\widehat{\theta}=2 / 5=0.4$

Parameter update, given observed data

- Posterior $\propto$ likelihood $\times$ prior
- $f(\theta \mid x) \propto f(x \mid \theta) f(\theta)$



## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Example

Assume for simplicity $\boldsymbol{\theta} \in\{\mathbf{0 . 0 , 0 . 1}, \mathbf{0} .2, \ldots, \mathbf{0 . 9}, \mathbf{1 . 0}\}=$ 11 values

| $\theta$-values | prior | likelihood | prior $\times$ <br> likelihood | posterior |
| :---: | :---: | :---: | :---: | :---: |
| 0 | 0.0909 | 0 | 0 | 0 |
| 0.1 | 0.0909 | 0.0729 | 0.0066 | 0.0437 |
| 0.2 | 0.0909 | 0.2048 | 0.0186 | 0.1229 |
| 0.3 | 0.0909 | 0.3087 | 0.0281 | 0.1852 |
| 0.4 | 0.0909 | 0.3456 | 0.0314 | 0.2074 |
| 0.5 | 0.0909 | 0.3125 | 0.0284 | 0.1875 |
| 0.6 | 0.0909 | 0.2304 | 0.0209 | 0.1383 |
| 0.7 | 0.0909 | 0.1323 | 0.0120 | 0.0794 |
| 0.8 | 0.0909 | 0.0512 | 0.0047 | 0.0307 |
| 0.9 | 0.0909 | 0.0081 | 0.0007 | 0.0049 |
| 1 | 0.0909 | 0 | 0 | 0 |
| Totals: | 1 |  | 0.1515 | 1 |



## CHALMERS

## Example

We can predict new values

$$
\begin{aligned}
& P(\operatorname{good} \text { bus tomorrow } \mid x)= \\
& =\sum_{\theta} P(\operatorname{good} \text { bus tomorrow } \mid \theta, x) p(\theta \mid x) \\
& =\sum_{\theta} \theta \cdot p(\theta \mid x) \\
& =0.429
\end{aligned}
$$



## CHALMERS

## The effect of different priors

- Prior 1: $\boldsymbol{U}(\mathbf{0}, 1)$
- $\boldsymbol{p}(\boldsymbol{\theta})=$ const
- Prior 2:
- $p(\theta) \propto \theta^{-\frac{1}{2}}(1-\theta)^{-\frac{1}{2}}$
- more weight on extreme values
- Prior 3:
- $p(\theta) \propto \theta^{\mathbf{1 0 0}}(1-\theta)^{\mathbf{1 0 0}}$
- most weight in the centre $\theta=0.5$



## CHALMERS

## The effect of different priors

- Prior 1: $\boldsymbol{U}(\mathbf{0}, 1)$
- $\boldsymbol{p}(\boldsymbol{\theta})=$ const
$\sim \operatorname{Beta}(1,1)$
- Prior 2:

- $p(\theta) \propto \theta^{-\frac{1}{2}}(1-\theta)^{-\frac{1}{2}}$
- more weight on extreme values
$\sim \operatorname{Beta}\left(\frac{1}{2}, \frac{1}{2}\right)$
- Prior 3:
- $p(\theta) \propto \theta^{100}(1-\theta)^{100}$
$\sim \operatorname{Beta}(101,101)$
- most weight around $\theta=0.5$


## The effect of different priors

- Prior 1: $\operatorname{Beta}(\mathbf{1}, \mathbf{1})$
- Prior 2: Beta $\left(\frac{1}{2}, \frac{1}{2}\right)$
- Prior 3: $\operatorname{Beta}(101,101)$
- Posterior 1: $\operatorname{Beta}(3,4)$
- Posterior 2: $\operatorname{Beta}(2.5,2.5)$
- Posterior 3: Beta $(103,104)$

```
Beta-prior + binomial likelihood }=>\mathrm{ Beta-posterior
Beta}(\boldsymbol{\alpha},\boldsymbol{\beta})+"x\mathrm{ of }\boldsymbol{n}\mathrm{ successes" }=>\operatorname{Beta}(\boldsymbol{\alpha}+\boldsymbol{x},\boldsymbol{\beta}+\boldsymbol{n}
```


## CHALMERS

## Example

- Prior 1: $P($ good bus tomorrow $\mid x) \approx 0.429$
- Prior 2: $P(\operatorname{good}$ bus tomorrow $\mid x) \approx 0.417$
- Prior 3: $\mathbf{P}($ good bus tomorrow $\mid x) \approx 0.498$



## CHALMERS

## The effect of different priors

- The more data, the less important the prior



## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Conjugate priors

- We have a sample of observed data: $x_{1}, \ldots, x_{n}$
- We have a corresponding likelihood function (or samling distribution): $\boldsymbol{f}(\boldsymbol{x} \mid \boldsymbol{\theta})$
- A prior $f(\theta)$ is called a conjugate prior if the corresponding posterior $f(\theta \mid x)$ belongs to the same family of distributions.

Bayes' theorem:

$$
f(\theta \mid x)=\frac{f(x \mid \theta) f(\theta)}{f(x)}
$$

## Conjugate priors

| Likelihood | Parameter | Prior | Posterior |
| :--- | :---: | :--- | :--- |
| Bernoulli |  |  |  |
| Binomial |  | Beta | Beta |
| Geometric |  |  |  |
| Negative binomial |  |  |  |
| Exponential | $\lambda$ |  | Gamma |

## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Conjugate priors

- Conjugacy is mutual, e.g.

Dirichlet $\propto$ Multinomial $\times$ Dirichlet
Multinomial $\propto$ Dirichlet $\times$ Multinomial

Bayes' theorem:

$$
f(\theta \mid x)=\frac{f(x \mid \theta) f(\theta)}{f(x)}
$$

## CHALMERS

UNIVERSITY OF TECHNOLOGY

## The exponential family of distributions

- The exponential family of distributions over $x$, given parameters $\eta$, takes the form

$$
f(\boldsymbol{x} \mid \boldsymbol{\eta})=h(\boldsymbol{x}) g(\boldsymbol{\eta}) \exp \left\{\boldsymbol{\eta}^{\mathrm{T}} u(\boldsymbol{x})\right\}
$$

- The function $u(x)$ is called a sufficient statistic for $\eta$, i.e. it contains all information needed to estimate $\eta$.
- All members of the exponential family has conjugate priors.
- Products of exponential family members also have conjugate priors.



## CHALMERS

## Example: the Bernoulli distribution

$$
\begin{array}{rlrl}
f(x \mid p) & =p^{x}(1-p)^{1-x} & & f(\boldsymbol{x} \mid \boldsymbol{\eta})=h(\boldsymbol{x}) g(\boldsymbol{\eta}) \exp \left\{\boldsymbol{\eta}^{\mathrm{T}} u(\boldsymbol{x})\right\} \\
& =\exp \left\{\ln \left(p^{x}(1-p)^{1-x}\right\}\right. & \\
& =\exp \{x \ln p+(1-x) \ln (1-p)\} & & \\
& =\exp \left\{x \ln \left(\frac{p}{1-p}\right)+\ln (1-p)\right\} & \text { substitute }\left[\eta=\ln \left(\frac{p}{1-p}\right)\right] \\
& =\exp \left\{x \eta-\ln \left(1+e^{\eta}\right)\right\} &
\end{array}
$$

## Examples of exponential family members

- Bernoulli
- Geometric
- Gamma
- Exponential
- Poisson
- Beta
- Normal
- Beta
- Dirichlet
- Chi-squared


## Also:

- Binomial, with fixed number of trials
- Multinomial, with fixed number of trials
- Negative binomial, with fixed number of failures



## Uninformative priors

- When nothing is known, we may want to play equal weights to all parameter values
$\Rightarrow$ Uniform distribution
+ Gives the same parameter estimate as Maximum Likelihood
- Not invariant under parameterization

$$
X \sim U[a, b], Y=f(X) \nRightarrow Y \sim U[f(a), f(b)]
$$

$\Rightarrow$ Large variation in posterior

## CHALMERS

## Jeffrey's prior

- Uninformative prior
- Invariant under transformation
- Given by

$$
\boldsymbol{p}(\boldsymbol{\theta}) \propto \sqrt{\operatorname{det}(\mathcal{J}(\boldsymbol{\theta}))}
$$

where $\mathcal{J}(\theta)$ is the Fisher information

$$
\mathcal{J}(\theta)=-E_{\theta}\left[\frac{d^{2} \log f(X \mid \theta)}{d \theta^{2}}\right]
$$

Flat prior


## CHALMERS

## Fisher information

- For a random variable $X$ with density $f(x \mid \theta)$ :

The Fisher information =
$=$ "information content of $X$ in terms of estimating $\theta$ "


## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Example: Jeffrey's prior

Let $X \sim \operatorname{Bin}(\boldsymbol{n}, \boldsymbol{p})$. We want a prior for $\boldsymbol{p}$.

$$
\begin{aligned}
& f(x \mid p)=\binom{n}{x} p^{x}(1-p)^{n-x} \\
& \log f(x \mid p)=x \log p+(n-x) \log (1-p) \\
& \frac{d}{d p} \log f(x \mid p)=\frac{x}{p}-\frac{n-x}{1-p} \\
& \frac{d^{2}}{d p^{2}} \log f(x \mid p)=-\frac{x}{p^{2}}-\frac{n-x}{(1-p)^{2}} \\
& \mathcal{J}(p)=-E_{p}\left[\frac{d^{2}}{d p^{2}} \log f(x \mid p)\right]=-\frac{n p}{p^{2}}-\frac{n-n p}{(1-p)^{2}}=\frac{n}{p(1-p)}
\end{aligned}
$$



$$
E_{p}[X]=n p
$$

## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Example: Jeffrey's prior

$$
f(p) \propto \sqrt{\mathcal{J}(p)} \propto p^{-1 / 2}(1-p)^{-1 / 2} \sim \operatorname{Beta}\left(\frac{1}{2}, \frac{1}{2}\right)
$$

Note: Jeffrey's prior is generally not conjugate.

Flat prior


## CHALMERS

## Reference priors

- Maximize the distance between the prior and the posterior
- Kullback-Leibler divergence
- Hellinger distance



## Inference in classical statistic

Based on a sample $x_{1}, \ldots, x_{n}$ from some density $f_{\theta}(x)$.
Parameter estimation
Estimate the parameter $\theta$ using Maximum Likelihood

$$
\widehat{\theta}=\operatorname{argmax}_{\theta} \log L_{\theta}\left(x_{1}, \ldots, x_{n}\right)
$$

Confidence intervals
A 95\% confidence interval for $\theta$ is an interval $\left(\theta^{\text {lo }}, \theta^{\text {up }}\right)$ such that

$$
P\left(\theta^{\mathrm{lo}} \leq \theta \leq \theta^{\text {up }}\right)=0.95 .
$$

Note: $\theta^{\text {lo }}$ and $\theta^{\text {up }}$ are random variables, not $\theta$.

## Inference in classical statistic

Based on a sample $x_{1}, \ldots, x_{n}$ from some density $f_{\theta}(x)$.
Hypothesis testing
We want to test the hypothesis

$$
\begin{aligned}
& \boldsymbol{H}_{0}: \boldsymbol{\theta}=\boldsymbol{\theta}_{\mathbf{0}} \\
& \boldsymbol{H}_{1}: \boldsymbol{\theta} \neq \boldsymbol{\theta}_{\mathbf{0}}
\end{aligned}
$$

using som test statistic $T$ (function of the sample). We reject $H_{0}$ on $5 \%$ significance level if $\widehat{\boldsymbol{\theta}} \geq \boldsymbol{\theta}^{\text {up }}$ or $\widehat{\boldsymbol{\theta}} \leq \boldsymbol{\theta}^{\text {lo }}$ where again

$$
P\left(\theta^{\mathrm{lo}} \leq \boldsymbol{\theta} \leq \boldsymbol{\theta}^{\text {up }}\right)=\mathbf{0 . 9 5}
$$

## Bayesian inference: parameter estimation

Based on a sample $D=\left\{x_{1}, \ldots, x_{n}\right\}$ from some density $f(x \mid \theta)$.
Parameter estimation
The most likely estimate of $\theta$ is the maximum of the posterior

$$
\widehat{\boldsymbol{\theta}}=\operatorname{argmax}_{\boldsymbol{\theta}} \boldsymbol{P}(\boldsymbol{\theta} \mid \boldsymbol{D})
$$

Note: if the prior $\mathbf{P}(\boldsymbol{\theta})$ is the uniform distribution, then this estimate is the same as the Maximum Likelihood estimate.

## Bayesian inference: credible intervals

Based on a sample $D=\left\{x_{1}, \ldots, x_{n}\right\}$ from some density $f(x \mid \theta)$.
Credible intervals
A 95\% credible interval for $\boldsymbol{\theta}$ is an interval ( $\left.\boldsymbol{\theta}^{\text {lo }}, \boldsymbol{\theta}^{\text {up }}\right)$ such that the posterior

$$
\begin{gathered}
P\left(\theta^{\mathrm{lo}} \leq \theta \leq \theta^{\mathrm{up}} \mid D\right)=0.95 \\
\int_{\theta}^{\theta^{\mathrm{up}}} p(\theta \mid D) d \theta=0.95
\end{gathered}
$$

Note: Now $\boldsymbol{\theta}$ is a random variable with prior $\mathbf{P}(\boldsymbol{\theta})$.

## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Bayesian inference: credible intervals

However, the credible interval is not unique. We need additional conditions.
For an ( $1-\alpha$ )-interval

- Equal-tailed interval (ETI)

$$
P\left(\theta \leq \theta^{\mathrm{lo}} \mid D\right)=P\left(\theta \geq \theta^{\text {up }}\right)=\alpha / 2
$$



## CHALMERS

UNIVERSITY OF TECHNOLOGY

## Bayesian inference: credible intervals

However, the credible interval is not unique. We need additional conditions.
For an ( $1-\alpha$ )-interval

- Equal-tailed interval (ETI)

$$
P\left(\theta \leq \theta^{\mathrm{lo}} \mid D\right)=P\left(\theta \geq \theta^{\text {up }}\right)=\alpha / 2
$$

- Highest density interval (HDI)

$$
\begin{aligned}
& \mathcal{C}=\{\theta: p(\theta \mid D) \geq k\} \text { where } \\
& \int_{\theta: p(\theta \mid D) \geq k} p(\theta \mid D) d \theta=1-\alpha
\end{aligned}
$$



## Confidence intervals vs credible intervals

- A 95\% credible interval contains the true value $\boldsymbol{\theta}$ with probability $95 \%$.
- i.e. based on data there is a $95 \%$ probability that the interval contains $\boldsymbol{\theta}$
- Statement after data is collected
- A 95\% confidence interval contains the true value of $\boldsymbol{\theta} 95 \%$ of the time.
- i.e. $95 \%$ of the samples we draw will cover the true value of $\boldsymbol{\theta}$
- Statement before data is collected


## Bayesian inference: hypothesis testing

## Bayes factor

We want to test the hypothesis

$$
\begin{aligned}
& H_{0}: \theta=\theta_{0} \\
& H_{1}: \theta \neq \theta_{0}
\end{aligned}
$$

The Bayes factor is the ratio of the posteriors

$$
\frac{P\left(H_{1} \mid D\right)}{P\left(H_{0} \mid D\right)}=\frac{P\left(D \mid H_{1}\right)}{P\left(D \mid H_{0}\right)}, \frac{P\left(H_{1}\right)}{P\left(H_{0}\right)}
$$

## Summary

- Bayesianism versus frequentism
- The choice of priors
- Conjugate priors
- Uninformative priors

- Jeffrey's prior
- Reference priors
- Exponential family
- Frequentist versus Bayesian inference
- Parameter estimation
- Confidence intervals - credible intervals
- Hypothesis testing - Bayes factor

