STAT 712 fa 2022 Lec 11 slides

Sundry bivariate nuggets, some inequalities, hierarchical models

Karl B. Gregory

University of South Carolina

These slides are an instructional aid; their sole purpose is to display, during the lecture, definitions, plots, results, etc. which take too much time to write by hand on the blackboard. They are not intended to explain or expound on any material.

1. **Covariance and correlation**

2. Cauchy-Schwarz, Hölder's, Minkowski's, and Jensen's inequalities

3. Hierarchical models

Covariance

The covariance between two rvs X and Y is defined as

$$\text{Cov}(X, Y) = \mathbb{E}(X - \mu_X)(Y - \mu_Y) =: \sigma_{XY},$$

where $\mu_X = \mathbb{E}X$ and $\mu_Y = \mathbb{E}Y$.

Useful expression: $\text{Cov}(X, Y) = \mathbb{E}XY - \mathbb{E}X\mathbb{E}Y$

👍

Exercise: Derive the useful expression for computing covariances.

Correlation

The correlation between two rvs X and Y is defined as

$$\text{corr}(X, Y) = \frac{\text{Cov}(X, Y)}{\sqrt{\text{Var}\,X}\sqrt{\text{Var}\,Y}} = \frac{\sigma_{XY}}{\sigma_X \sigma_Y} =: \rho_{XY},$$

where $\sigma_X = \sqrt{\text{Var}\,X}$ and $\sigma_Y = \sqrt{\text{Var}\,Y}$.

Theorem (Correlation between minus 1 and 1, cf. Thm 4.5.7 in CB)

For any rvs X and Y,

1. $-1 \leq \text{corr}(X, Y) \leq 1$
2. $\text{corr}(X, Y) = \pm 1$ *iff there exist $a \neq 0$ and b such that $P(Y = aX + b) = 1$.*

We will prove the first part of this result later.

Exercise: Let (X, Y) be a pair of rvs with joint pdf given by

$$f(x, y) = \frac{1}{8}(x + y) \cdot \mathbf{1}(0 < x < 2, 0 < y < 2).$$

1. Find $\text{Cov}(X, Y)$.
2. Find $\text{corr}(X, Y)$.

Covariance and correlation of linearly transformed rvs

For any two rvs X and Y and constants $a, b, c, d \in \mathbb{R}$, we have

$$\text{Cov}(aX + b, cY + d) = ac \cdot \text{Cov}(X, Y)$$
$$\text{corr}(aX + b, cY + d) = \text{sign}(ac) \cdot \text{corr}(X, Y).$$

Exercise: Prove the result.

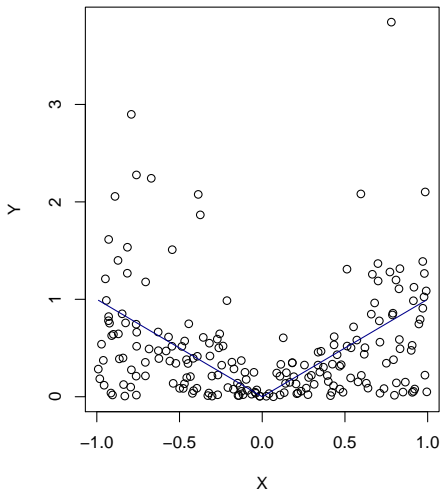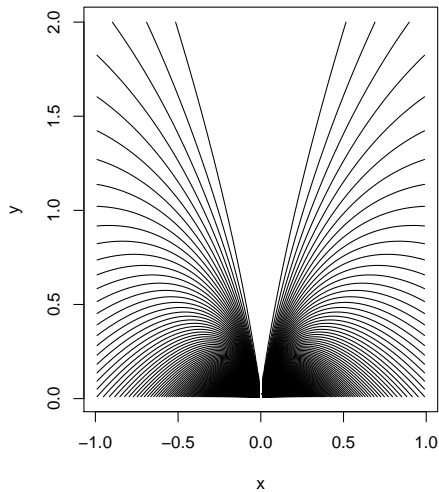Theorem (independence implies covariance equal to zero)

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

If $\text{Cov}(X, Y) = 0$, it does not mean that X and Y are independent!

Exercise: Let (X, Y) be a pair of rvs with joint pdf given by

$$f(x, y) = \frac{1}{2|x|} e^{-y/|x|} \mathbf{1}(x \in (-1, 1) \setminus \{0\}, y > 0).$$

① Check whether X and Y are independent.
② Compute $\text{Cov}(X, Y)$.

Theorem (Variance/Covariance of linear combinations of rvs)

Let $X_1, \ldots, X_n$ and $Y_1, \ldots, Y_m$ be rvs and let $a_1, \ldots, a_n, b_1, \ldots, b_m \in \mathbb{R}$. Then

1. $\mathrm{Cov}(\sum_{i=1}^n a_i X_i, \sum_{j=1}^m b_j Y_j) = \sum_{i=1}^n \sum_{j=1}^m a_i b_j \, \mathrm{Cov}(X_i, Y_j)$.
2. $\mathrm{Var}\left(\sum_{i=1}^n a_i X_i\right) = \sum_{i=1}^n a_i^2 \, \mathrm{Var}(X_i) + 2 \sum_{i<j} a_i a_j \, \mathrm{Cov}(X_i, X_j)$.

Simple cases of the above are

1. $\mathrm{Cov}(aX + bY, cU + dV) =$
   $ac \, \mathrm{Cov}(X, U) + ad \, \mathrm{Cov}(X, V) + ac \, \mathrm{Cov}(Y, U) + bd \, \mathrm{Cov}(Y, V)$.
2. $\mathrm{Var}(aX + bY) = a^2 \, \mathrm{Var}\, X + b^2 \, \mathrm{Var}\, Y + 2ab \, \mathrm{Cov}(X, Y)$.

Exercise: Prove the general results.

Exercise: Let $Z_1, \ldots, Z_n$ have unit variance and suppose

$$\text{corr}(Z_i, Z_j) = \rho \in (-1, 1) \text{ for } i \neq j.$$

Find $\text{Var} \, \bar{Z}$, where $\bar{Z} = n^{-1} \sum_{i=1}^{n} Z_i$.

Exercise: Let $Y_1, \ldots, Y_n$ be independent rvs such that

$$Y_i \sim \text{Normal}(\mu, \sigma_i^2), \quad i = 1, \ldots, n$$

and let

$$\bar{Y} = \frac{1}{n} \sum_{i=1}^{n} Y_i \quad \text{and} \quad \tilde{Y} = \frac{\sum_{i=1}^{n} \sigma_i^{-2} Y_i}{\sum_{j=1}^{n} \sigma_j^{-2}}.$$

1. Find $\mathbb{E}\bar{Y}$.
2. Find $\text{Var}\,\bar{Y}$.
3. Find $\mathbb{E}\tilde{Y}$.
4. Find $\text{Var}\,\tilde{Y}$.
5. Consider case $\sigma_1^2 = \cdots = \sigma_n^2$.
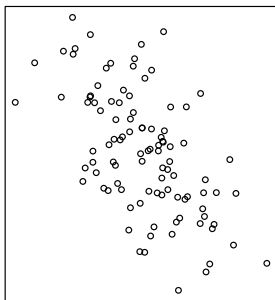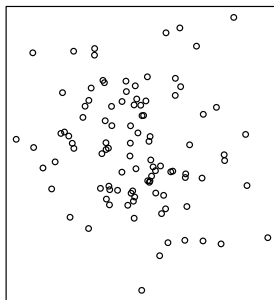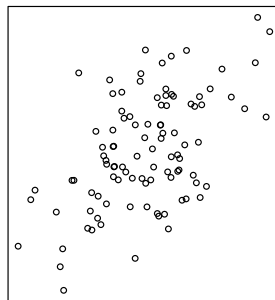
Bivariate Normal distribution

The rvs (X, Y) have the bivariate Normal distribution if they have joint pdf

$$
f(x, y; \mu_X, \mu_Y, \sigma_X, \sigma_Y, \rho) = \frac{1}{2\pi} \frac{1}{\sigma_X \sigma_Y \sqrt{1 - \rho^2}}
$$
$$
\times \exp\left[-\frac{1}{2(1 - \rho^2)} \left(\left[\frac{X - \mu_X}{\sigma_X}\right] - 2\rho \left[\frac{X - \mu_X}{\sigma_X}\right] \left[\frac{Y - \mu_Y}{\sigma_Y}\right] + \left[\frac{Y - \mu_Y}{\sigma_Y}\right]^2 \right) \right],
$$

where

- μ_X and μ_Y are mean of X and Y.
- σ_X^2 and σ_Y^2 are variance of X and Y.
- ρ is corr(X, Y).

Exercise: Show biv-Normal (X, Y) are independent iff corr$(X, Y) = 0$.

rho = −0.5   rho = 0   rho = 0.5

Exercise: Show that with $\rho = \sigma_{XY}/(\sigma_X \sigma_Y)$, we can write the biv-Normal pdf as

$$f(x,y) = \frac{\exp\left[-\frac{1}{2}\begin{pmatrix} X - \mu_X \\ Y - \mu_Y \end{pmatrix}^T \begin{pmatrix} \sigma_X^2 & \sigma_{XY} \\ \sigma_{XY} & \sigma_Y^2 \end{pmatrix}^{-1} \begin{pmatrix} X - \mu_X \\ Y - \mu_Y \end{pmatrix}\right]}{2\pi \cdot \begin{vmatrix} \sigma_X^2 & \sigma_{XY} \\ \sigma_{XY} & \sigma_Y^2 \end{vmatrix}^{1/2}}.$$

Multivariate Normal distribution

Let $\mathbf{X} = (X_1, \ldots, X_d)^T$ be a vector of rvs with joint pdf

$$f(\mathbf{x}) = \frac{1}{(2\pi)^{d/2}|\mathbf{\Sigma}|^{1/2}} \exp\left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right],$$

where $\boldsymbol{\mu} \in \mathbb{R}^d$ and $\mathbf{\Sigma}$ a symmetric positive definite $d \times d$ matrix.

Then we say $\mathbf{X}$ has the *multivariate Normal distribution* with mean vector $\boldsymbol{\mu}$ and covariance matrix $\mathbf{\Sigma}$, where $\mathbb{E}X_j = \mu_j$ and $\mathrm{Cov}(X_j, X_k) = \mathbf{\Sigma}_{jk}$, $1 \leq j, k \leq d$.

1 Covariance and correlation

2 Cauchy-Schwarz, Hölder's, Minkowski's, and Jensen's inequalities

3 Hierarchical models

Theorem (Cauchy-Schwarz Inequality)

For any rvs X and Y we have $|\mathbb{E}XY| \leq \mathbb{E}|XY| \leq \sqrt{\mathbb{E}X^2}\sqrt{\mathbb{E}Y^2}$.

Exercise:

1. Prove the inequality.
2. Use to prove $\text{corr}(X, Y) \in [-1, 1]$ for any rvs X, Y.

Theorem (Hölder's inequality)

For any two rvs X and Y and any $p, q \geq 1$ such that $1/p + 1/q = 1$, we have

$$|\mathbb{E}XY| \leq \mathbb{E}|XY| \leq (\mathbb{E}|X|^p)^{1/p}(\mathbb{E}|Y|^q)^{1/q}$$

The Cauchy-Schwarz is a special case of Hölder's with $p = q = 2$.

Use fact that for any $a > 0$, $b > 0$ and $p, q \geq 1$ such that $1/p + 1/q = 1$, we have

$$\frac{a^p}{p} + \frac{b^q}{q} \geq ab.$$

Theorem (Minkowski's inequality)

For any rvs X and Y and any $p \in [1, \infty)$ we have

$$(\mathbb{E}|X + Y|^p)^{1/p} \leq (\mathbb{E}|X|^p)^{1/p} + (\mathbb{E}|Y|^p)^{1/p}.$$

Exercise: Let Z_1 and Z_2 be standard Normals rvs.

1. Give an upper bound for $(\mathbb{E}|Z_1 - Z_2|^2)^{1/2}$ using Minkowski's inequality.
2. Give $(\mathbb{E}|Z_1 - Z_2|^2)^{1/2}$ exactly, letting $\rho = \text{corr}(Z_1, Z_2)$.

Theorem (Jensen's inequality)

For any rv X and any convex function g we have $g(\mathbb{E}X) \leq \mathbb{E}g(X)$.

Exercise: Prove Jensen's inequality. Recall that g is convex if

$$g(\lambda x + (1 - \lambda)x') \leq \lambda g(x) + (1 - \lambda)g(x') \quad \text{for all } \lambda \in (0, 1), \ x, x' \in \mathbb{R}.$$

We often use Jensen's to claim $e^{\mathbb{E}X} \leq \mathbb{E}e^X$ and $\log \mathbb{E}X \geq \mathbb{E}\log X$.

1 Covariance and correlation

2 Cauchy-Schwarz, Hölder's, Minkowski's, and Jensen's inequalities

3 Hierarchical models

- Often (X, Y) relation is most clearly described by a conditional and marginal:

$$f(y|x) = \frac{f(x, y)}{f_X(x)} \iff f(x, y) = f(y|x)f_X(x).$$

- A *hierarchical model* describes the joint dist. of an rv pair (X, Y) in the form

$$Y|X \sim \text{Some distribution depending on } X$$
$$X \sim \text{Some distribution}$$

- Can use to get interesting marginal distributions for Y, which take the form

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

Poisson-Binomial hierarchical model example

Let

$X = \#$ customers entering a store in a day

$Y = \#$ customers who make purchases

We might assume the following hierarchical model for Y:

$$Y|X \sim \text{Binomial}(X, p)$$
$$X \sim \text{Poisson}(\lambda).$$

Exercise: Find the following:

1. The joint pmf of (X, Y).
2. The marginal pmf of Y.
3. $\mathbb{E}Y$ and $\text{Var}\, Y$.

Theorem (iterated expectation and iterated variance)

For any random variables X and Y we have

- $\mathbb{E}Y = \mathbb{E}(\mathbb{E}[Y|X])$
- $\text{Var}\, Y = \mathbb{E}(\text{Var}[Y|X]) + \text{Var}(\mathbb{E}[Y|X])$

Exercise:

1. Prove the 1st result above.
2. Apply results to previous example.

Normal random-effects hierarchical model

Let an elementary-school pupil from the U.S. be chosen at random and let

$Y =$ score on a standardized test the selected pupil.

$A =$ average test score at the school of the selected pupil.

We might assume the following hierarchical model for Y:

$$Y|A \sim \text{Normal}(A, \sigma^2)$$
$$A \sim \text{Normal}(\mu_A, \sigma_A^2).$$

Exercise:

1. Find $\mathbb{E}Y$ and $\text{Var}\, Y$.
2. Find the marginal pdf of Y.

We can use hierarchical models to make interesting "mixture" distributions:

Mixture distribution induced by a hierarchical model

For a pair of rvs (X, Y), let p and f represent pmfs and pdfs, respectively. Then

- $Y|X \sim f(y|x)$ with $X \sim f_X(x)$ gives $f_Y(y) = \int_{\mathbb{R}} f(y|x) f_X(x) dx$.

- $Y|X \sim p(y|x)$ with $X \sim p_X(x)$ gives $p_Y(y) = \sum_{x \in \mathcal{X}} p(y|x) p_X(x) dx$.

- $Y|X \sim f(y|x)$ with $X \sim p_X(x)$ gives $f_Y(y) = \sum_{x \in \mathcal{X}} f(y|x) p_X(x) dx$

- $Y|X \sim p(y|x)$ with $X \sim f_X(x)$ gives $p_Y(y) = \int_{\mathbb{R}} p(y|x) f_X(x) dx$.

The marginal distributions of Y are called *mixture distributions*.

Beta-Binomial hierarchical model example

Let a spectator of a basketball game be chosen at random and let

$Y = \#$ freethrows made out of n attempts by the chosen spectator.

$P = $ free-throw success rate of the chosen spectator.

We might assume the following hierarchical model for Y:

$$Y|P \sim \text{Binomial}(n, P)$$
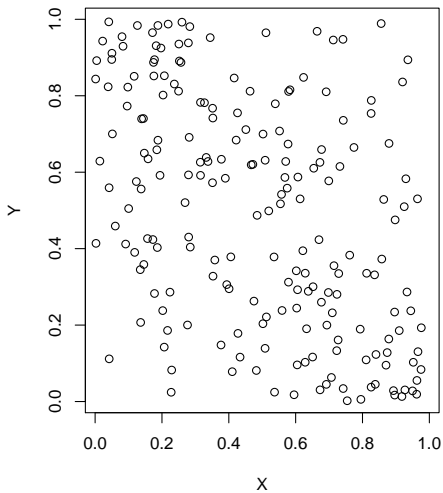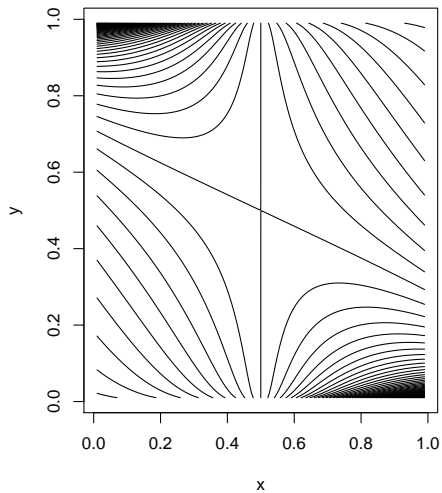$$P \sim \text{Beta}(\alpha, \beta).$$

Exercise: Find

1. $\mathbb{E}Y$ and $\text{Var } Y$.
2. The marginal distribution of Y.

Exercise: Let (X, Y) be a pair of rvs such that

$$Y|X \sim \text{Beta}(3/2 - X, 1/2 + X)$$
$$X \sim \text{Uniform}(0, 1)$$

1. Find $\text{Cov}(X, Y)$
2. Find $\text{Var}\, Y$.

Multinoulli trial

A *multinoulli trial* is an experiment in which there are K possible outcomes which occur with the probabilities $p_1, \ldots, p_K$, where $\sum_{k=1}^{K} p_k = 1$.

Extends Bernoulli trial (two outcomes) to two or more outcomes.

Multinoulli distribution

Let the random variables $X_1, \ldots, X_K$ encode the outcome of a multinoulli trial as

$$X_k = \begin{cases} 1 & \text{if outcome } k \text{ occurs} \\ 0 & \text{otherwise} \end{cases} \quad \text{for } k = 1, \ldots, K.$$

Then the set $(X_1, \ldots, X_K)$ of K rvs has the *multinoulli distribution* and we write

$$(X_1, \ldots, X_K) \sim \text{Multinoulli}(p_1, \ldots, p_K).$$

Exercise:

1. Write down the joint pmf of $(X_1, \ldots, X_K)$.
2. Find the marginal pmf of X_k for each $k = 1, \ldots, K$.
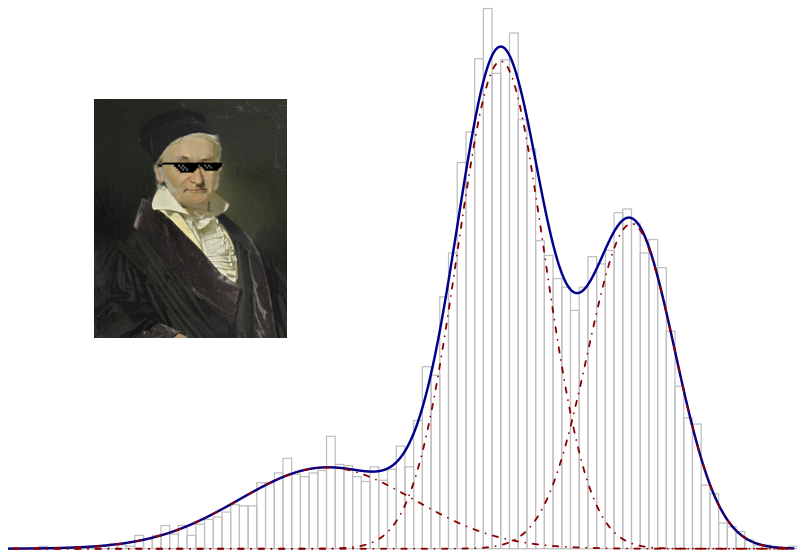3. Find $\mathbb{E}X_k$ and $\text{Cov}(X_k, X_{k'})$.

Gaussian mixture hierarchical model

Consider the hierarchical model

$$Y|(X_1, \ldots, X_K) \sim \text{Normal}\left(\sum_{k=1}^{K} X_k \mu_k, \sum_{k=1}^{K} X_k \sigma_k^2\right) \quad \text{and}$$

$$(X_1, \ldots, X_K) \sim \text{Multinoulli}(p_1, \ldots, p_K),$$

where $(\mu_1, \sigma_1^2), \ldots, (\mu_K, \sigma_K^2)$ are K mean and variance pairs.

Exercise: Find $\mathbb{E}Y$ and $\text{Var}\,Y$. Then find the marginal distribution of Y.

Multinomial distribution

For $k = 1, \ldots, K$, let Y_k be # times outcome k occurrs in n independent Multinoulli trials with the outcome probabilities $p_1, \ldots, p_K$.

Then the set of rvs $(Y_1, \ldots, Y_K)$ has the *multinomial distribution* and we write

$$(Y_1, \ldots, Y_K) \sim \text{Multinomial}(n, p_1, \ldots, p_K).$$

The joint pmf of $(Y_1, \ldots, Y_K)$ is given by

$$p(y_1, \ldots, y_K) = \left(\frac{n!}{y_1! \cdots y_K!} \right) p_1^{y_1} \cdots p_K^{y_K}$$

for $(y_1, \ldots, y_K) \in \{0, 1, \ldots, n\}^K$ st $\sum_{k=1}^{K} y_k = n$.

Note: For $k = 1, \ldots, K$ the marginal distribution of y_k is $\text{Binomial}(n, p_k)$.

Exercise: Let $Z \sim \text{Normal}(0, 1)$ and let $\mu \in \mathbb{R}$.

1. Find the pdf of $W = (Z + \mu)^2$. This has the *non-central chi-squared dist*.
2. Find the marginal pdf of U in the hierarchical model

$$U|K \sim \chi^2_{1+2K}$$
$$K \sim \text{Poisson}(\mu^2/2).$$

3. Compare these using the fact that

$$\frac{2^k}{\sqrt{2\pi}} \frac{k!}{(2k)!} = \frac{1}{\Gamma(k+1/2)2^{k+1/2}} \quad \text{for all } k = 0, 1, \dots$$