

Probability measure

- Given space Ω ("abstract space").
- Need a σ -algebra \mathcal{A} on Ω ("measurable events")
 - $A \in \mathcal{A} \Rightarrow A^c \in \mathcal{A}$
 - $(A_i)_{i \in \mathbb{N}} \subset \mathcal{A} \Rightarrow \bigcup_{i=1}^{\infty} A_i \in \mathcal{A}$ ("countable union")
 - $\emptyset, \Omega \in \mathcal{A}$
 - countable intersections

• A measure μ on (Ω, \mathcal{A}) is a function

$$\mu: \mathcal{A} \rightarrow [0, \infty]$$

that is countably additive: If $(A_i)_{i \in \mathbb{N}}$ is a sequence of pairwise disjoint sets, then

$$\mu\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mu(A_i).$$

A measure P on a measurable space (Ω, \mathcal{A}) is called a probability measure if $P(\Omega) = 1$.

The elements in \mathcal{A} are called events.

Then (Ω, \mathcal{A}, P) is called a probability space.

Example (1): Throw one die

$\Omega = \{1, 2, \dots, 6\}$, $\mathcal{A} = \mathcal{P}(\Omega)$ (σ -algebra generated by the "elementary events" $\{1\}, \{2\}, \dots, \{6\}$).

P can be defined uniquely by assigning

$$P(\{1\}) = P(\{2\}) = \dots = P(\{6\}) = \frac{1}{6}$$

For example $P(\{1, 5\}) = P(\{1\}) + P(\{5\}) = \frac{1}{3}$

Throw two dice:

$$\Omega = \{1, 2, \dots, 6\} \times \{1, 2, \dots, 6\}$$

$$\mathcal{A} = \mathcal{P}(\Omega)$$

$$P(\{(i, j)\}) = \frac{1}{36}$$

first die second
 $= \{(1, 1), (1, 2), (1, 3), \dots\}$

Example (2) Normal distribution

$$\Omega = \mathbb{R}$$



$\mathcal{A} =$ Borel- σ -algebra

$$f_{\mu, \sigma} : \mathbb{R} \rightarrow \mathbb{R},$$

$$x \mapsto \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$$

$$P : \mathcal{A} \rightarrow [0, 1]$$

$$P(A) := \int_A f_{\mu, \sigma}(x) dx$$



Different types of probability measures

Discrete measure:

$\Omega = \{x_1, x_2, \dots\}$ finite or at most countable.

$$\mathcal{A} = \mathcal{P}(\Omega)$$

We define a probability measure $P: \mathcal{A} \rightarrow [0, 1]$ by assigning probabilities to the "elementary events":

$$P(\{x_i\}) =: p_i$$

with $0 \leq p_i \leq 1$, $\sum_i p_i = 1$.

For $A \in \mathcal{A}$ we assign

$$P(A) = \sum_{\{i | x_i \in A\}} p_i.$$

Examples: toss a coin; distribution on \mathbb{Q}


Dirac measures:

For $x \in \mathbb{R}$, we define the Dirac measure δ_x on

$(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ by setting

$$\delta_x(A) = \begin{cases} 1 & \text{if } x \in A \\ 0 & \text{otherwise} \end{cases}$$

Sometimes this is called a point mass at a point x .



A discrete measure on \mathbb{R} can be written as a sum of Dirac measures. For example, throwing a die can be described as

$$\frac{1}{6} (\delta_1 + \delta_2 + \dots + \delta_6)$$

Measures with a density

Consider $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$ and the Lebesgue measure λ .


Consider a function $f: \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ that is measurable and satisfies $\int f d\lambda = 1$. (= $\int f(x) dx$)

Then we define a measure ν on \mathbb{R}^n by setting, for all $A \in \mathcal{A}$,

$$\nu(A) := \int_A f(x) dx.$$

ν is the probability measure on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$ with density f .

Notation: $\nu = f \cdot \lambda$




Question? Can we describe every prob measure on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$ in terms of a density? Answer: no!

Counterexample: δ_0 Dirac measure

Def A prob. measure ν on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$ is called absolutely continuous with respect to another measure μ on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$ if every μ -null set is also a ν -null set:


$$\forall B \in \mathcal{B}(\mathbb{R}^n): \mu(B) = 0 \Rightarrow \nu(B) = 0.$$

Notation: $\nu \ll \mu$



$$\mu(A) = 0 \Rightarrow \underbrace{\int_A f d\mu}_{\nu(A)} = 0$$

Example: $N(0,1) \ll \lambda$



Example: $\delta_0 \not\ll \lambda$ because

$$\lambda(\{0\}) = 0 \text{ but } \delta_0(\{0\}) = 1.$$

Theorem (Radon-Nikodym)

Consider two prob. measures ν, μ on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$. Then the following two statements are equivalent:

(see next page)

(1) ν has a density wrt μ .

(2) ν is absolutely continuous wrt μ .

Proof idea

(1) \Rightarrow (2) easy

(2) \Rightarrow (1) We need to construct a density!

Consider the set \mathcal{G} of all functions g with the following properties:

$$(*) \left\{ \begin{array}{l} \bullet g \text{ is measurable, } g \geq 0 \\ \bullet g \cdot \mu \leq \nu, \text{ that is} \\ \forall A \in \mathcal{F}(\mathbb{R}^n): \int_A g d\mu \leq \nu(A). \end{array} \right.$$

- Observe: $g \equiv 0$ satisfies $(*)$, so \mathcal{G} is not empty.
- If g, h both satisfy $(*)$, then $\sup(g, h)$ satisfies $(*)$.
- Define $\gamma := \sup_{g \in \mathcal{G}} \int g d\mu$ and construct a sequence $(g_n)_{n \in \mathbb{N}}$ such that $\lim \int g_n d\mu = \gamma$.
- Define "density" $f := \sup g_n$
- Now prove: f does the job. ▣

Def μ, ν measures on (Ω, \mathcal{A}) . ν is called singular wrt μ if there exists $A \in \mathcal{A}$ such that $\mu(A) = 0$ but $\nu(A^c) = 0$. Notation: $\mu \perp \nu$.




Example: $\lambda \perp \delta_0$

Theorem (Decomposition by Lebesgue)

μ, ν prob. measures on (Ω, \mathcal{A}) . Then there exists a unique decomposition $\nu = \nu_1 + \nu_2$ such that $\nu_1 \ll \mu$ and $\nu_2 \perp \mu$.

Example: $\nu = \frac{1}{2}(N(0,1) + \delta_0)$



$\nu = \nu_1 + \nu_2$ where $\nu_1 = \frac{1}{2} N(0,1)$, $\nu_2 = \frac{1}{2} \delta_0$.

Proof Let \mathcal{N}_μ be the set of all null-sets wrt μ . $\subset \mathcal{A}$.

$$\alpha := \sup \{ \nu(A) \mid A \in \mathcal{N}_\mu \}$$

Can construct a countable sequence $(A_n)_{n \in \mathbb{N}}$, $A_n \in \mathcal{N}_\mu$,

such that $\nu(A_n) \nearrow \alpha$. By countable additivity \bullet

we get $\nu\left(\underbrace{\bigcup_{n \in \mathbb{N}} A_n}_{=: N}\right) = \alpha$.

Define $\nu_1: A \mapsto \nu(A \cap N^c)$

$\nu_2: A \mapsto \nu(A \cap N)$

Does the job. ■

Cantor-distribution: non-trivial distribution that
is singular w.r.t λ

Construct the Cantor set:

• Start with $C_0 := [0, 1]$

"Remove middle part"



• $C_1 := [0, 1/3] \cup [2/3, 1]$

"Remove middle parts from all intervals"



• $C_2 =$




⋮



The Cantor set is the limit in this process.

Now construct a probability distribution:

Consider the cdf of the sets C_0, C_1, C_2, \dots



\vdots

Take limit. $\nearrow \nu$ Can prove many strange properties:

- Cantor-set is compact, non-empty, empty interior.
- The cdf of " ν " is continuous. ν is a prob. meas.
- But: $\lambda(C) = 0$.

$$\Rightarrow \lambda \perp \nu$$

Cumulative distribution function

Let P is a prob. measure on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$.

Define the function $F: \mathbb{R} \rightarrow \mathbb{R}$, $x \mapsto P([-\infty, x])$.

We say that F is a cumulative distribution function (cdf), that is it satisfies the following ~~prob~~ properties:

(i) F is monotonically increasing,


$$\lim_{x \rightarrow -\infty} F(x) = 0, \quad \lim_{x \rightarrow +\infty} F(x) = 1.$$


(ii) F is continuous from the right:

$(x_n)_n$ sequence with $x_n \searrow x$

(i.e. $x_n \geq x_{n+1}$ and $x_n \rightarrow x$) then also

$$F(x_n) \rightarrow F(x).$$





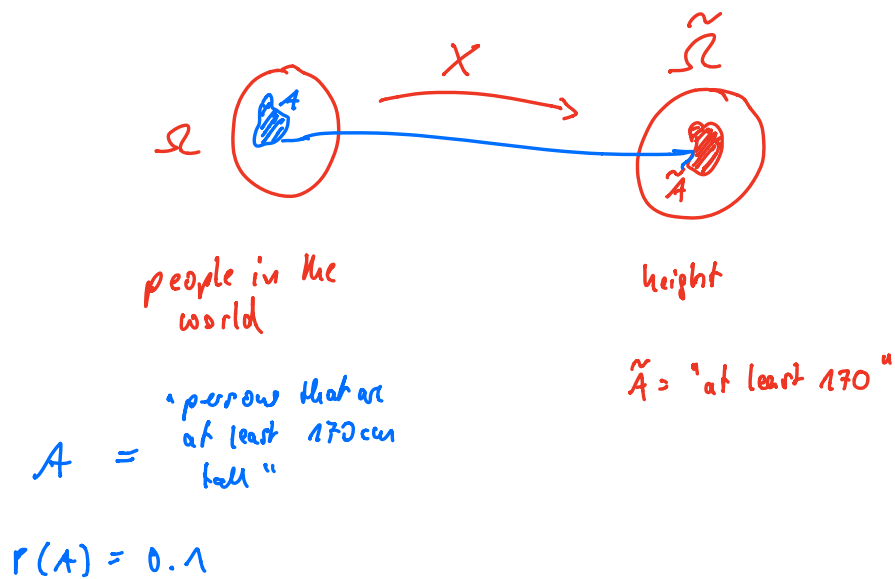
Let $F: \mathbb{R} \rightarrow \mathbb{R}$ be a function with properties (i) and (ii).
 Then there exist a unique prob. measure P on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$
 such that $P((-\infty, x]) := F(x)$.

Random variable

Def Let (Ω, \mathcal{A}, P) be a probability space, $(\tilde{\Omega}, \tilde{\mathcal{A}})$ be another measurable space. A mapping: $X: \Omega \rightarrow \tilde{\Omega}$

is called a random variable if X is measurable, i.e.

$$\forall \tilde{A} \in \tilde{\mathcal{A}} : X^{-1}(\tilde{A}) := \{\omega \in \Omega \mid X(\omega) \in \tilde{A}\} \in \mathcal{A}.$$



Example: sum of two dice

$$\Omega = \{(i, j) \mid i, j \in \{1, \dots, 6\}\}$$

$$\mathcal{A} = \mathcal{P}(\Omega)$$

$$P(\{(i, j)\}) = \frac{1}{36}$$

X "sum of the two values"

$$X: \Omega \rightarrow \{2, \dots, 12\}, (i, j) \mapsto i + j$$

is measurable.

$$\tilde{\Omega} = \{2, \dots, 12\}$$

$$\tilde{\mathcal{A}} = \mathcal{P}(\tilde{\Omega})$$

Def A random variable $X: \Omega \rightarrow \tilde{\Omega}$ induces a measure on the target space:

For $\tilde{A} \in \tilde{\mathcal{F}}$ we define

$$\underline{P}_X(\tilde{A}) := P(X^{-1}(\tilde{A}))$$

This is a probability measure on $(\tilde{\Omega}, \tilde{\mathcal{F}})$ and it is called the distribution of X .

Def $X: (\Omega, \mathcal{F}, P) \rightarrow (\tilde{\Omega}, \tilde{\mathcal{F}})$. Then the family


$$\sigma(X) := \{X^{-1}(\tilde{A}) \mid \tilde{A} \in \tilde{\mathcal{F}}\}$$

is a σ -algebra on Ω and it is called the σ -algebra induced by X


(it is the smallest σ -algebra on Ω that makes X measurable).

Conditional probabilities

Notation: $P(A \cap B) = P(\text{"A and B"})$



$P(A \cup B) = P(\text{"A or B"})$



Def (Ω, \mathcal{A}, P) probability space,
 $A, B \in \mathcal{A}$, $P(B) > 0$. Then

$P(A | B) := \frac{P(A \cap B)}{P(B)}$ is called the


conditional probability of A given B.

Theorem The mapping $P_B: \mathcal{A} \rightarrow [0, 1]$, $A \mapsto P(A|B)$ is a probability measure on (Ω, \mathcal{A}) , it is called the conditional distribution of P with respect to B .

Example: $\Omega = \text{all persons on earth}$
 $\mathcal{A} = \mathcal{P}(\Omega)$
 $P = \text{"uniform"}$

Event $A := \text{"person has been vaccinated"}$
 $B := \text{"person has disease"}$

$P(\text{disease} \mid \text{vaccinated})$



$P(\text{vaccinated} \mid \text{disease})$

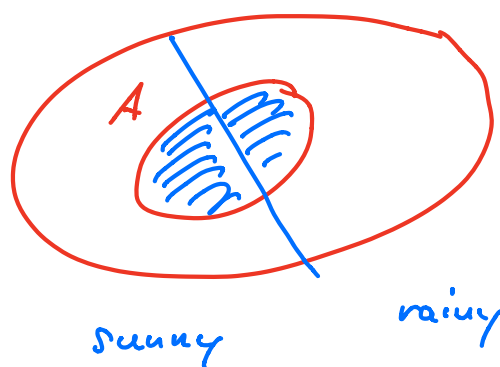
Example: two dice

$P(\text{"sum is 10"} \mid \text{"first die was 5"})$

Bayes formula

Law of total probability: Let B_1, B_2, \dots, B_k be a disjoint partition of Ω with $B_i \in \mathcal{A}$ for all i , and $A \in \mathcal{A}$. Then

$$P(A) = \sum_{i=1}^n P(A|B_i) \cdot P(B_i) = \sum_{i=1}^n P(A \cap B_i)$$



Bayes formula:

$$P(B_i | A) = \frac{P(A|B_i) \cdot P(B_i)}{\sum_i P(A|B_i) \cdot P(B_i)} = \frac{P(A \cap B_i)}{P(A)}$$

Example: breast cancer screening

Assume 1% of all women above 40 have breast cancer.

90% of women with breast cancer will be test positive. ("true positives")

8% of women without breast cancer will receive a positive result as well ("false positives")

Given that a woman ~~receives~~ receives a positive test result, what is the likelihood that she has breast cancer?

$$P(\text{cancer} | \text{positive}) = \frac{P(\text{positive} | \text{cancer}) \cdot P(\text{cancer})}{P(\text{pos.} | \text{cancer}) P(\text{cancer}) + P(\text{pos} | \text{not cancer}) \cdot P(\text{not cancer})}$$

$$= \frac{0.9 \cdot 0.01}{0.9 \cdot 0.01 + 0.08 \cdot 0.99} \approx 10\%$$

Independence

Consider a probability space (Ω, \mathcal{F}, P) . Two events $A, B \in \mathcal{F}$ are called independent if

$$P(A \cap B) = P(A) \cdot P(B)$$

Observation: A is independent of $B \Leftrightarrow P(A|B) = P(A)$

A family of events $(A_i)_{i \in I}$ is called independent if for all finite subsets $J \subset I$ we have

$$P\left(\bigcap_{i \in J} A_i\right) = \prod_{i \in J} P(A_i).$$

(Family is called pairwise independent if $\forall i, j \in I$:
 $P(A_i \cap A_j) = P(A_i) \cdot P(A_j)$. This does not
imply independence!)

Two random variables $X: \Omega \rightarrow \Omega_1$, $Y: \Omega \rightarrow \Omega_2$
are called independent if their induced σ -algebras $\sigma(X)$, $\sigma(Y)$
are independent:

$$\forall A \in \sigma(X), B \in \sigma(Y): P(A \cap B) = P(A) \cdot P(B).$$

Notation for independence:

$A \perp B$

$X \perp Y$

Expectation (discrete case)

Considers a discrete random variable $X: \Omega \rightarrow \mathbb{R}$
(that is, $X(\Omega)$ is at most countable).

Definition (Ω, \mathcal{A}, P) prob. space, $S \subset \mathbb{R}$ at most countable, $X: \Omega \rightarrow S$ random variable.

If $\sum_{r \in S} |r| \cdot P(X=r) < \infty$, then

$E(X) := \sum_{r \in S} r \cdot P(X=r)$ is called the expectation of X .

(sometimes people write EX , $\mathbb{E}X$ or $\mathbb{E}(X)$).

Examples

- Toss a coin. $\Omega = \{\text{head}, \text{tail}\}$, $\mathcal{A} = \mathcal{P}(\Omega)$, $P(\text{head}) = p$
 $P(\text{tail}) = 1-p$.
 $0 < p < 1$.

$X: \Omega \rightarrow \{0, 1\}$, head $\mapsto 1$, tail $\mapsto 0$.

$$E(X) = 0 \cdot \underbrace{P(X=0)}_{1-p} + 1 \cdot \underbrace{P(X=1)}_p = p.$$

- Test error of a classifier.

Def A rv is called "centered" if $E(X) = 0$.

Important properties:

• Linear: $E(a \cdot X + b \cdot Y) = a \cdot E(X) + b \cdot E(Y)$.

$\begin{matrix} \text{rv} & & \text{rv} \\ \downarrow & & \downarrow \\ \in \mathbb{R} & & \in \mathbb{R} \end{matrix}$

• X, Y independent $\Rightarrow E(X \cdot Y) = E(X) \cdot E(Y)$

$$\sum_{i,j} |x_i y_j| P(X=x_i, Y=y_j) \stackrel{\text{ind.}}{=} \sum_{i,j} |x_i y_j| P(X=x_i) \cdot P(Y=y_j)$$

$$= \sum_{i,j} \underbrace{|x_i y_j|}_{|x_i| |y_j|} P(X=x_i) \cdot P(Y=y_j)$$

$$= \underbrace{\left(\sum_i |x_i| P(X=x_i) \right)}_{< \infty} \underbrace{\left(\sum_j |y_j| P(Y=y_j) \right)}_{< \infty}$$

Variance, covariance, correlation (discrete case)


Def $X, Y: (\Omega, \mathcal{A}, \mathbb{P}) \rightarrow \mathbb{R}$ discrete rvs with $E(X^2) < \infty$, $E(Y^2) < \infty$.

Then $\text{Var}(X) := E((X - E(X))^2)$

is called the variance of X

and $\sqrt{\text{Var}(X)} =: \sigma_X$

is called the standard deviation.



high variance
moderate variance

$\text{Cov}(X, Y) := E((X - E(X)) \cdot (Y - E(Y)))$ is called the covariance of X and Y .

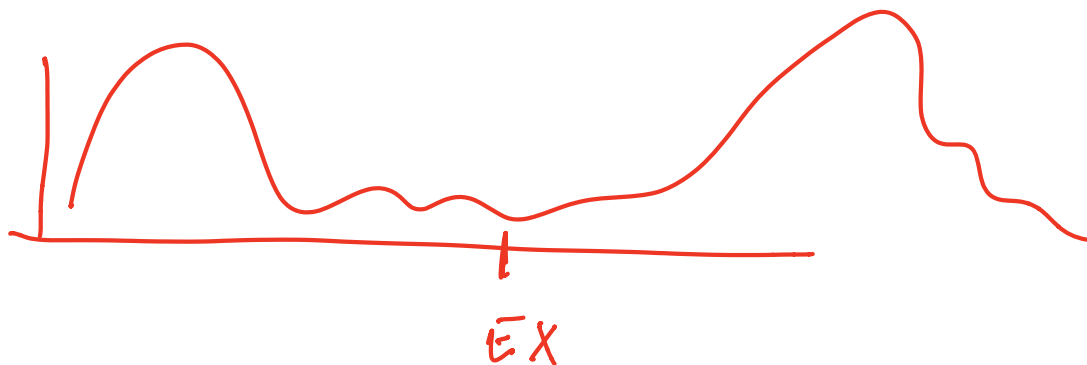
$\rho_{XY} := \frac{\text{Cov}(X, Y)}{\sigma_X \cdot \sigma_Y} \in [-1, 1]$ is called the correlation coefficient.

If $\text{Cov}(X, Y) = 0$, then X and Y are called uncorrelated.

More generally, for $k \in \mathbb{N}$ we define the terms


$E(X^k)$ ("k-th moment"),

$E((X - E(X))^k)$ ("k-th centered moment")




intuition about covariance


$$\text{Cov}(X, Y) = E((X - E(X)) \cdot (Y - E(Y)))$$



positive, large covariance
 $\rho \approx 0.9$




negative cov,
 large in absolute values
 $\rho \approx -0.9$



$\text{Cov} \approx 0$
 (uncorrelated).



Uncorrelated $\not\Rightarrow$ independence!



Properties

- $\text{Var}(X) = E(X^2) - (E(X))^2$
- $\text{Cov}(X, Y) = E(X \cdot Y) - E(X) \cdot E(Y)$
- $E(aX + b) = a \cdot E(X) + b$
- $\text{Var}(a \cdot X + b) = a^2 \text{Var}(X)$
- $\text{Cov}(X, Y) = \text{Cov}(Y, X)$
- $\text{Var}(X + Y) = \text{Var}(X) + \text{Var}(Y) + 2 \text{Cov}(X, Y)$
- X, Y independent $\Rightarrow \text{Cov}(X, Y) = 0$
 $\not\Leftarrow$
- X, Y independent $\Rightarrow \text{Var}(X + Y) = \text{Var}(X) + \text{Var}(Y)$.

Expectation and variance in the general setting

$$L^k(\Omega, \mathcal{A}, P) := \left\{ X: \Omega \rightarrow \mathbb{R} \mid X \text{ measurable and } \int_{\Omega} |X^k| dP < \infty \right\}$$

(Ω, \mathcal{A}, P) prob. space, $X: \Omega \rightarrow \mathbb{R}$ with distribution

$P_X = X(P)$, $X \in L^1(\Omega, \mathcal{A}, P)$. The **expectation**

of X is then defined as

$$E(X) := \int_{\Omega} X dP = \int_{\mathbb{R}} x dP_X(x)$$

(case of density f :

$$\int_{\mathbb{R}} x f(x) dx)$$

If $X^k \in L^1(\Omega, \mathcal{A}, P)$ then

$E(X^k) = \int X^k dP$ is called the k -th moment of X .

If $X \in L^2(\Omega, \mathcal{A}, P)$ we define

$$\text{Var}(X) = E((X - E(X))^2)$$

$$\text{Cov}(X, Y) = E((X - E(X)) \cdot (Y - E(Y))).$$

Markov and Chebyshev inequalities

Cauchy-Schwarz - inequality

$X, Y \in L^2(\Omega, \mathcal{A}, P)$. Then:

$$E(X \cdot Y)^2 \leq E(X^2) \cdot E(Y^2)$$

Markov inequality: $\varepsilon > 0$, $f: [0, \infty[\rightarrow [0, \infty[$,

f monotonically increasing. Then

$$P(|Y| > \varepsilon) \leq \frac{E(f(|Y|))}{f(\varepsilon)}$$

In particular,

$$P(|Y| > \varepsilon) \leq \frac{E(|Y|)}{\varepsilon}$$

Chebyshev inequality: $\varepsilon > 0$, $X \in L^2(\Omega, \mathcal{A}, P)$. Then:

$$P(|X - E(X)| > \varepsilon) \leq \frac{\text{Var}(X)}{\varepsilon^2}$$

key quantity in learning theory

Examples of probability distributions

Discrete distributions

• Uniform distr. on $\{1, \dots, n\}$: $P(\{i\}) = \frac{1}{n}$

• Binomial distribution on $\{0, \dots, n\}$

Toss a coin n times, independently, each time with probability p of observing head. ~~Denote~~ Denote head = 1, tail 0,

$X := \# \text{ heads}$

$$P(X=k) := \binom{n}{k} p^k (1-p)^{n-k}$$

• Poisson distribution on \mathbb{N}

Parameter $\lambda > 0$

$$P(X=k) = \frac{\lambda^k e^{-\lambda}}{k!}$$

Intuition: number of incoming calls at a hotline.

Continuous distributions


Uniform distribution on $[a, b]$: constant density



Normal distribution on \mathbb{R}

Density: parameter μ (mean), σ (std. deviation)

$$f_{\mu, \sigma}(x) := \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$$



Notation: $N(\mu, \sigma^2)$

Some first properties:

- $X \sim N(\mu_1, \sigma_1^2)$, $Y \sim N(\mu_2, \sigma_2^2)$, X, Y independent.
Then $X+Y \sim N(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2)$

Normal distribution in higher dimension

$$X: \Omega \rightarrow \mathbb{R}^n, \quad X = \begin{pmatrix} X_1 \\ \vdots \\ X_n \end{pmatrix}, \quad \mu_i \in E(X_i), \quad \mu = \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_n \end{pmatrix}$$


$\Sigma \in \mathbb{R}^{n \times n}$ with $\Sigma_{ij} = \text{Cov}(X_i, X_j)$, called covariance matrix.

$$f_{\mu, \Sigma}(x) = \frac{1}{(2\pi)^{n/2} (\det \Sigma)^{1/2}} \exp\left(-\frac{1}{2} (x-\mu)^t \Sigma^{-1} (x-\mu)\right)$$

Notation: $N(\mu, \Sigma)$


Prop Σ is psd and symmetric.

Consequence: Σ has real-valued, non-negative eigP.



Contour lines of $f_{\mu, \Sigma}$

• X_1, \dots, X_n are independent $\Leftrightarrow \Sigma = \begin{pmatrix} \sigma_1^2 & & 0 \\ & \sigma_2^2 & \\ 0 & & \ddots \\ & & & \sigma_n^2 \end{pmatrix}$




- $X \sim N(\mu_1, \Sigma_1)$, $Y \sim N(\mu_2, \Sigma_2)$, independent, then
 $X + Y \sim N(\mu_1 + \mu_2, \Sigma_1 + \Sigma_2)$

Mixture of Gaussians

Consider $\pi_1, \pi_2, \dots, \pi_k$ with $0 \leq \pi_i \leq 1$ and $\sum \pi_i = 1$

Consider the following density:

$$f(x) = \sum_{i=1}^k \pi_i \cdot f_{\mu_i, \Sigma_i}(x)$$




Convergence of random variables

Consider rv $X_i: \Omega \rightarrow \mathbb{R}$, $i \in \mathbb{N}$, $X: \Omega \rightarrow \mathbb{R}$,
 (Ω, \mathcal{A}, P) a probability space.

(1) $(X_i)_{i \in \mathbb{N}}$ converges to X almost surely : \Leftrightarrow

$$P\left(\left\{\omega \in \Omega \mid \lim_{i \rightarrow \infty} X_i(\omega) = X(\omega)\right\}\right) = 1$$

Notation: $X_i \rightarrow X$ a.s.



(2) $(X_i)_{i \in \mathbb{N}}$ converges to X in probability : \Leftrightarrow

$$\forall \varepsilon > 0 \quad P\left(\left\{\omega \in \Omega \mid |X_i(\omega) - X(\omega)| > \varepsilon\right\}\right) \rightarrow 0$$

Let us check that these definitions make sense. We need to prove that the events in (1) and (2) are in fact in \mathcal{A} .

Case (1):

$$\lim_{i \rightarrow \infty} X_i(\omega) = X(\omega)$$

$$\Leftrightarrow \forall k \in \mathbb{N} \exists N \in \mathbb{N} \forall n > N: |X_n(\omega) - X(\omega)| < \frac{1}{k}$$

So we get:

$$\{\omega \mid X_n(\omega) \rightarrow X(\omega)\} =$$

$$= \underbrace{\bigcap_{k \in \mathbb{N}} \bigcup_{N \in \mathbb{N}} \bigcap_{n \geq N}}_{\text{countable unions and intersections}} \left\{ \omega \mid |X_n(\omega) - X(\omega)| < \frac{1}{k} \right\} \in \mathcal{A}$$

countable unions
and intersections

X_n, X are measurable \Rightarrow
 $|X_n - X|$ is measurable

so $\{\dots\} \in \mathcal{A}$


$$(3) \quad X_n \rightarrow X \text{ in } L^p \text{ ("in the } p\text{-th mean")} \quad : \Leftrightarrow$$

$$X_n, X \in L^p \text{ and } \|X_n - X\|_p \rightarrow 0.$$

(4) Let $M^1(\mathbb{R}^n)$ be the set of all probability measures on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$. Assume $(\mu_n) \subset M^1(\mathbb{R}^n)$, $\mu \in M^1(\mathbb{R}^n)$.
 $C_b(\mathbb{R}^n) :=$ space of bounded continuous functions.

$$\mu_n \rightarrow \mu \text{ weakly} \quad : \Leftrightarrow$$

$$\forall f \in C_b(\mathbb{R}^n) : \int f d\mu_n \rightarrow \int f d\mu$$



Excursion:

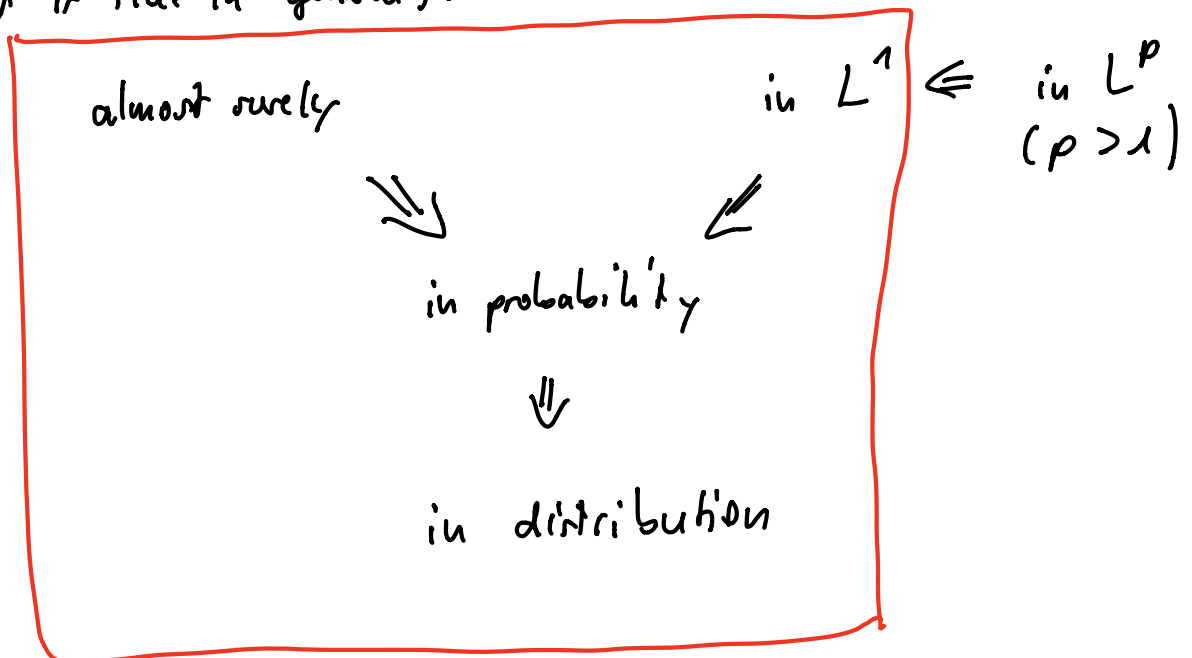
In functional analysis, a sequence $(x_n)_n$ in a Banach space B converges weakly if for all bounded lin. functionals f , we have that $f(x_n) \rightarrow f(x)$. (i.e. for all $f \in B'$).

Space $M^1(\mathbb{R}^n)$ itself is not a Banach space, but $C M(\mathbb{R}^n)$, space of all bounded measures.

The dual space of $M(\mathbb{R}^n)$ is $C_b(\mathbb{R}^n)$.


(5) $X_i, X : (\Omega, \mathcal{A}, P) \rightarrow \mathbb{R}^n$. The sequence X_n converges in distribution to X : \Leftrightarrow
the distributions P_{X_n} converge to P_X weakly.

We have the following implications (but none of the missing implications is true in general):



Example (convergence a.s., in prob., but not in L^1)

$$X_n : \mathbb{R} \rightarrow \mathbb{R}, \quad X_n(\omega) = \begin{cases} n & \text{for } 0 \leq x \leq \frac{1}{n} \\ 0 & \text{otherwise} \end{cases}$$



$$\forall x > 0: X_n(x) \rightarrow 0.$$


Can formally see: a.s., in prob.

But: no convergence in L^1 .


Example (convergence in prob., in L^1 , but not a.s.)

"sliding blocks"



$$f_1 = \mathbb{1}_{[0, 1]}$$



$$f_2 = \mathbb{1}_{[0, 1/2]}, \quad f_3 = \mathbb{1}_{[1/2, 1]}$$



$$f_4 = \mathbb{1}_{[0, 1/3]}, \quad f_5 = \mathbb{1}_{[1/3, 2/3]}, \quad f_6 = \mathbb{1}_{[2/3, 1]}$$



Example (Conv. in distribution, but not in prob.)

• $X_n : [0, 1] \rightarrow \mathbb{R}$, $X_1 = X_2 = \dots = \mathbb{1}_{[0, \frac{1}{2}]}$

• $X = \mathbb{1}_{[\frac{1}{2}, 1]}$

Obviously $X_n \not\rightarrow X$ in prob., but:

$$P_{X_1} = \frac{1}{2}(\delta_0 + \delta_1) = P_{X_2} = P_{X_3} = \dots = P_X$$

so $X_n \rightarrow X$ in distribution.

Theorem of Borel - Cantelli

(Ω, \mathcal{A}, P) prob. space, $(A_n)_n$ sequence of events in \mathcal{A} .

$$P(A_n \text{ infinitely often}) := P(A_n \text{ i.o.})$$

$$= P(\{\omega \in \Omega \mid \omega \in A_n \text{ for infinitely many } n\})$$

Proposition: X_n, X r.v. on (Ω, \mathcal{A}, P) .

$$X_n \rightarrow X \text{ a.s.} \Leftrightarrow$$

$$\forall \varepsilon > 0 : P(\{|X_n - X| > \varepsilon \text{ inf. often}\}) = 0$$

Proof intuition:

$$\{\lim X_n = X\}$$

$$= \{\forall k : |X_n - X| > \frac{1}{k} \text{ at most finitely often}\}$$

$$= \bigcap_{k \in \mathbb{N}} \{|X_n - X| > \frac{1}{k} \text{ at most fin. often}\}$$

$$\left(\bigcup_{k \in \mathbb{N}} \{|X_n - X| > \frac{1}{k} \text{ inf. often}\} \right)^{\text{complement}}$$

Theorem: Consider a sequence of events $(A_n)_n \subset \mathcal{A}$.

(1) If $\sum_{n=1}^{\infty} P(A_n) < \infty$, then $P(A_n \text{ i.o.}) = 0$.

(2) If $\sum_{n=1}^{\infty} P(A_n) = \infty$, and if $(A_n)_n$ are independent,
then $P(A_n \text{ i.o.}) = 1$.

Application in learning theory:

Assume that $P(|X_n - x| > \frac{1}{n}) < \delta_n$, and

assume that $\sum_{n=1}^{\infty} \delta_n < \infty$. Then you can use

Borel - Cantelli to prove that

$$P(|X_n - x| > \frac{1}{n} \text{ i.o.}) = 0,$$

thus $X_n \rightarrow x$ a.s.

Limit Theorems: LLN and CLT

Strong law of large numbers

$X_n : (\Omega, \mathcal{A}, P) \rightarrow \mathbb{R}$ iid (independently distributed and independent). Assume the mean $\mu := E(X_1) < \infty$, and $\text{Var}(X_1) =: \sigma^2 < \infty$. Then:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n X_i = \mu \quad \text{a.s. and in } L^2.$$

Remarks:

- Many versions of this theorem exist. (slightly relaxing iid)
- "Strong law" \Leftrightarrow convergence a.s.
- "Weak law" \Leftrightarrow convergence in probability

Central Limit Theorem

$(X_i)_{i \in \mathbb{N}}$ iid rv with mean μ , variance $\sigma^2 < \infty$.

Consider the rv $S_n := \sum_{i=1}^n X_i$. We normalize it to



$$Y_n := \frac{S_n - n\mu}{\sqrt{n}\sigma}$$

(which has mean 0 and standard dev. 1).

Then $Y_n \rightarrow Y$ in distribution where $Y \sim \mathcal{N}(0, 1)$.

(Illustration: X_i coin, head $\hat{=} 1$, tail $\hat{=} 0$)


$$S_n = \sum X_i \in [0, n]$$



Concentration inequalities

Motivation: random projections

$\dots \mathbb{R}^d$, d large
 \dots want to project in \mathbb{R}^l , l "small"



Theorem of Johnson-Lindenstrauss:

Can guarantee (for certain parameters ϵ, k)

$$(1-\epsilon) \|x_i - x_j\|_{\mathbb{R}^d} \leq \|\pi(x_i) - \pi(x_j)\|_{\mathbb{R}^l}$$

$$\leq (1+\epsilon) \|x_i - x_j\|_{\mathbb{R}^d}$$

Construction / proof steps:

(1) Assume you know $\|x_i - x_j\|_{\mathbb{R}^d} = 1$.

Compute $E(\|\pi(x_i) - \pi(x_j)\|_{\mathbb{R}^l})$, "easy".

(2) $P(\|\pi(x_i) - \pi(x_j)\| - E(\dots) > t) ?$


Hoeffding inequality

Theorem (Hoeffding): $(\Omega, \mathcal{F}, P) \rightarrow (\mathbb{R}, \mathcal{B})$
 x_1, \dots, x_n i.i.d. r.v., independent,

assume that $x_i \in [a_i, b_i]$ a.s. for $i=1, \dots, n$.

Let $S_n := \sum_{i=1}^n (x_i - E(x_i))$. Then for any $t > 0$,

$$P(S_n \geq t) \leq \exp\left(-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right).$$



Application of Hoeffding: SLLN

Prop $(x_i)_{i \in \mathbb{N}}$ i.i.d. r.v., $a \leq x_i \leq b$, let X have the same distribution as the x_i .

Then: $\frac{1}{n} \sum_{i=1}^n x_i \rightarrow E(X)$ a.s.

Proof Hoeffding \Rightarrow

$$\bullet P\left(\frac{1}{n} \sum x_i - E(X) > t\right) \leq \exp\left(-\frac{2nt^2}{(b-a)^2}\right)$$

$$\bullet P\left(\frac{1}{n} \sum x_i - E(X) < -t\right)$$

$$= P\left(\frac{1}{n} \sum (-x_i) - E(-x) > t\right) \leq \exp\left(-\frac{2nt^2}{(b-a)^2}\right)$$

Combined we get

$$P\left(\left|\frac{1}{n} \sum x_i - E(x)\right| > t\right) \leq 2 \exp\left(-\frac{2nt^2}{(b-a)^2}\right).$$

Now want to apply Borel-Cantelli to get a.s. convergence:

$$Z_n := \frac{1}{n} \sum_{i=1}^n x_i$$

$$\sum_{n=0}^{\infty} P(Z_n - E(x) > t) \leq 2 \cdot \underbrace{\sum_{n=0}^{\infty} \exp\left(-\frac{2nt^2}{(b-a)^2}\right)}_{=: \text{sum}} \stackrel{!}{\leq} \infty$$

⊛ Substitute: $r := \exp\left(-\frac{2t^2}{(b-a)^2}\right) \in [0, 1[$

observe: $\exp\left(-\frac{2nt^2}{(b-a)^2}\right) = r^n$

$$\text{sum} = 2 \sum_{n=0}^{\infty} r^n = 2 \cdot \frac{1}{1-r} < \infty.$$

Now Borel-Cantelli gives almost sure convergence. \square

Remark: Hoeffding is tight (cannot be improved without further assumptions). For fair coin tosses it is tight.

But: not tight if coin is biased \leadsto need other inequalities

Bernstein inequality

Theorem (Bernstein): x_1, \dots, x_n independent with 0 mean,
 $|x_i| < 1$ a.s. Let $\sigma^2 := \frac{1}{n} \sum_{i=1}^n \text{Var}(x_i)$. Then
for all $t > 0$,

$$P\left(\frac{1}{n} \sum_{i=1}^n x_i > t\right) \leq \exp\left(-\frac{nt^2}{2(\sigma^2 + t/3)}\right)$$

Concentration inequality for functions with bounded differences

Consider a function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ (or more generally,
 $f: \mathcal{X}^n \rightarrow \mathbb{R}$ for some "arbitrary" space \mathcal{X}).

We say that f has the bounded differences property if
there exist constants c_1, \dots, c_n such that

$$\textcircled{*} \quad \sup_{\substack{x_1, \dots, x_n \in \mathcal{X} \\ \tilde{x}_i \in \mathcal{X}}} \left| f(x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_n) \right. \\ \left. - f(x_1, \dots, x_{i-1}, \tilde{x}_i, x_{i+1}, \dots, x_n) \right| \leq c_i$$

Example: $f(x_1, \dots, x_n) = \sum_{i=1}^n x_i$, and $a \leq x_i \leq b \forall i$, then
 f satisfies $\textcircled{*}$ with $c_i = b - a$.


Glivenko - Cantelli Theorem

F cdf : $F(a) = P(X \leq a)$

$X_1, \dots, X_n \sim F$, iid

$F_n : \mathbb{R} \rightarrow [0, 1]$

$$F_n(a) := \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{X_i \leq a\}}$$



Now fix one particular $a_0 \in \mathbb{R}$.

$F_n(a_0) \rightarrow F(a_0)$ by the law of large numbers.

Because $\mathbb{1}_{\{X_i \leq a_0\}}$ is a Binomial rv with

$$p = P(X_i \leq a_0).$$


So it is clear that $F_n \rightarrow F$ ^{a.s.} pointwise (i.e. $\forall a_0$)

Now let's look at uniform convergence.

Theorem X_1, \dots, X_n iid random variables with cdf F .

Let F_n be the empirical cdf induced by the sample. Then:

$$P\left(\sup_{a \in \mathbb{R}} |F_n(a) - F(a)| > \varepsilon\right) \leq \varepsilon \cdot (n+1) \cdot \exp\left(-\frac{n\varepsilon^2}{32}\right).$$



In particular, $\sup |F_n - F| \rightarrow 0$ a.s.,

i.e. $F_n \rightarrow F$ uniformly a.s.

Proof

Observe: $LLN \Rightarrow P(|F_n(a_0) - F(a_0)| > \epsilon) \rightarrow 0$
for any fixed a_0 .

Problem: need to look at

$$P\left(\sup_{a \in \mathbb{R}} |F_n(a) - F(a)| > \epsilon\right)$$


difficult because \mathbb{R} is uncountable

If we take a supremum over a finite set,
it is easier:

$$\begin{aligned}
P\left(\max_{i=1, \dots, n} |U_i| > \epsilon\right) &= \\
&= P(|U_1| > \epsilon \text{ or } |U_2| > \epsilon \text{ or } \dots \text{ or } |U_n| > \epsilon) \\
&\leq \sum_{i=1}^n P(|U_i| > \epsilon)
\end{aligned}$$

Trick of the proof: convert $\sup_{a \in \mathbb{R}}$ to something "finite".

How could we achieve this?



|red - green|

$$|\text{red} - \text{blue}| \leq 2 |\text{green} - \text{blue}|$$

Step 1: Symmetrization by ghost sample

Assume $X_1', \dots, X_n' \sim F$ independently ("ghost sample"),

Denote by F_n' the empirical cdf induced by ghost sample

Now it is easy to prove:

$$P\left(\sup_a | \underline{F_n(a)} - \underline{F(a)} | > \varepsilon \right)$$

$$\leq 2 P\left(\sup_a | \underline{F_n(a)} - \underline{F_n'(a)} | > \frac{\varepsilon}{2} \right)$$

Step 2: Want to split this in two terms

$$| F_n(a) - F_n'(a) | = \left| \frac{1}{n} \sum_{i=1}^n \left(\mathbb{1}_{\{X_i \leq a\}} - \mathbb{1}_{\{X_i' \leq a\}} \right) \right|$$

Introduce Rademacher random variables $\sigma_1, \dots, \sigma_n$:

$$\sigma_i(\{-1\}) = \sigma_i(\{1\}) = 1/2.$$

Distribution of \otimes is the same as the distr. of the following:

$$\left| \frac{1}{n} \sum_{i=1}^n \sigma_i \left(\mathbb{1}_{\{X_i \leq a\}} - \mathbb{1}_{\{X_i' \leq a\}} \right) \right| = \otimes \otimes$$

Now we have:

$$2P\left(\sup_a |F_n(a) - F_n'(a)| > \frac{\varepsilon}{2}\right)$$

$$= 2P\left(\sup_a \left| \frac{1}{n} \sum \sigma_i (\mathbb{1}_{X_i \leq a} - \mathbb{1}_{X_i' \leq a}) \right| > \frac{\varepsilon}{2}\right)$$

$$\leq 2P\left(\sup_a \left| \underbrace{\frac{1}{n} \sum \sigma_i \mathbb{1}_{X_i \leq a}}_u \right| > \frac{\varepsilon}{4}\right) + 2P\left(\sup_a \left| \underbrace{\frac{1}{n} \sum \sigma_i \mathbb{1}_{X_i' \leq a}}_v \right| > \frac{\varepsilon}{4}\right)$$

Observe:

$$P(|u-v| > \frac{\varepsilon}{2}) \leq P(|u| > \frac{\varepsilon}{4} \text{ or } |v| > \frac{\varepsilon}{4})$$

↑ right side is necessary for left side

$$= 4 \cdot P\left(\sup_a \left| \frac{1}{n} \sum \sigma_i \mathbb{1}_{\{X_i \leq a\}} \right| > \frac{\varepsilon}{4}\right)$$

Step 3

Exploit "finite structure":

Fix x_1, \dots, x_n (\equiv condition on x_1, \dots, x_n)

We look at $\mathbb{1}_{X_i \leq a}$ 

The var $\mathbb{1}_{\{x_1 \leq a\}}, \dots, \mathbb{1}_{\{x_n \leq a\}}$ For fixed a can only

have $n+1$ realizations

$$P\left(\sup_a \frac{1}{n} \left| \sum \sigma_i \mathbb{1}_{X_i \leq a} \right| > \frac{\varepsilon}{4} \mid x_1, \dots, x_n\right) \leq$$

$$\leq (n+1) \sup_a \underbrace{P\left(\frac{1}{n} \left| \sum \sigma_i \mathbb{1}_{\{X_i \leq a\}} \right| > \frac{\varepsilon}{4} \mid x_1, \dots, x_n\right)}$$

use Hoeffding (#)

Step 4 apply
Hoeffding to (4)

W.a.:

$$P\left(\frac{1}{n} \sum \sigma_i \mathbb{1}_{x_i \leq a} > \frac{\varepsilon}{4} \mid x_1, \dots, x_n\right)$$

$$\leq 2 \exp\left(-\frac{n \varepsilon^2}{32}\right)$$

Combining everything gives the theorem.

Product space, joint distributions

Considers two measurable spaces $(\Omega_1, \mathcal{A}_1)$, $(\Omega_2, \mathcal{A}_2)$.

Define the product space $(\Omega_1 \times \Omega_2, \mathcal{A}_1 \otimes \mathcal{A}_2)$ with

$$\Omega_1 \times \Omega_2 = \{(\omega_1, \omega_2) \mid \omega_1 \in \Omega_1, \omega_2 \in \Omega_2\}$$

$$\mathcal{A}_1 \otimes \mathcal{A}_2 = \{A_1 \times A_2 \mid A_1 \in \mathcal{A}_1, A_2 \in \mathcal{A}_2\}.$$

Considers two rvs $X_1: (\Omega, \mathcal{A}, P) \rightarrow (\Omega_1, \mathcal{A}_1)$
 $X_2: (\Omega, \mathcal{A}, P) \rightarrow (\Omega_2, \mathcal{A}_2)$.

$$X := (X_1, X_2): (\Omega, \mathcal{A}, P) \rightarrow (\Omega_1 \times \Omega_2, \mathcal{A}_1 \otimes \mathcal{A}_2)$$

$$(X_1, X_2)(\omega) = (X_1(\omega), X_2(\omega)).$$

The distribution $P_{(X_1, X_2)}$ on $(\Omega_1 \times \Omega_2, \mathcal{A}_1 \otimes \mathcal{A}_2)$ is called

the joint distribution of X_1 and X_2 .

Example in ML: (X, Y) where X is the input data, Y is the label

Product measure: $(\Omega_1, \mathcal{A}_1, P_1)$, $(\Omega_2, \mathcal{A}_2, P_2)$ two

prob. spaces. We define the product measure $P_1 \otimes P_2$ on

the product space $(\Omega_1 \times \Omega_2, \mathcal{A}_1 \otimes \mathcal{A}_2)$ as

$$(P_1 \otimes P_2)(A_1 \times A_2) := P_1(A_1) \cdot P_2(A_2).$$

Theorem Two rvs X_1, X_2 are independent if and only if their joint distribution coincides with the product distribution:

$$P_{(X_1, X_2)} = P_1 \otimes P_2.$$

Marginal distribution

Consider the joint distribution $P_{(X_1, X_2)}$ of two rvs $X := (X_1, X_2)$. The marginal distribution of X wrt X_1 is the original distribution of X_1 on $(\Omega_1, \mathcal{A}_1)$, namely P_{X_1} . Similarly for P_{X_2} .

Example in the discrete case:

$Y \setminus X$	x_1	x_2	x_3	Σ
Y_1	p_1	p_2	p_3	$p_1 + p_2 + p_3 = P(Y = Y_1)$
Y_2	p_4	p_5	p_6	$p_4 + p_5 + p_6 = P(Y = Y_2)$
	$p_1 + p_4$ $= P(X = x_1)$...		$\hat{=}$ marginal distribution wrt Y .

marginal wrt X

Marginal distributions in case of densities

$X, Y: (\Omega, \mathcal{A}, P) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$, $Z := (X, Y)$. Assume that the joint distribution of Z has a density f on \mathbb{R}^2 . Then the following statements hold:

(1) Both X and Y have densities on $(\mathbb{R}, \mathcal{D}(\mathbb{R}))$ given by

$$f_X(x) = \int_{-\infty}^{\infty} f(x,y) dy$$

joint density (pointing to $f(x,y)$)
sum over Y (pointing to dy)

$$f_Y(y) = \int_{-\infty}^{\infty} f(x,y) dx$$

(2) X and Y are independent iff

$$f(x,y) = f_X(x) \cdot f_Y(y) \quad \text{a.s.}$$

Mixed cases

For example, consider X a continuous rv with density and Y a discrete rv.

See, $X = \text{income} \in \mathbb{R}$

$Y = \text{"yes" or "no"}$, discrete


Special case: marginals of multivariate normal distributions

2 dim Consider a 2-dim normal rv $X = \begin{pmatrix} X_1 \\ X_2 \end{pmatrix}$ with mean

$$\mu = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} \in \mathbb{R}^2 \quad \text{and} \quad \text{cov. } \Sigma = \begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{21} & \sigma_2^2 \end{pmatrix}.$$

Then the marginal distribution of X w.r.t X_1 is again

a normal distribution with mean μ_1 and var σ_1^2 .



n-dim

$x = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \in \mathbb{R}^n$. Group the variables: $\begin{matrix} x_1 \\ \vdots \\ x_k \end{matrix} \} \tilde{x} \in \mathbb{R}^k$
 $\begin{matrix} x_{k+1} \\ \vdots \\ x_n \end{matrix} \} x^\# \in \mathbb{R}^{n-k}$

Want to look at the marginal of x wrt \tilde{x} .

$\mu = \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_n \end{pmatrix}$ mean, $\tilde{\mu} := \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_k \end{pmatrix}$, $\mu^\# = \begin{pmatrix} \mu_{k+1} \\ \vdots \\ \mu_n \end{pmatrix}$

$\Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix} \} k$

$\underbrace{\hspace{2cm}}_k$

Now the marginal of x wrt \tilde{x} is a normal distr. on \mathbb{R}^k with mean $\tilde{\mu}$ and cov. Σ_{11} .

Conditional distributions

Discrete case:

Know conditional probabilities: $P(A | B)$

defined for events $A, B \in \mathcal{A}$, and $P(B) > 0$.

Let $X, Y: (\Omega, \mathcal{A}, P) \rightarrow \mathbb{R}$ be discrete rv, $y \in \mathbb{R}$ such that $P(Y=y) > 0$. Then we can define the conditional probability

measure $P_{X|Y=y}: A \mapsto P(X \in A | Y=y)$.

This is a probability measure.

For general rv this is surprisingly complicated!

\leadsto "regular conditional probabilities" \leadsto skipped

Conditional distributions in case of densities

Assume $Z := (X, Y)$ has a joint density $f: \mathbb{R}^2 \rightarrow \mathbb{R}$,
and marginal densities $f_X, f_Y: \mathbb{R} \rightarrow \mathbb{R}$. Then the function

$$f_{X|Y=y}(x) := \frac{f(x, y)}{f_Y(y)}$$

is then also a density on \mathbb{R} , called the conditional density of X given $Y=y$.

Example: normal distributions

$$\mu = \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_n \end{pmatrix} \begin{matrix} \} \tilde{\mu} \\ \} \mu^\# \end{matrix} \quad \Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}$$

If $X = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \sim N(\mu, \Sigma)$, then the conditional distributions

of $\tilde{X} = \begin{pmatrix} x_1 \\ \vdots \\ x_k \end{pmatrix}$ wrt $x^\# = \begin{pmatrix} x_{k+1} \\ \vdots \\ x_n \end{pmatrix}$ is given by

$$P_{\tilde{X} | X^\#} \sim N \left(\mu_1 + \Sigma_{12} \Sigma_{22}^{-1} (x^\# - \tilde{\mu}), \right. \\ \left. \Sigma_{11} - \Sigma_{12} \Sigma_{22}^{-1} \Sigma_{21} \right).$$



Conditional expectation

Def (discrete case) $X, Y: (\Omega, \mathcal{A}, P) \rightarrow \mathbb{R}$

assume X takes finitely (countably) many values

$x_1, \dots, x_n \in \mathbb{R}$, Y takes finitely (countably) many values

$y_1, \dots, y_m \in \mathbb{R}$, always with a positive probability.

$$E(Y | X = x_i) := \sum_{j=1}^m y_j \underbrace{P(Y = y_j | X = x_i)}_{\text{well defined}}$$

Example: two dice, $X =$ first one, $Y =$ second one, independent

$$E(\text{sum} | X = 1) = \sum_{i=1}^{12} i \cdot P(\text{sum} = i | X = 1)$$

$$= \sum_{k=1}^6 (1+k) \cdot P(Y = k | X = 1)$$

$$= \sum_{k=1}^6 (1+k) P(Y = k) = \sum_{k=1}^6 (1+k) \cdot \frac{1}{6} = 4.5$$

So far we defined $E(Y | X = x_i)$, but often we want to consider the "function" $E(Y | X)(\omega)$. This is a rv:

$$E(Y | X): (\Omega, \mathcal{A}, P) \rightarrow (\mathbb{R}, \mathcal{B}).$$

Leads to the following:

Def (direct case) X, Y as before. Then the conditional expectation is defined as follows:

$$E(Y|X) := f(X) \quad \text{with}$$

$$f(x) = \begin{cases} E(Y|X=x) & \text{if } P(X=x) > 0 \\ \text{arbitrary, say } 0 & \text{otherwise} \end{cases}$$


⚠ $E(Y|X)$ is only defined a.s.

Now we want to move to the more general case.

Sketch: X continuous rv

Y discrete rv $\leadsto Y_1, \dots, Y_5$

Want to look at $E(X|Y)$



Want to "define" $E(X|Y) := \sum_{i=1}^5 E(X|Y=y_i) \cdot \mathbb{1}_{B_i}(\omega)$

But need to make sure that it is measurable wrt $\sigma(Y)$.
(“the bür”)

rv

Def (Conditional expectation on L_1)

Consider rv $X: (\Omega, \mathcal{F}_0, P) \rightarrow \mathbb{R}$, $X \in L_1(\Omega, \mathcal{F}_0, P)$.

Let \mathcal{F} be a sub- σ -algebra of \mathcal{F}_0 . (Intuition: \mathcal{F}_0 will be the σ -alg. generated by the variable Y we want to condition on).


We now define the cond. exp. of X given \mathcal{F}

$E(X | \mathcal{F})$ as any random variable Z that satisfies

(1) Z is measurable wrt \mathcal{F}

(2) For all $A \in \mathcal{F}$ we have

$$\int_A X dP = \int_A Z dP$$



• Existence of $E(X | \mathcal{F})$ is not clear ^A a priori, it needs to be proved.

• $E(X | \mathcal{F}) := E(X | \sigma(Y))$

Examples (extreme cases)

• $X = Y$. Then $E(X | \mathcal{F}) = X$ (a.s.)

• $X \perp\!\!\!\perp Y$. $E(X | \mathcal{F}) = E(X)$ (a.s.)

Case of joint densities

$X, Z: \Omega \rightarrow \mathbb{R}$ have a joint density $f(x, z)$.

Let $g: \mathbb{R} \rightarrow \mathbb{R}$ bounded, set $Y := g(Z)$. Assume we want to compute $E(Y|X) = E(\underbrace{g(Z)}_Y | X)$.

Recall X has density $f_X(x) = \int f(x, z) dz$.

The conditional density of Z given $X=x$ is

$$f_{X=x}(z) = \frac{f(x, z)}{f_X(x)} \quad (\text{if } f_X(x) \neq 0)$$

Now consider $h(x) := \int \underbrace{g(z)}_Y f_{X=x}(z) dz$, now define

$$E(Y|X) = h(X).$$