Data analysis
Principal component analysis

Olivier Roustant

February 14, 2022

Principal Component Analysis (PCA): Outline

The aim: To reduce dimension

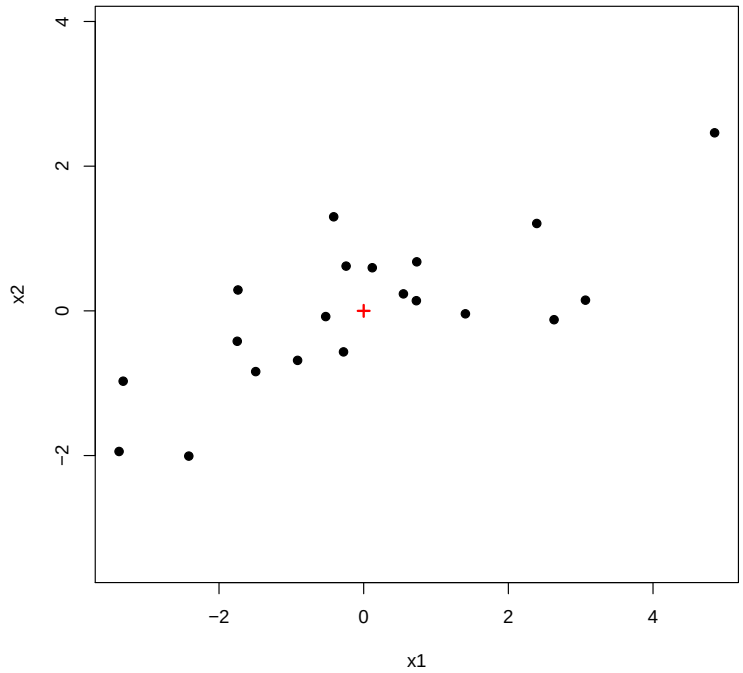

This is a $2D$ cloud of points, centered at 0.
Can you find a 1D axis 'containing' the maximum of information?

Inertia

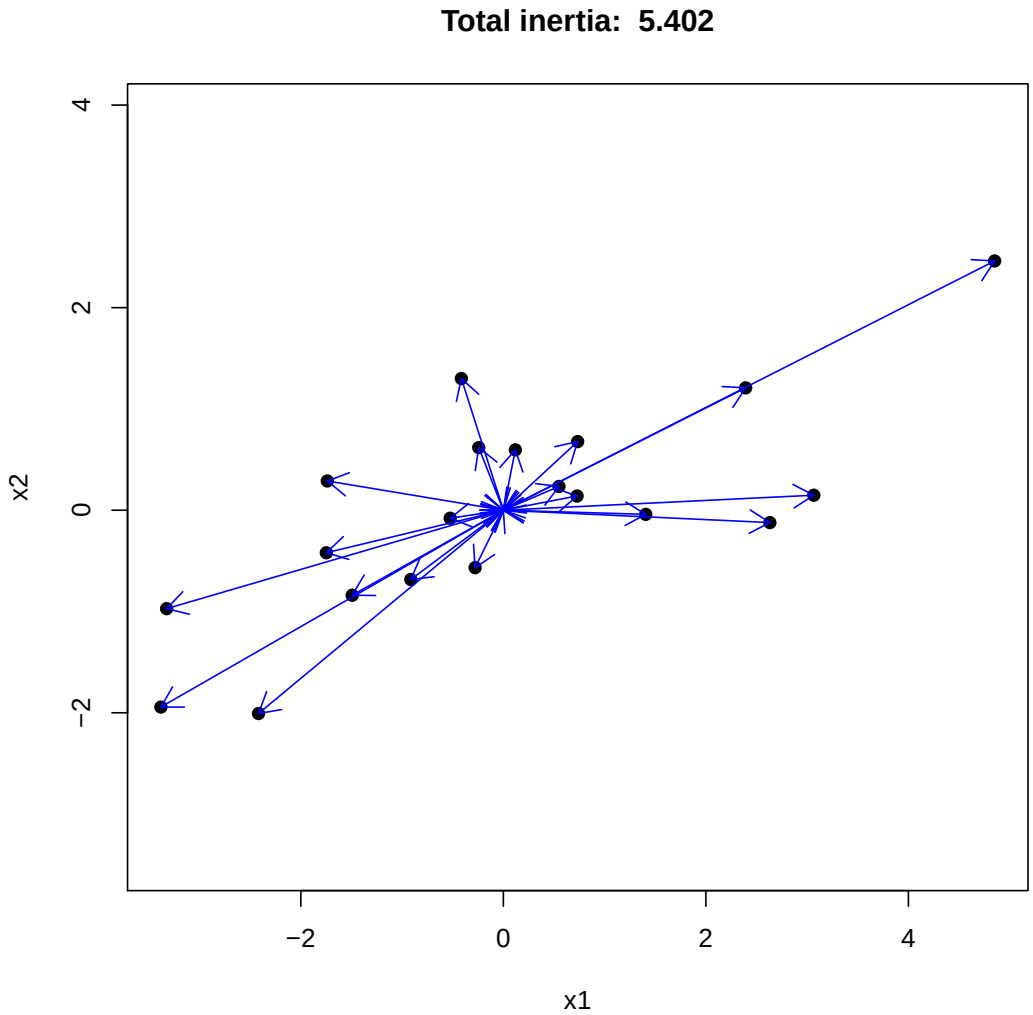

Total inertia: mean square of distances to the center.

Inertia

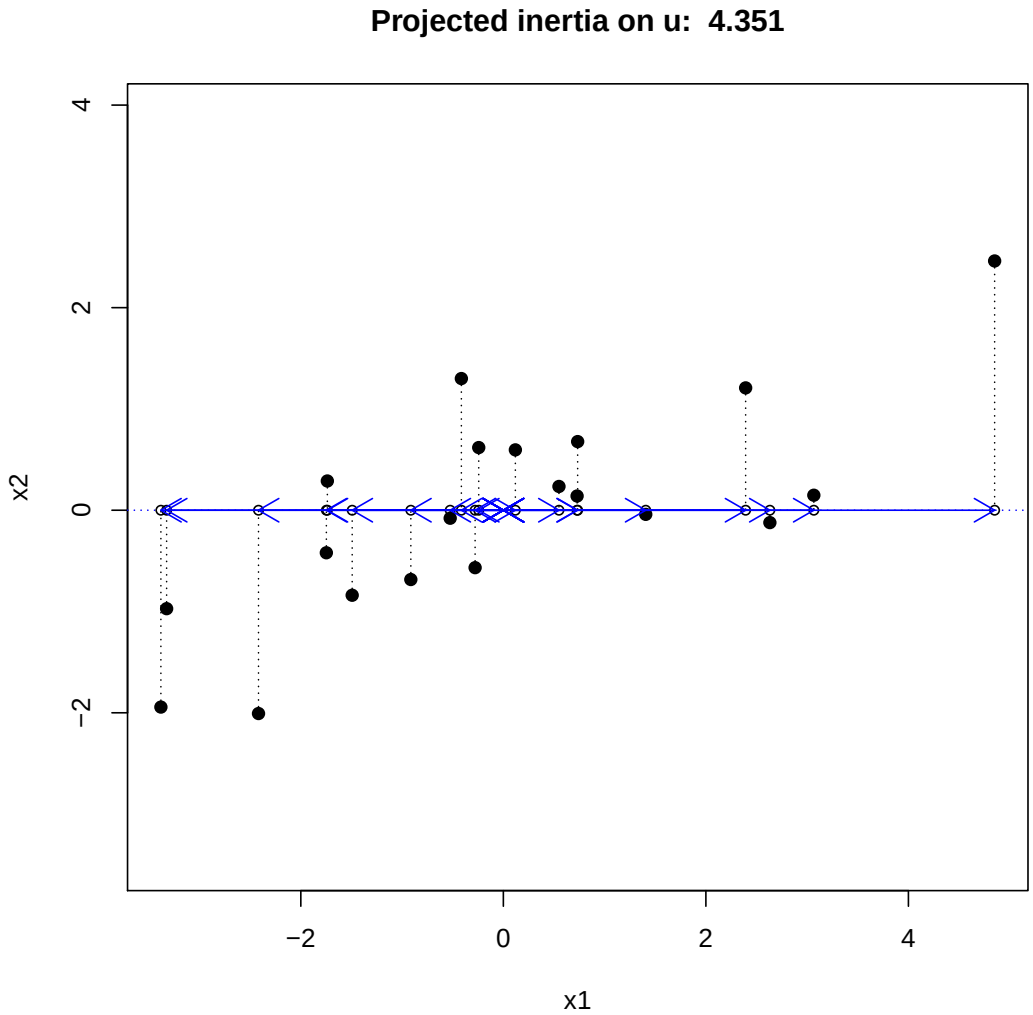

Projected inertia: inertia of projections. How much do we lose?

Maximizing the projected inertia

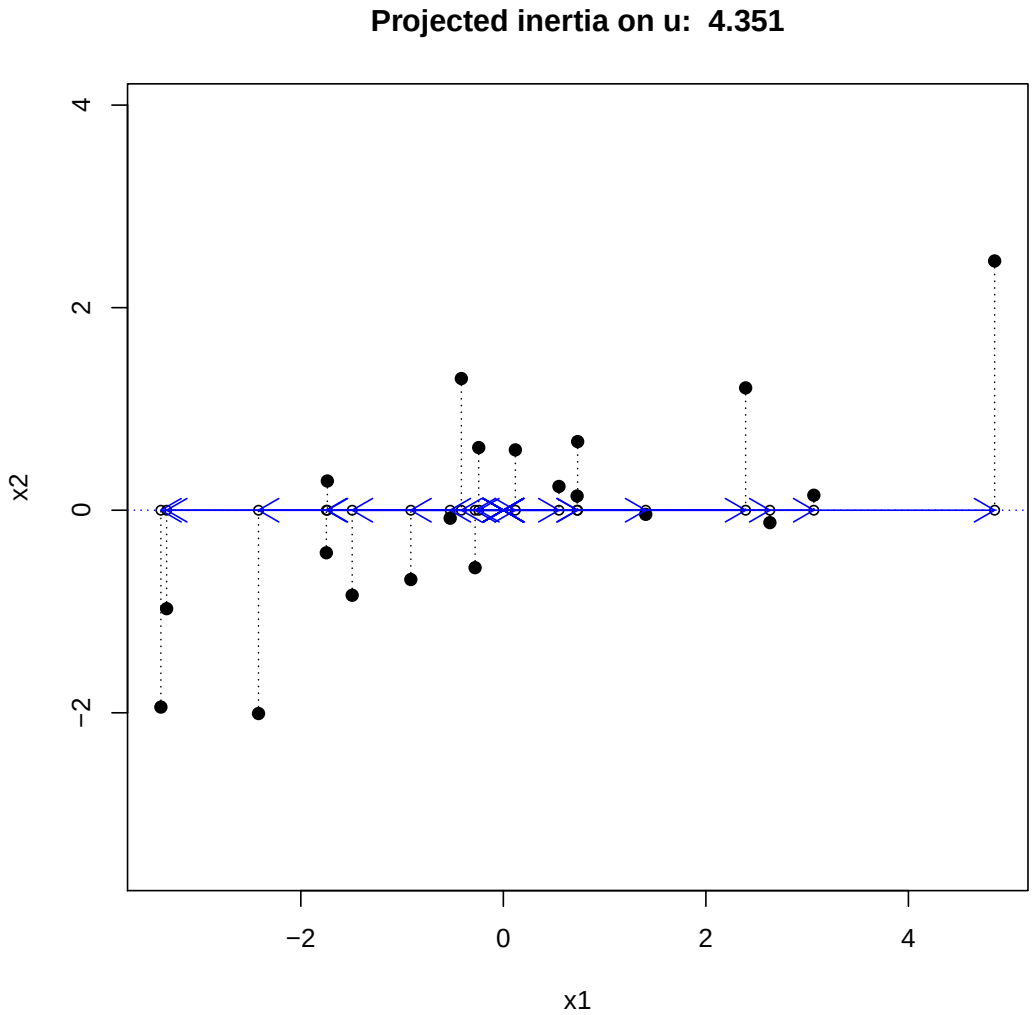

Projected inertia: For what axis is it maximal?

Maximizing the projected inertia

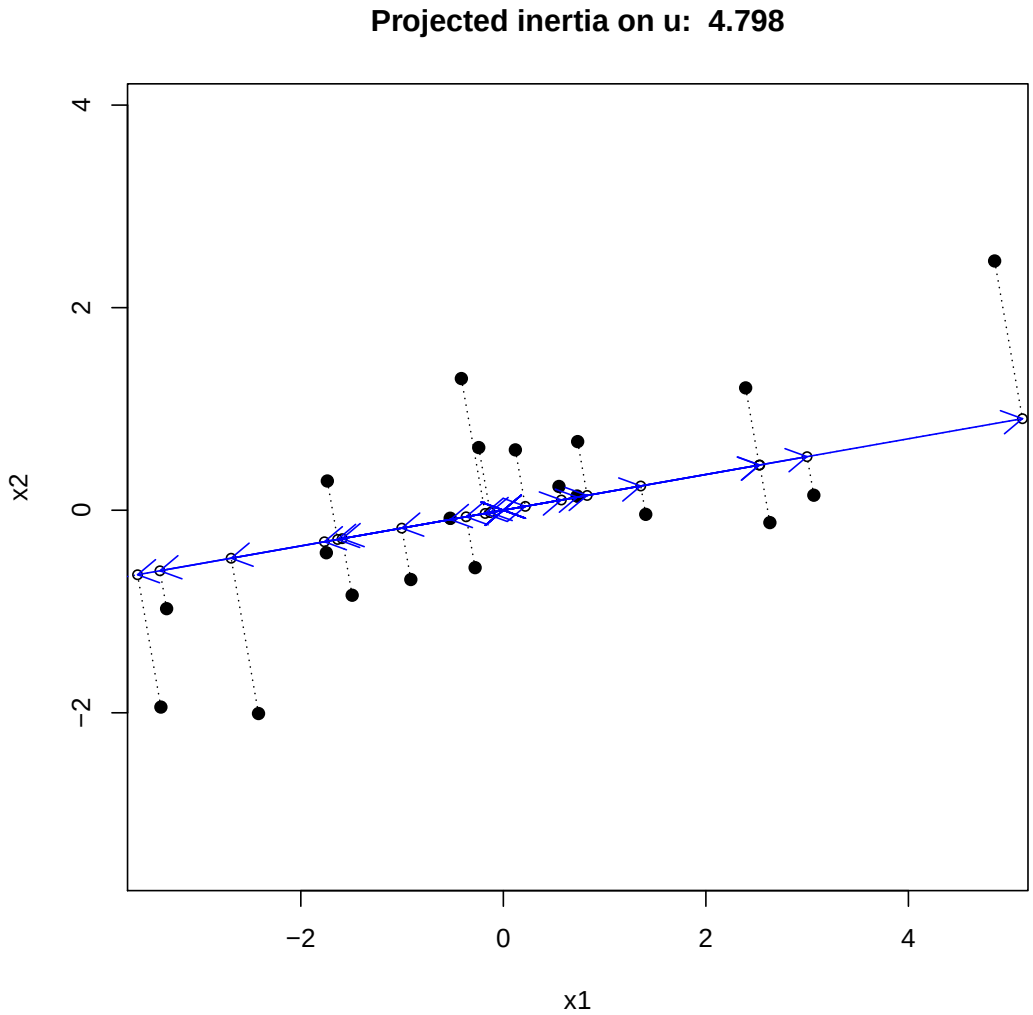

Projected inertia: For what axis is it maximal?

Maximizing the projected inertia

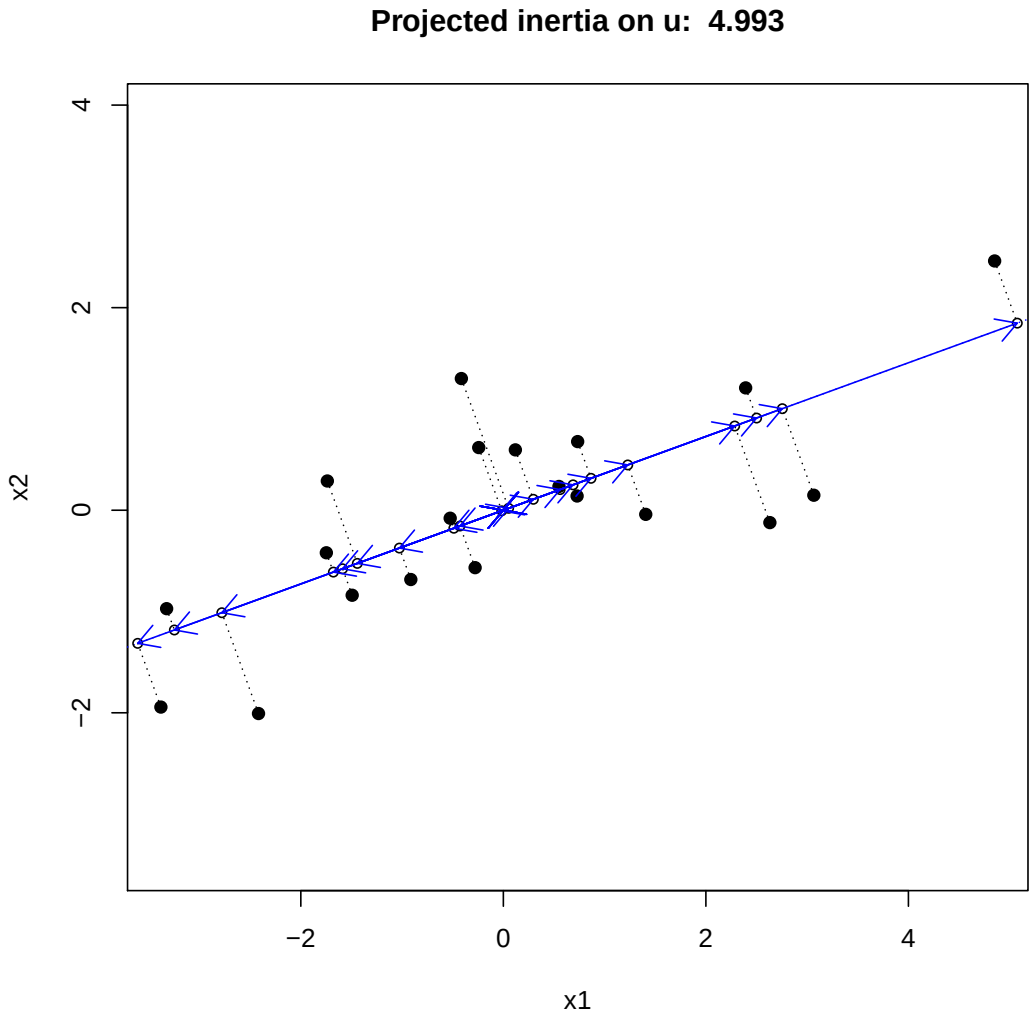

Projected inertia: For what axis is it maximal?

Maximizing the projected inertia

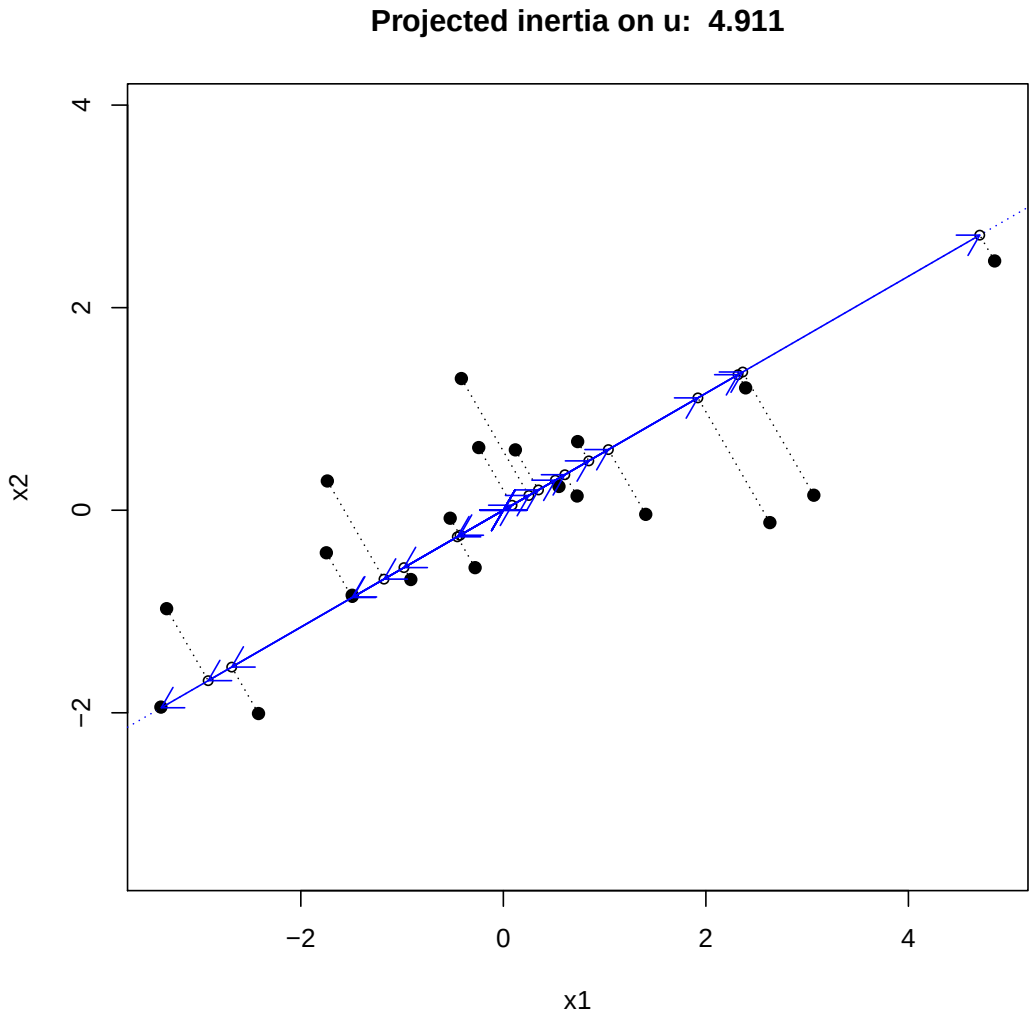

Projected inertia: For what axis is it maximal?

Maximizing the projected inertia

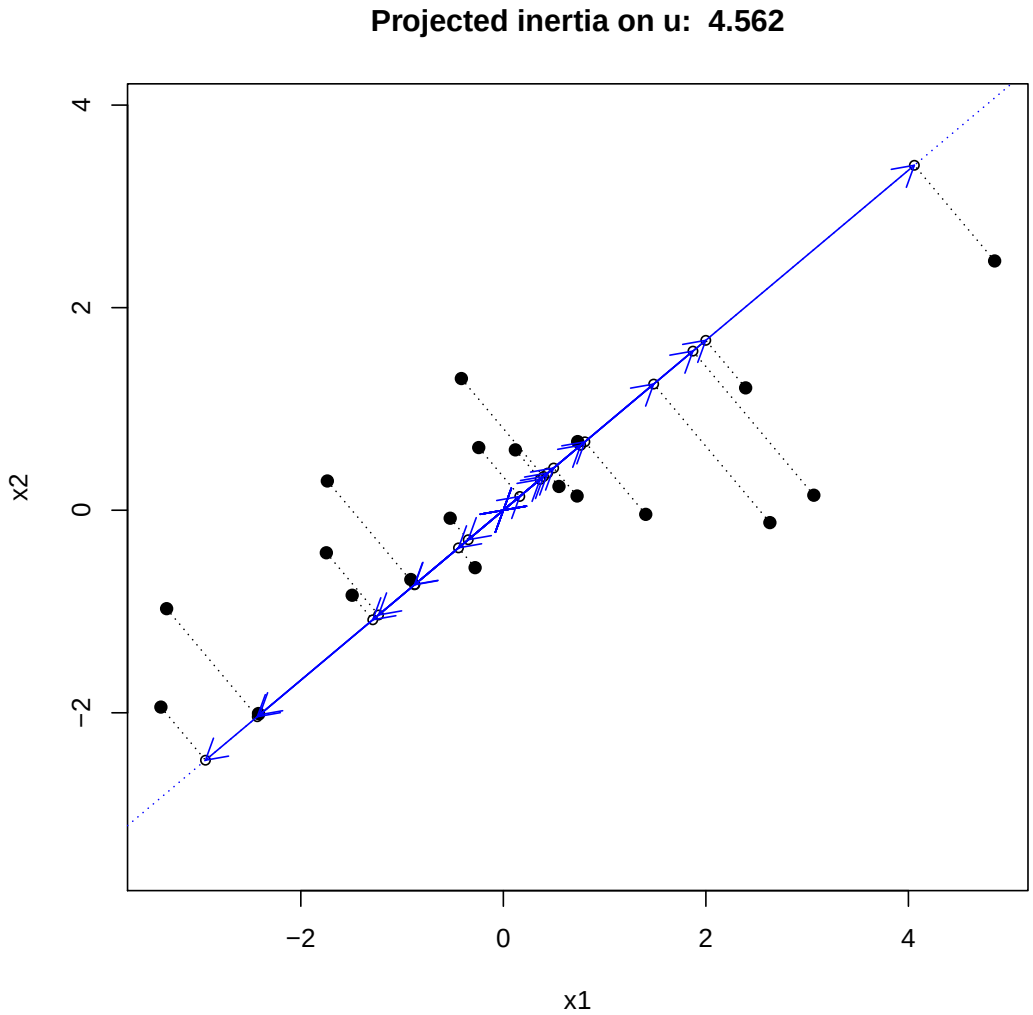

Projected inertia: For what axis is it maximal?

Maximizing the projected inertia

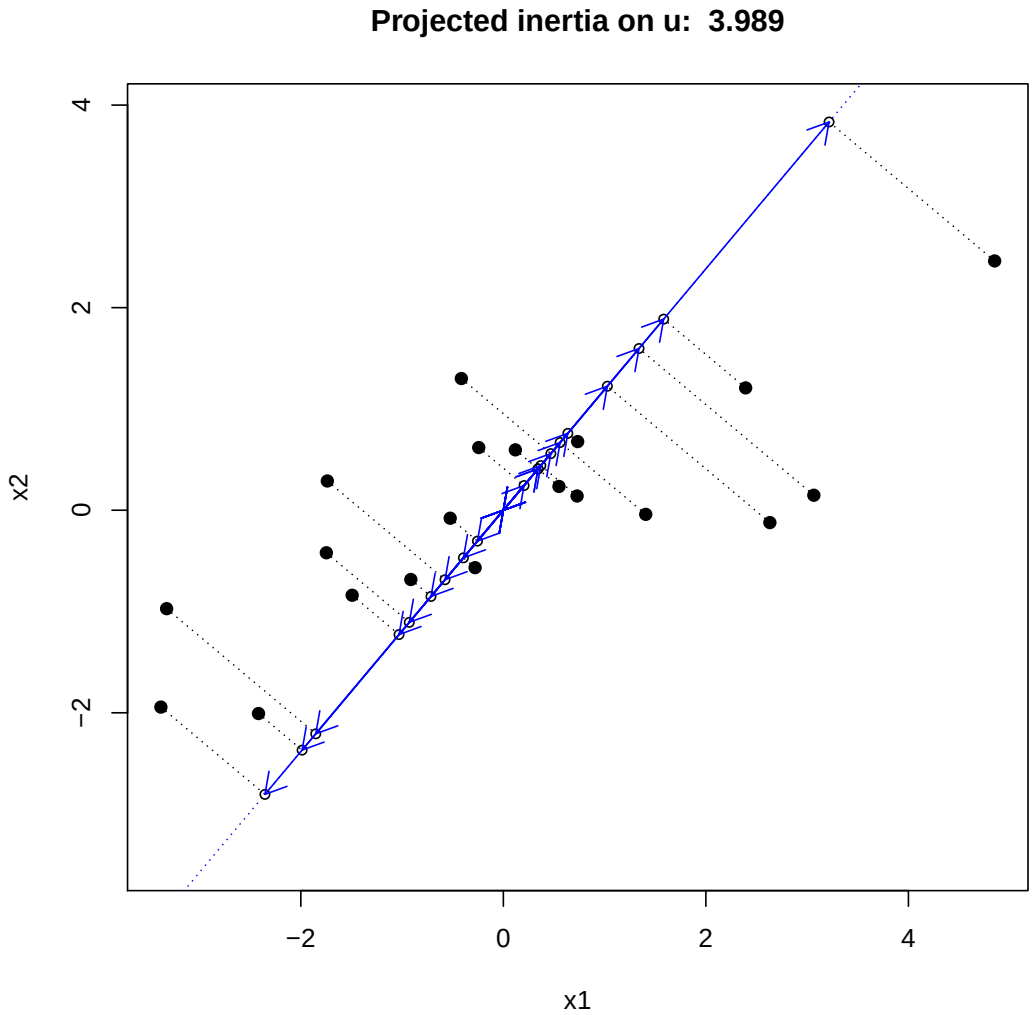

Projected inertia: For what axis is it maximal?

Maximizing the projected inertia

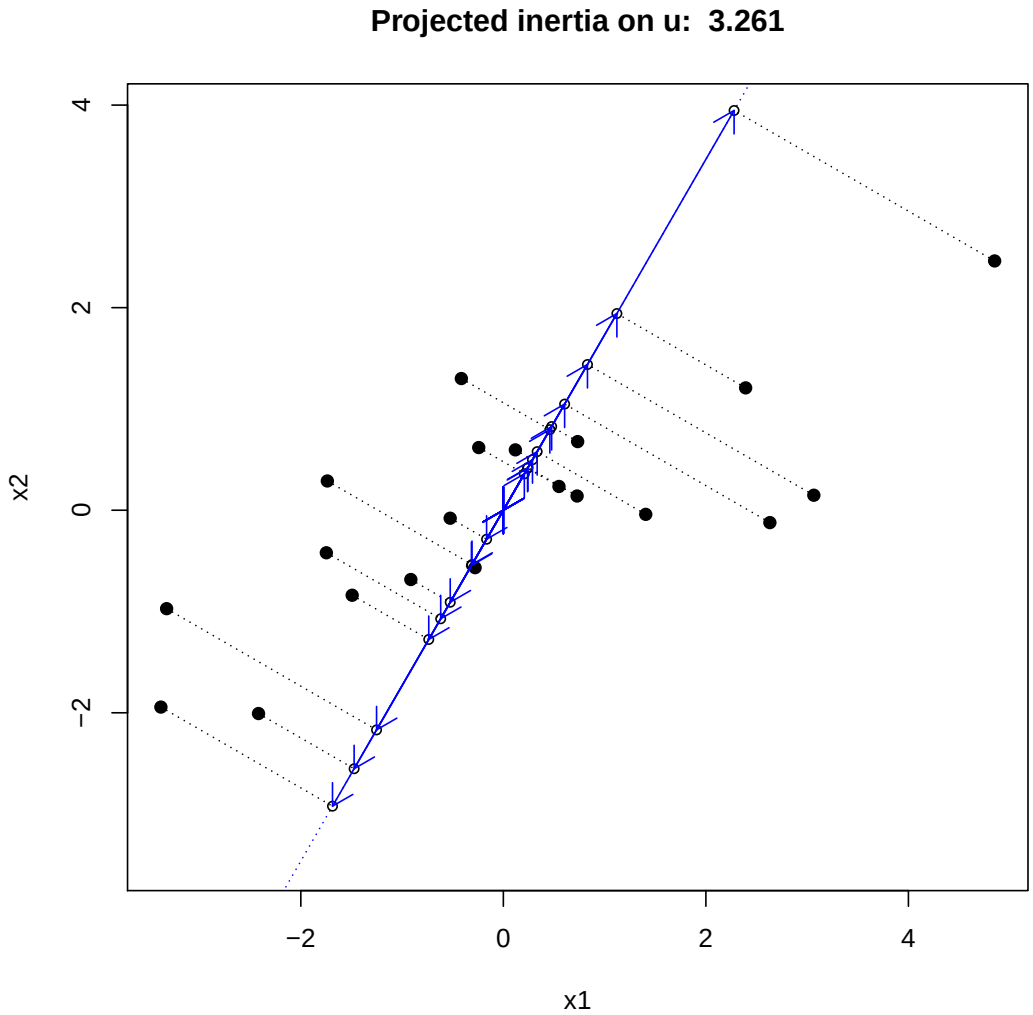

Projected inertia: For what axis is it maximal?

Maximizing the projected inertia

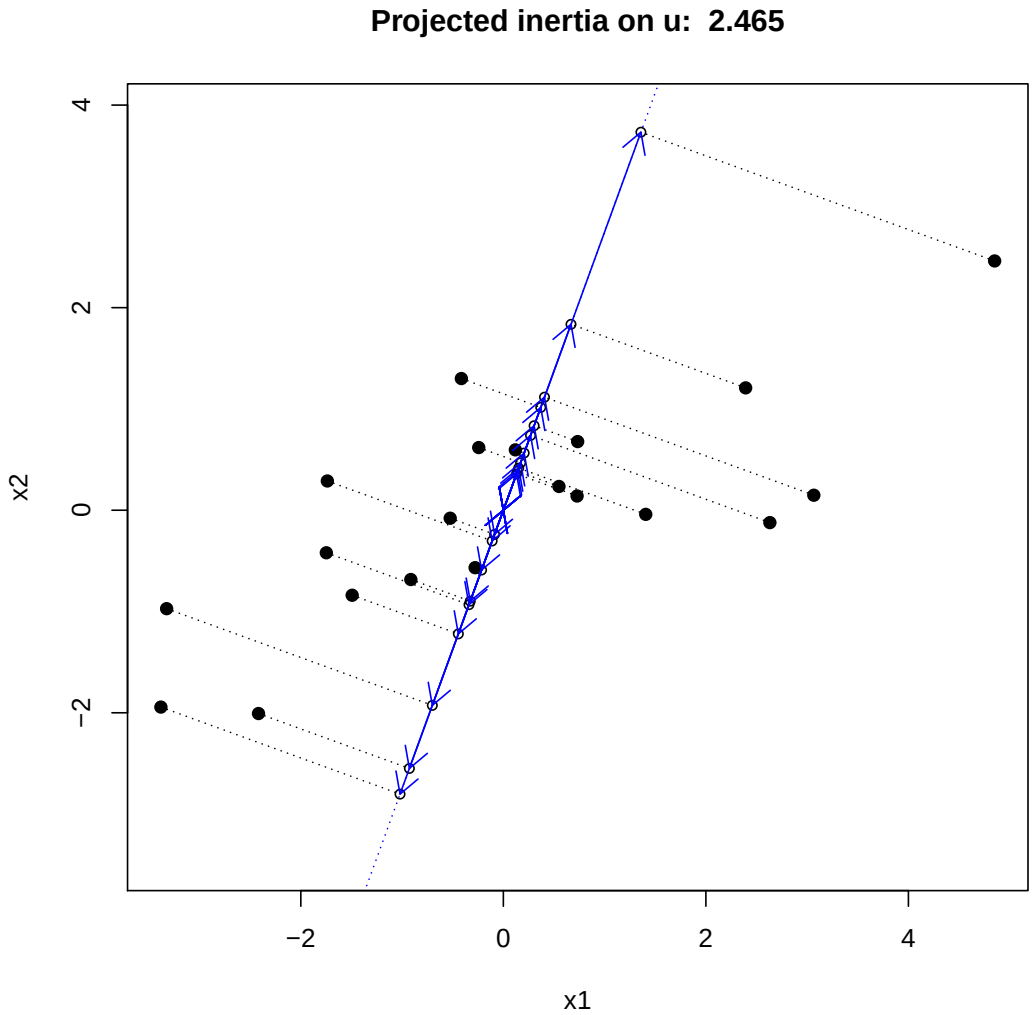

Projected inertia: For what axis is it maximal?

Maximizing the projected inertia

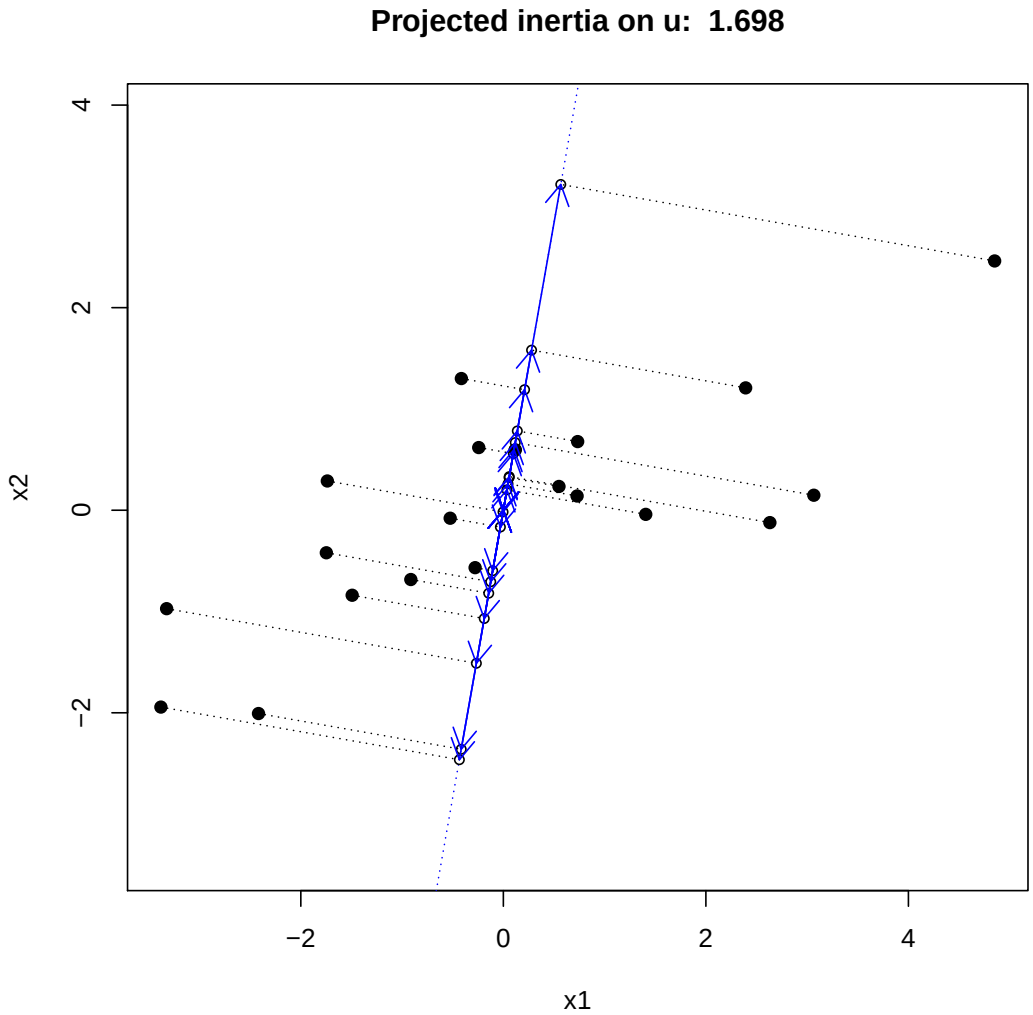

Projected inertia: For what axis is it maximal?

Maximizing the projected inertia

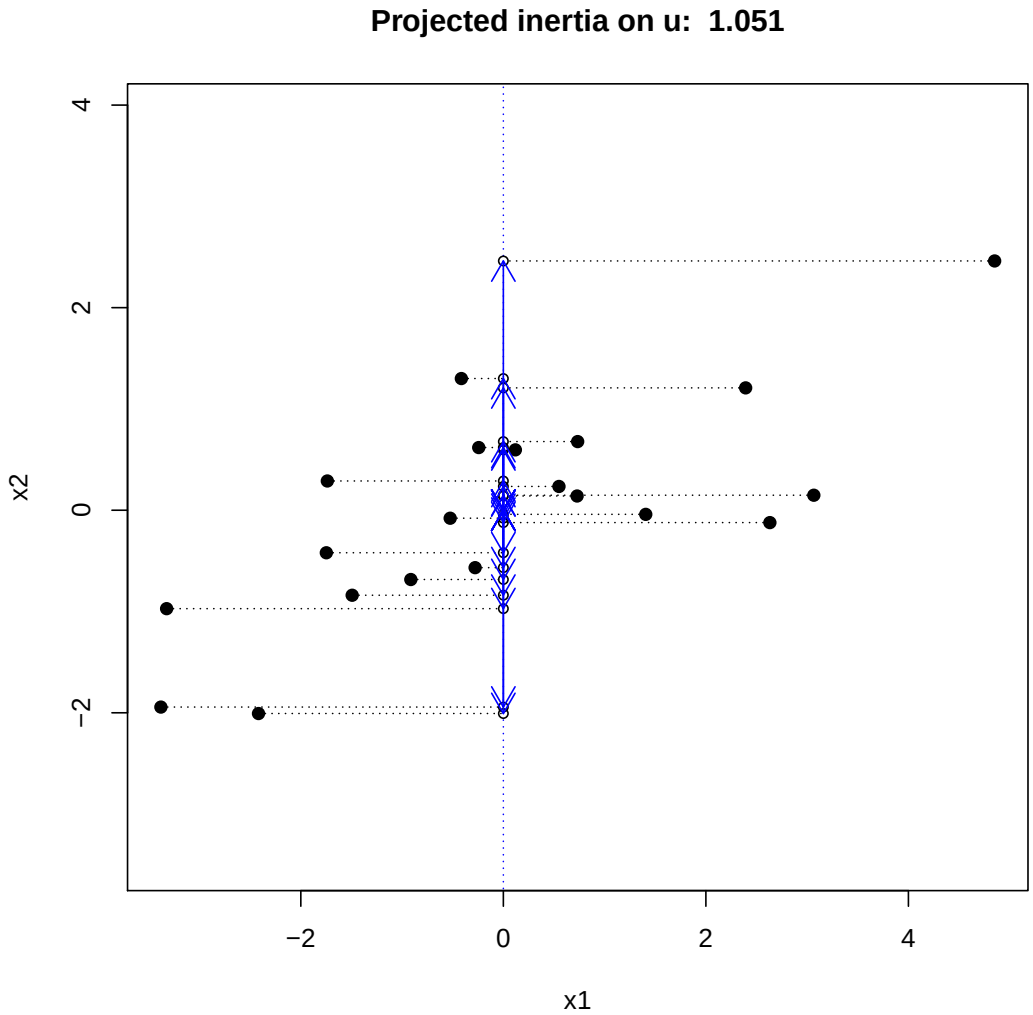

Projected inertia: For what axis is it maximal?

Maximizing the projected inertia

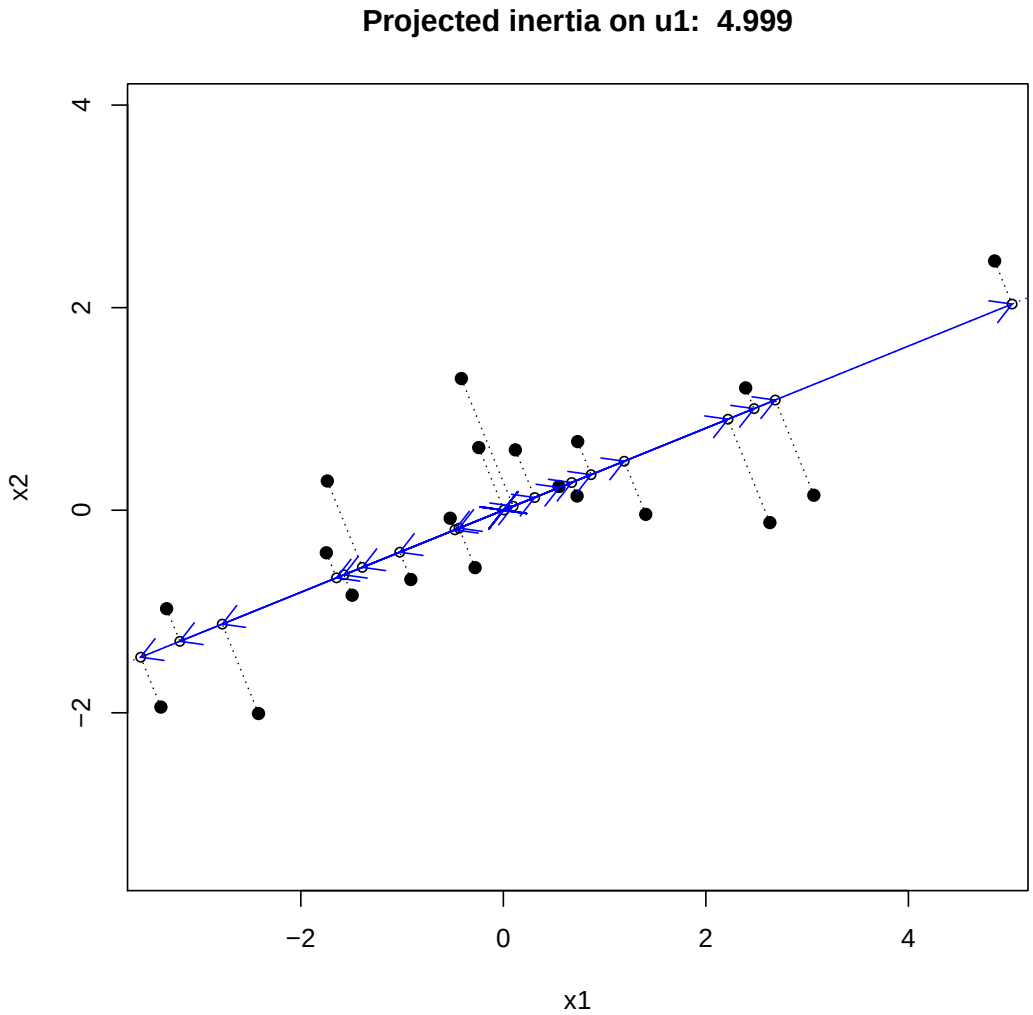

Projected inertia: Maximal for the largest eigenvalue of the covariance matrix

Maximizing the projected inertia, recursion

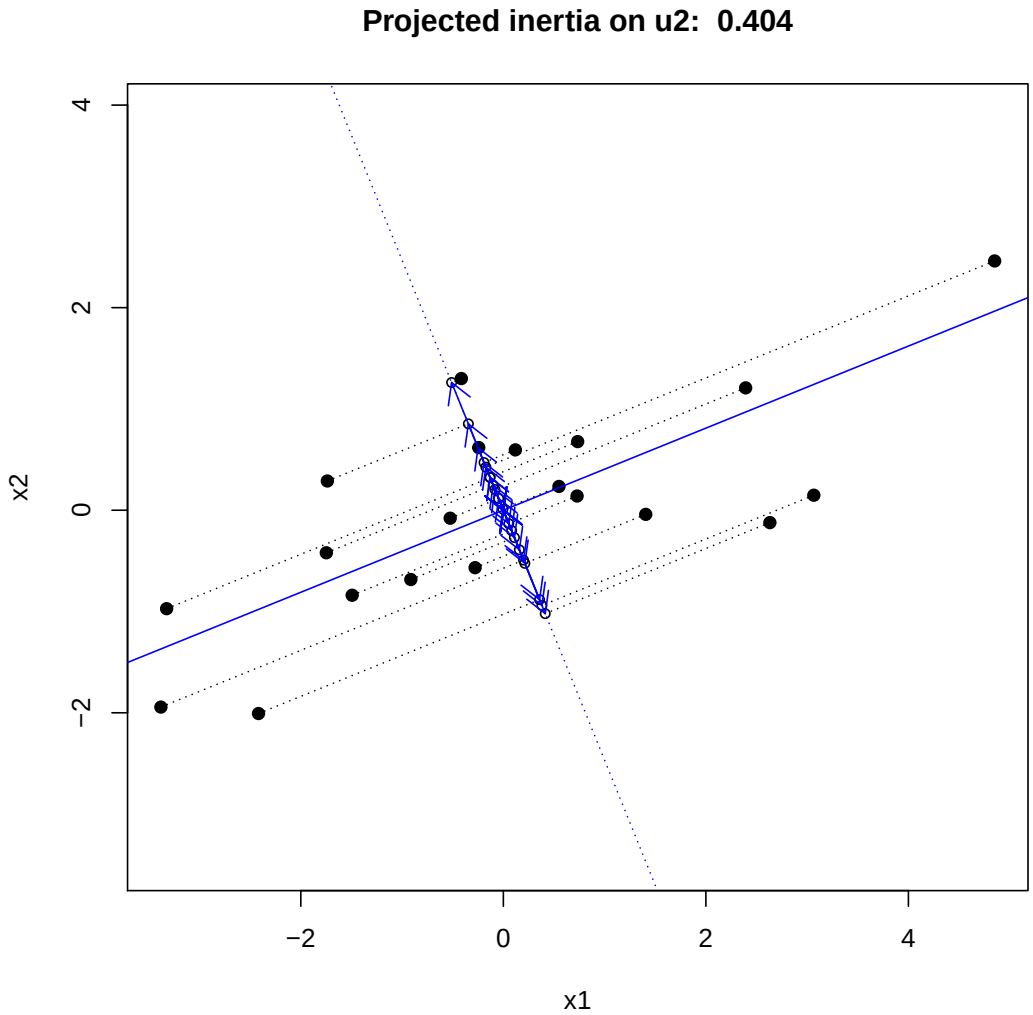

The second largest eigenvalue maximizes the projected inertia in the orthogonal of the first

Maximizing the projected inertia, summary

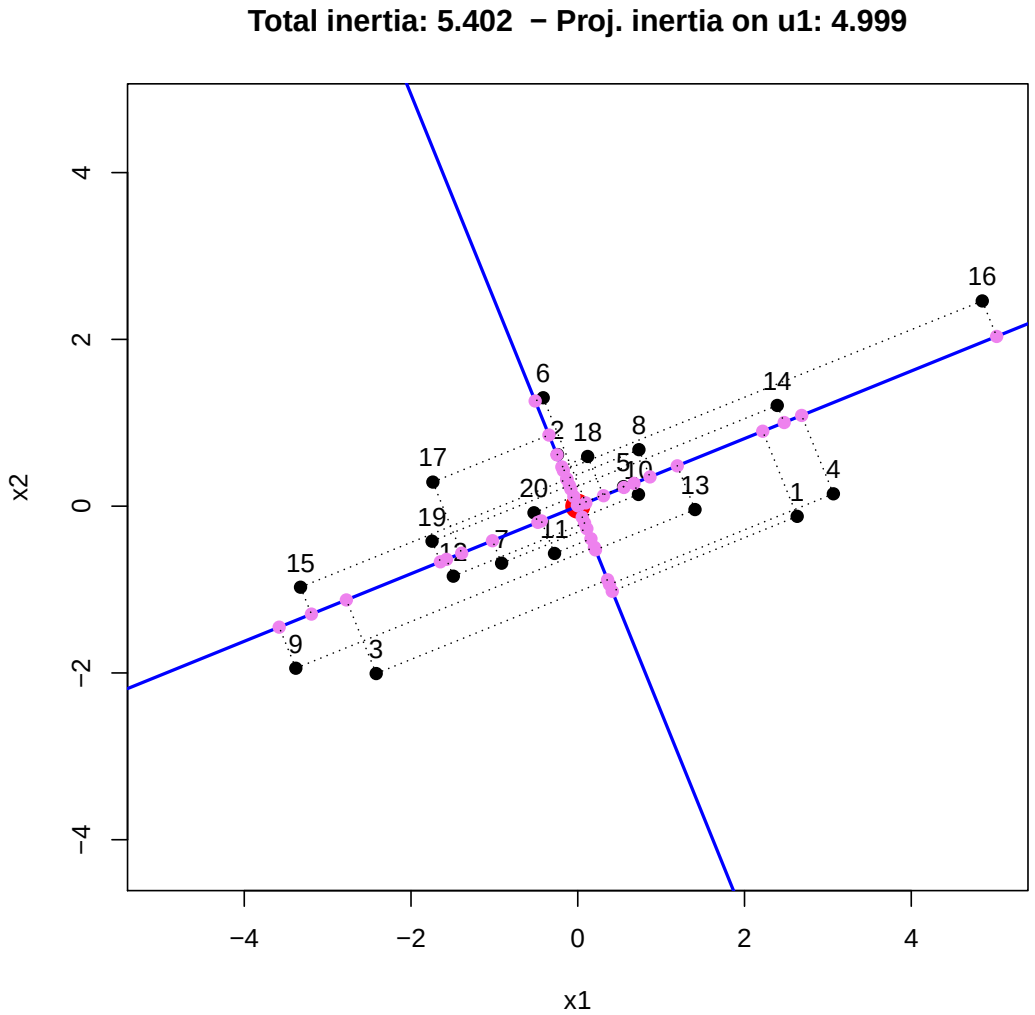

Projected points on the first two ‘principal components’

Maximizing the projected inertia, summary

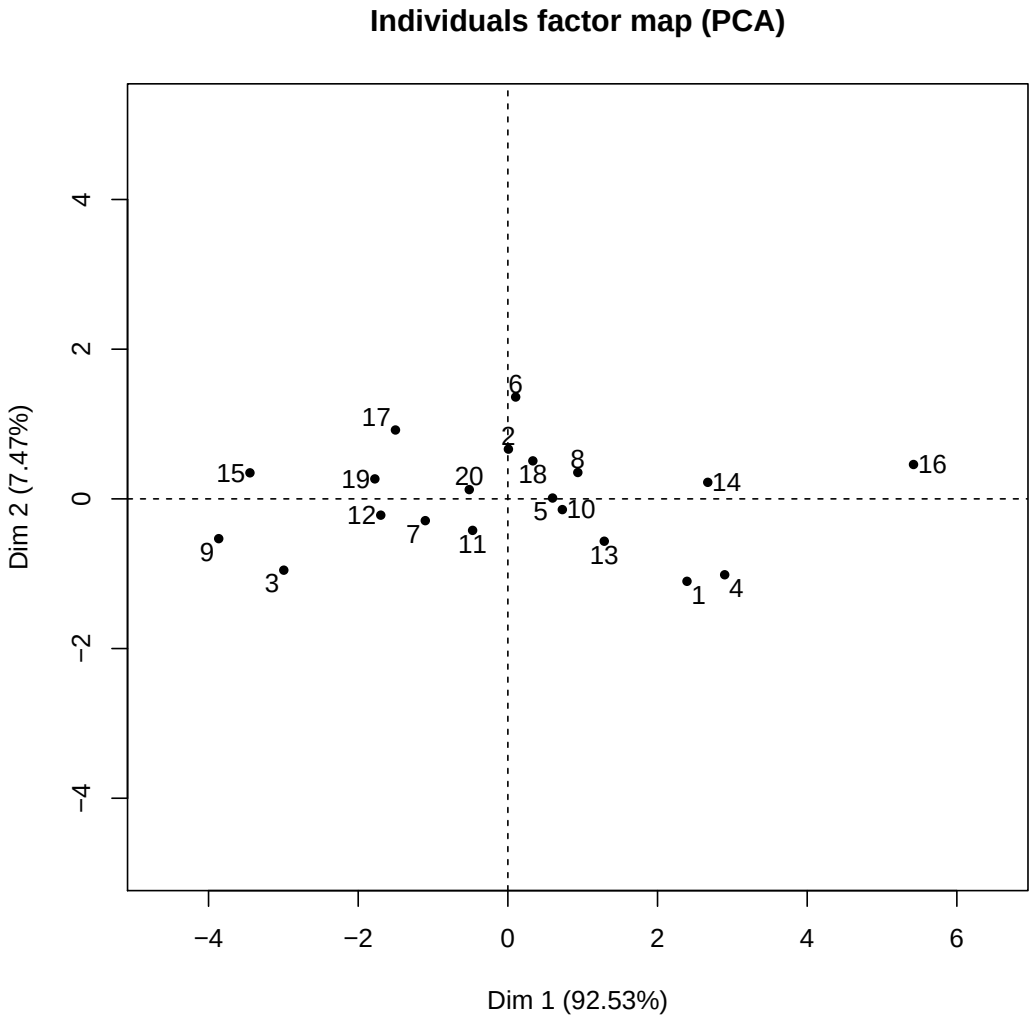

Representation with package FactoMineR. Percentages are inertia ratio w.r.t. total inertia

Theory

Notations and assumption

- $\mathbf{X}$: a matrix of size $n \times p$, representing the data:

	$\mathbf{x}^1$	$\dots$	$\mathbf{x}^j$	$\dots$	$\mathbf{x}^p$
$\mathbf{x}_1$	x_1^1	$\dots$	x_1^j	$\dots$	x_1^p
$\vdots$	$\vdots$		$\vdots$		$\vdots$
$\mathbf{x}_i$	x_i^1	$\dots$	x_i^j	$\dots$	x_i^p
$\vdots$	$\vdots$		$\vdots$		$\vdots$
$\mathbf{x}_n$	x_n^1	$\dots$	x_n^j	$\dots$	x_n^p

- $\mathbf{g}$: center of gravity (empirical mean), $\mathbf{g} = \bar{\mathbf{x}} = \frac{1}{n}\sum_{i=1}^{n} \mathbf{x}_i \, (\in \mathbb{R}^p)$.

$\mathbf{g}$	$\overline{\mathbf{x}^1}$	$\dots$	$\overline{\mathbf{x}^j}$	$\dots$	$\overline{\mathbf{x}^p}$

We assume that $\mathbf{g} = 0$, i.e. the data have been centered.

Notations and assumption

- The rows of $\mathbf{X}$ lie in $\mathbb{R}^p$, and form the **indivuals space**.
  It is an Euclidean space, equipped with the usual ℓ^2 norm $\|.\|$.

- The columns of $\mathbf{X}$ lie in $\mathbb{R}^n$, and form the **variables space**.
  It is an Euclidean space. Instead of choosing the usual ℓ^2 norm, we rescale it by $1/n$. Indeed, as the data are centered, it corresponds to the empirical covariance:

$$\langle \mathbf{x}^j, \mathbf{x}^k \rangle_{\mathbb{R}^n} := \frac{1}{n} \sum_{i=1}^{n} x_i^j x_i^k = \widehat{\mathrm{cov}}(\mathbf{x}^j, \mathbf{x}^k).$$

  Notice that **orthogonal variables = uncorrelated variables**.
  Γ denotes the $p \times p$ empirical covariance matrix:

$$\Gamma = \left(\widehat{\mathrm{cov}}(\mathbf{x}^j, \mathbf{x}^k) \right)_{1 \leq j,k \leq p} = \frac{1}{n} \mathbf{X}^\top \mathbf{X} = \frac{1}{n} \sum_{i=1}^{n} \mathbf{x}_i \mathbf{x}_i^\top .$$

Notations and assumption

- **Inertia**: mean squared distance of the data to their center (here **0**),

$$\mathcal{I} = \frac{1}{n}\sum_{i=1}^{n} \|\mathbf{x}_i\|^2$$

- **Projected inertia** on a subspace $F \subseteq \mathbb{R}^p$. Same definition for the projected points onto F (we denote by Π_F the projection operator):

$$\mathcal{I}_F = \frac{1}{n}\sum_{i=1}^{n} \|\Pi_F(\mathbf{x}_i)\|^2$$

Properties of inertia

Link with variance, and inertia decomposition.

Consider a $1D$ axis spanned by a unit vector $\mathbf{a}$, and denote $\mathcal{I}_{\mathbf{a}} = \mathcal{I}_{\mathbb{R}\mathbf{a}}$. Then:

$$\mathcal{I}_{\mathbf{a}} = \mathbf{a}^\top \Gamma \mathbf{a}, \qquad \text{and} \qquad \mathcal{I} = \mathcal{I}_{\mathbf{a}} + \mathcal{I}_{\mathbf{a}^\perp}$$

Moreover, $\mathcal{I}_{\mathbf{a}}$ and $\mathcal{I}$ are interpreted in terms of variances:

- $\mathcal{I}_{\mathbf{a}}$ is the empirical variance of the projected points onto $\mathbb{R}\mathbf{a}$,
- $\mathcal{I}$ is the sum of the empirical variances of the p variables:

$$\mathcal{I}_{\mathbf{a}} = \frac{1}{n}\sum_{i=1}^{n} \langle \mathbf{x}_i, \mathbf{a} \rangle^2, \qquad \mathcal{I} = \sum_{j=1}^{p} \hat{\sigma}_j^2, \quad \text{with} \quad \hat{\sigma}_j^2 = \frac{1}{n}\sum_{i=1}^{n} (x_i^j)^2$$

Remark: The empirical variances are computed here by dividing by n the sum of squares, contrarily to unbiased statistical estimates (division by $n-1$).

Properties of inertia (proofs)

Left to exercise.

Main result

Theorem (principal component analysis)

As the covariance matrix Γ is real symmetric, it admits a spectral decomposition in orthogonal eigenspaces. Denote $\lambda_1 \geq \cdots \geq \lambda_p \geq 0$ the eigenvalues, and $\mathbf{v}_1, \ldots, \mathbf{v}_p$ orthogonal eigenvectors. Then:

- $\mathbf{v}_1$ maximizes $\mathcal{I}_{\mathbf{a}}$ over $\mathbf{a}$, which is then equal to λ_1.
- $\mathbf{v}_2$ maximizes $\mathcal{I}_{\mathbf{a}}$ over $\mathbf{a}$ in $(\mathbf{v}_1)^{\perp}$, which is then equal to λ_2.
- $\mathbf{v}_3$ maximizes $\mathcal{I}_{\mathbf{a}}$ over $\mathbf{a}$ in $(\mathbf{v}_1, \mathbf{v}_2)^{\perp}$, which is then equal to λ_3.
- ...

Furthermore the inertia (called *total inertia*) is decomposed:

$$\mathcal{I} = \mathcal{I}_{\mathbf{v}_1} + \cdots + \mathcal{I}_{\mathbf{v}_p} = \lambda_1 + \cdots + \lambda_p$$

Main result (proof)

Left to exercise.
Hint: Use the decomposition of **a** in the basis of eigenvectors.

Principal components

- The eigenvectors $\mathbf{v}_1, \ldots, \mathbf{v}_p$ define a new orthonormal basis in $\mathbb{R}^p$.
- The change of variables is defined by:

$$\mathbf{C} = \mathbf{XP}, \qquad \text{with} \quad \mathbf{P} = [\mathbf{v_1}, \ldots, \mathbf{v_p}] .$$

The $n \times p$ matrix $\mathbf{C}$ is called **matrix of principal components**. The columns of $\mathbf{C}$ are called **principal variables**. They contain the coordinates of the individuals in the new space.

Principal components

- Principal variables are **linear combinations of the original variables**, with coefficients given by the eigenvectors:

$$\mathbf{C}^j = \mathbf{X}\mathbf{P}_j = \sum_{k=1}^{p} (\mathbf{v}_j)_k \mathbf{x}^k$$

- Principal variables are **uncorrelated and** $\widehat{\mathrm{var}}(\mathbf{C}^k) = \lambda_k$:

$$\left(\widehat{\mathrm{cov}}(\mathbf{C}^j, \mathbf{C}^k)\right)_{1 \leq j,k \leq p} = \frac{1}{n}\mathbf{C}^\top \mathbf{C} = \mathbf{P}^\top \Gamma \mathbf{P} = \mathrm{diag}(\lambda_1, \ldots, \lambda_p).$$

Remark: singular value / spectral decomposition

PCA can be done with **Singular Value Decomposition (SVD)**, which decomposes a rectangular matrix $n \times m$ or rank r as

$$\mathbf{X} = \mathbf{U}\Lambda^{1/2}\mathbf{V}^\top,$$

where Λ is the diagonal matrix containing the r non-zero eigenvalues of $\mathbf{X}^\top\mathbf{X}$ (or $\mathbf{X}\mathbf{X}^\top$), ranked by decreasing order, and $\mathbf{U}$ (resp. $\mathbf{V}$) is an orthogonal matrix for $\|.\|_{\mathbb{R}^n}$ (resp. for $\|.\|_{\mathbb{R}^m}$) containing the eigenvectors of $\mathbf{X}\mathbf{X}^\top$ (resp. $\mathbf{X}^\top\mathbf{X}$).

In the frequent case when $p = r$ (e.g. $n > p$), we have:

$$\mathbf{V} = \mathbf{P}, \qquad \Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_n).$$

(In the general case, $\mathbf{V}$ contains the r columns of $\mathbf{P}$ corresponding to non-zero eigenvalues.) Further, due to our definition of the scalar product in $\mathbb{R}^n$, we have $\frac{1}{n}\mathbf{U}^\top\mathbf{U} = I_p$. Then, you can recover all the formulas of the textbook, e.g.:

$$\mathbf{C} = \mathbf{X}\mathbf{P} = \mathbf{U}\Lambda^{1/2}\mathbf{P}^\top\mathbf{P} = \mathbf{U}\Lambda^{1/2}.$$

Variations (metric, weights)

Changing the metric in the individuals space

Consider a new norm on $\mathbb{R}^p$, called **metric**, defined by a positive definite matrix $\mathbf{M}$, of size p:

$$\|\mathbf{x}\|_M^2 = \mathbf{x}^\top \mathbf{M}\mathbf{x}.$$

Let $\mathbf{R}$ be an invertible matrix s.t. $\mathbf{R}^\top\mathbf{R} = \mathbf{M}$ (e.g. square root, Choleski decomposition). Then, the map

$$\mathbf{R} : \begin{array}{ccc} (\mathbb{R}^p, \|.\|_M) & \to & (\mathbb{R}^p, \|.\|) \\ \mathbf{x} & \mapsto & \mathbf{Rx} \end{array}$$

is an isometry, and thus preserves distances and orthogonality.

Indeed: $\|\mathbf{Rx}\|^2 = (\mathbf{Rx})^\top(\mathbf{Rx}) = \mathbf{x}^\top\mathbf{M}\mathbf{x} = \|\mathbf{x}\|_M^2$.

Changing the metric in the individuals space

Due to the isometry property, we deduce immediately:

PCA with / without metric

$\mathbf{v}$ max. projected inertia for original data $\mathbf{x}_1, \ldots, \mathbf{x}_n$ with metric $\|.\|_M$

$\Leftrightarrow$

$\mathbf{Rv}$ max. proj. inertia for transformed data $\mathbf{Rx}_1, \ldots, \mathbf{Rx}_n$ with $\|.\|$

$\Leftrightarrow$

$\mathbf{Rv}$ is an eigenvector of $\frac{1}{n}\sum_{i=1}^{n}(\mathbf{Rx}_i)(\mathbf{Rx}_i)^\top = \mathbf{R}\left(\frac{1}{n}\mathbf{X}^\top\mathbf{X}\right)\mathbf{R}^\top$

$\Leftrightarrow$

$\mathbf{v}$ is an eigenvector of $\left(\frac{1}{n}\mathbf{X}^\top\mathbf{X}\right)\mathbf{M} = \Gamma\mathbf{M}$

Changing the metric in the individuals space

Recall that the data are assumed to be centered.

Example. Standardize (centered) data.

$$\mathbf{M} = \mathrm{diag}\left(\frac{1}{\hat{\sigma}_1^2}, \ldots, \frac{1}{\hat{\sigma}_p^2}\right)$$

Then we can choose $\mathbf{R} = \mathrm{diag}\left(\frac{1}{\hat{\sigma}_1}, \ldots, \frac{1}{\hat{\sigma}_p}\right)$. Thus doing PCA with the metric $\mathbf{M}$ is equivalent to doing usual PCA on the standardized data.

Changing the weights in the variable space

In the standard formulation, each individual $\mathbf{x}_1, \ldots, \mathbf{x}_n$ has weight $\frac{1}{n}$.

Obviously, one can use positive weights $\omega_1, \ldots, \omega_n$ that sum to one. It can be useful if some individuals have more importance.

This can be viewed as an isometric transformation in the space $\mathbb{R}^n$ by the diagonal matrix containing the square roots of ω_i.
The theory is immediately adapted, by modifying the definitions, e.g.:

$$\mathcal{I} = \sum_{i=1}^{n} \omega_i \|\mathbf{x}_i\|^2, \qquad \Gamma = \sum_{i=1}^{n} \omega_i \mathbf{x}_i \mathbf{x}_i^\top.$$

Link between the notations slides / textbook

In this presentation, we started from the simplest case (centered data, standard Euclidean metric, same weights), and explained how to obtain the general formula.

In the textbook, the general case is considered.

To obtain the formula of the slideshow from the textbook, you should simply use:

$$\bar{\mathbf{X}} = \mathbf{X}, \qquad \mathbf{M} = \mathbf{I}_p, \qquad \mathbf{D} = \frac{1}{n}\mathbf{I}_n, \qquad \mathbf{V} = \mathbf{P}, \qquad \mathbf{S} = \Gamma$$

Results interpretation

Example on a temperature dataset

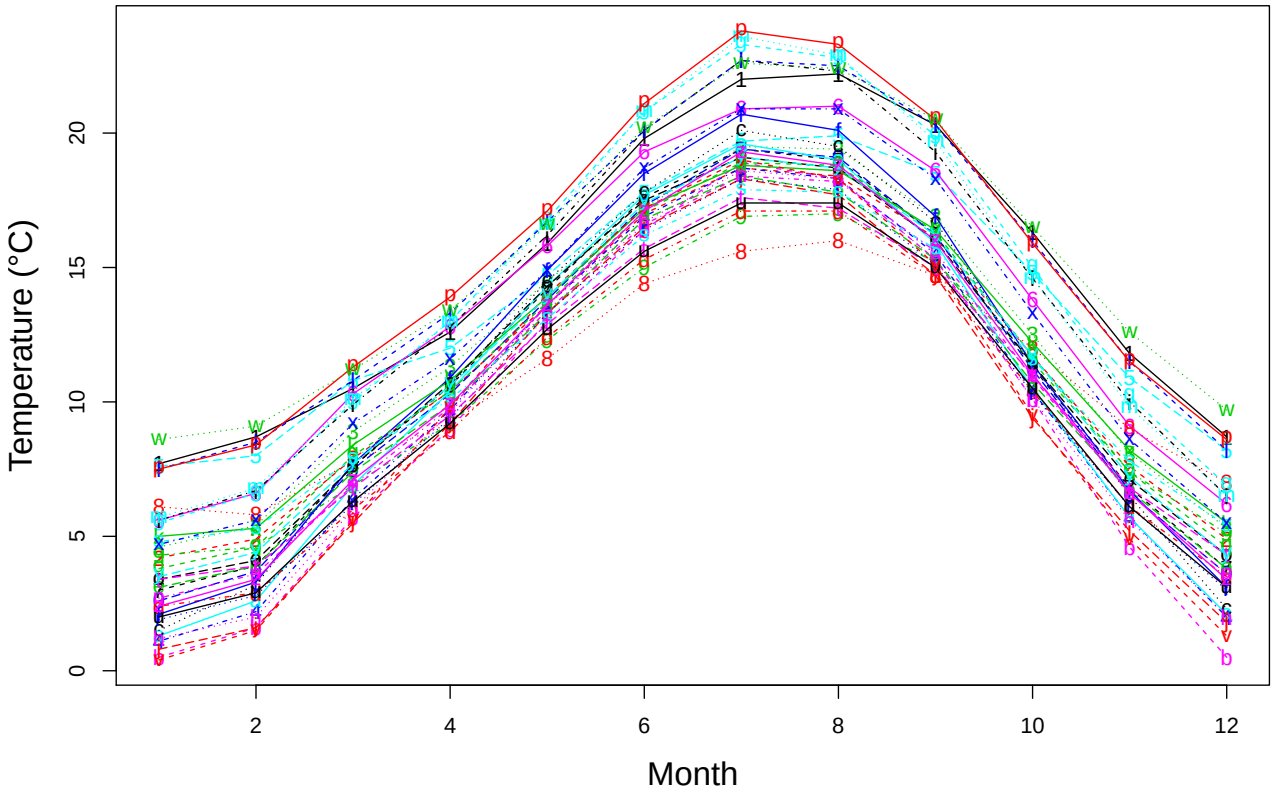

Dataset: Temperature at $n = 36$ cities (individuals) for $p = 12$ months (variables).

Dimension reduction

Here the variables are highly correlated, and a strong dimension reduction is expected. The decrease of inertia shows that 2 dimensions explain approx 99% of the variance.

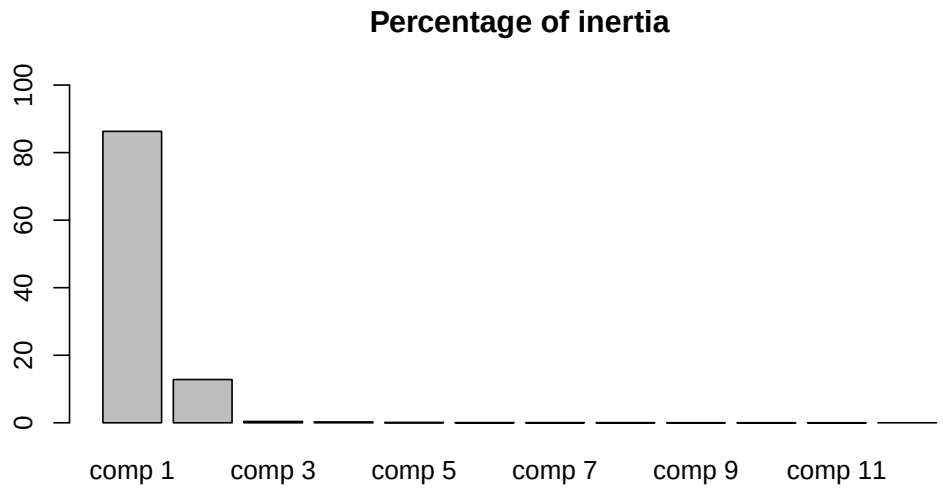

Interpretation of principal components

Remember that $\mathbf{C}^j = \sum_{k=1}^{p} (\mathbf{v}_j)_k \mathbf{x}^k$.

To interpret $\mathbf{C}^j$, look at $\mathbf{v}_j$. Here we can plot them as a function of time.

- $\mathbf{C}^1 \propto (\mathbf{x}^1 + \cdots + \mathbf{x}^{12})$, proportional to the annual temperature
- $\mathbf{C}^2 \propto (\mathbf{x}^5 + \ldots + \mathbf{x}^8) - (\mathbf{x}^1 + \mathbf{x}^2 + \mathbf{x}^{11} + \mathbf{x}^{12})$, contrast summer/winter

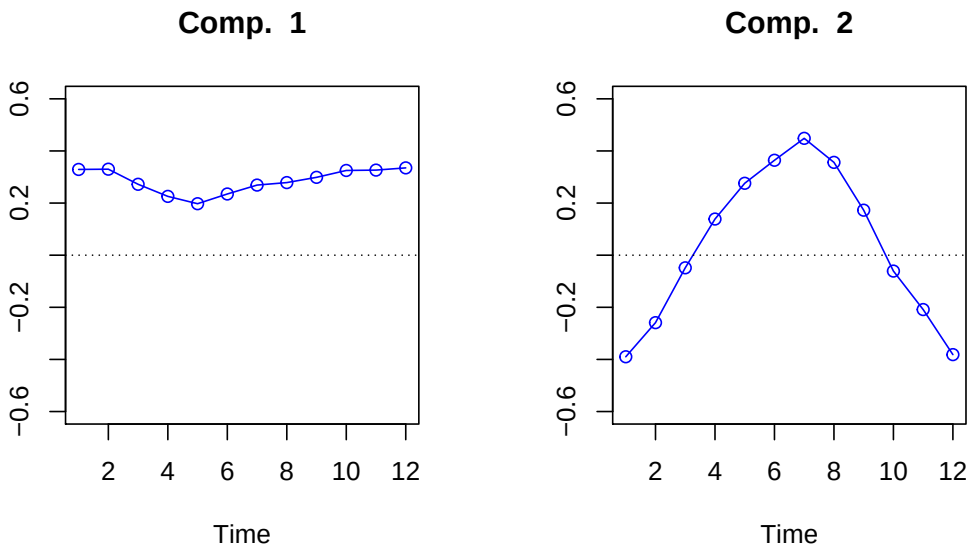

Coordinates of the first 2 eigenvectors in $\mathbb{R}^{12}$.

Graphics for individuals

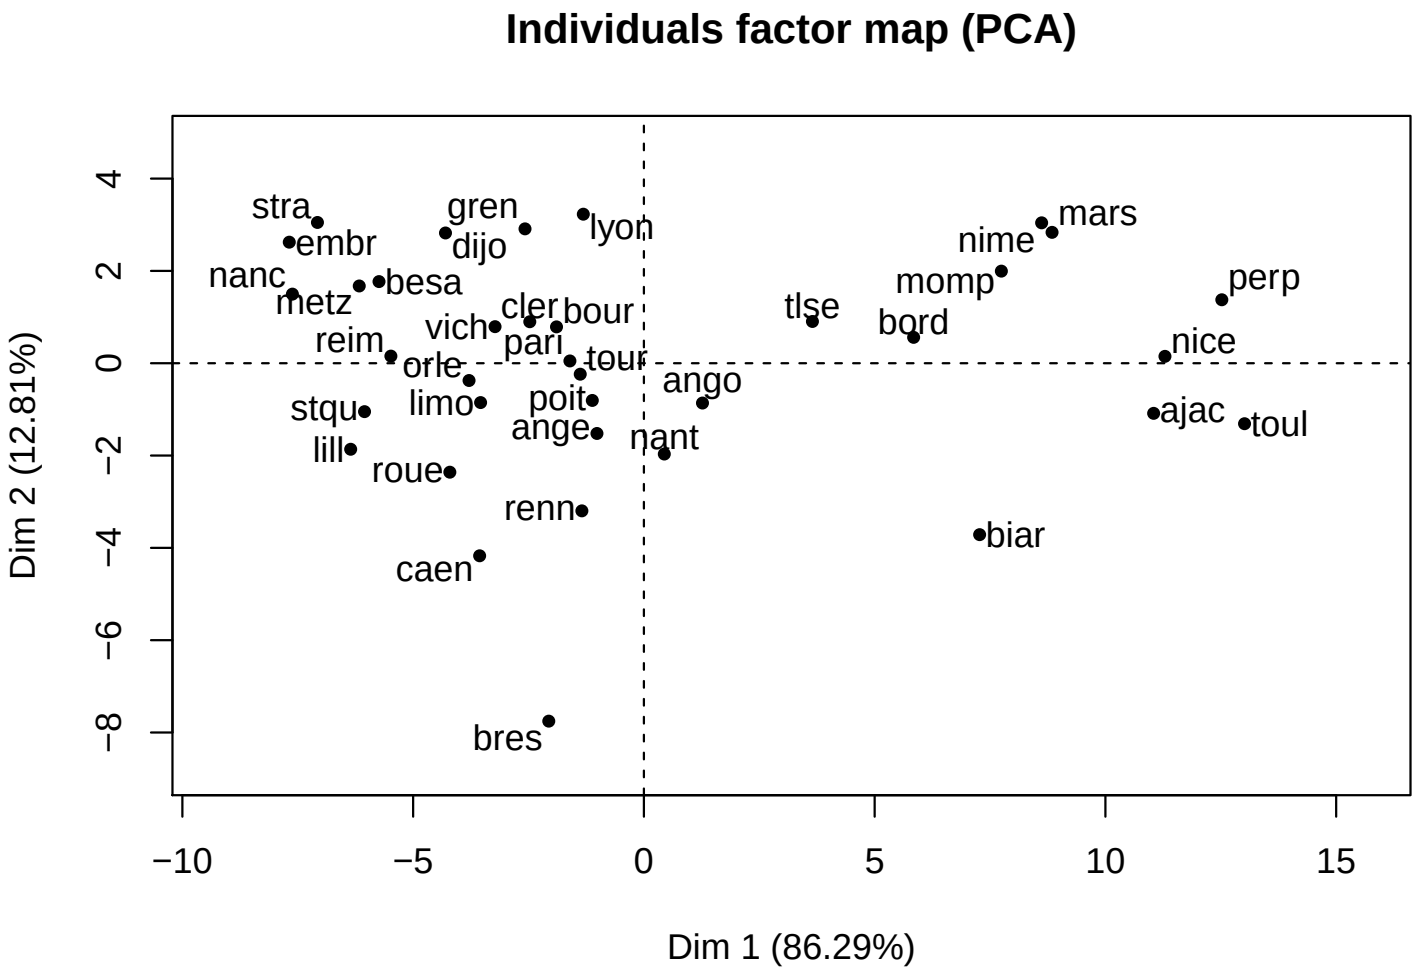

PCA: Projection on the first 2 principal axis.

Graphics for variables

- The principal variables $\mathbf{C}^k$ are orthogonal with variance λ_k. Thus, they define an orthonormal basis $\mathbf{U}^k = \mathbf{C}^k / \sqrt{\lambda_k}$.
- Consider the coordinates $a_{j,k}$ of the original variables in this basis

$$a_{j,k} = \mathrm{cov}(\mathbf{X}^j, \mathbf{U}^k).$$

  We thus have, $\|\mathbf{x}^j\|^2_{\mathbb{R}^n} = \hat{\sigma}_j^2 = \sum_k a_{j,k}^2$.

- The idea is to plot these coordinates for two principal components.

Graphics for variables, case of unit variance

- When the variables have been normalized (unit variance),

$$a_{j,k} = \mathrm{cor}(\mathbf{X}^j, \mathbf{U}^k) = \cos(\widehat{\mathbf{X}^j, \mathbf{U}^k})$$

  and $\sum_{k=1}^{p} a_{j,k}^2 = 1$.

- Thus the coordinates $(a_{j,k})_k$ belong to a p-dimensional sphere.

- Further $(a_{j,1}, a_{j,2})$ belongs to the unit disk: $a_{j,1}^2 + a_{j,2}^2 \leq 1$.
  It is closed to the unit circle if $a_{j,3}, \ldots, a_{j,p}$ are nearly zero.
  In that case, $\mathbf{X}^j$ is well-represented by $\mathbf{C}^1, \mathbf{C}^2$.
  This is the **circle of correlations** for components (1, 2).

Interpretation of principal components

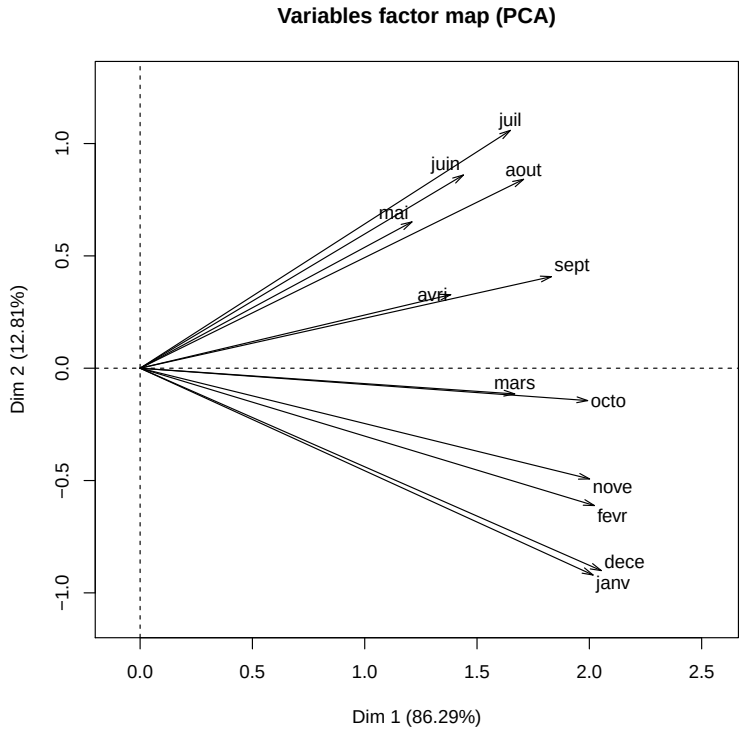

Coordinates of the variables in the orthonormal basis of principal variables. We see again that Axis 1 weigths all months nearly equally, whereas Axis 2 exhibits a contrast summer / winter.

Interpretation of principal components

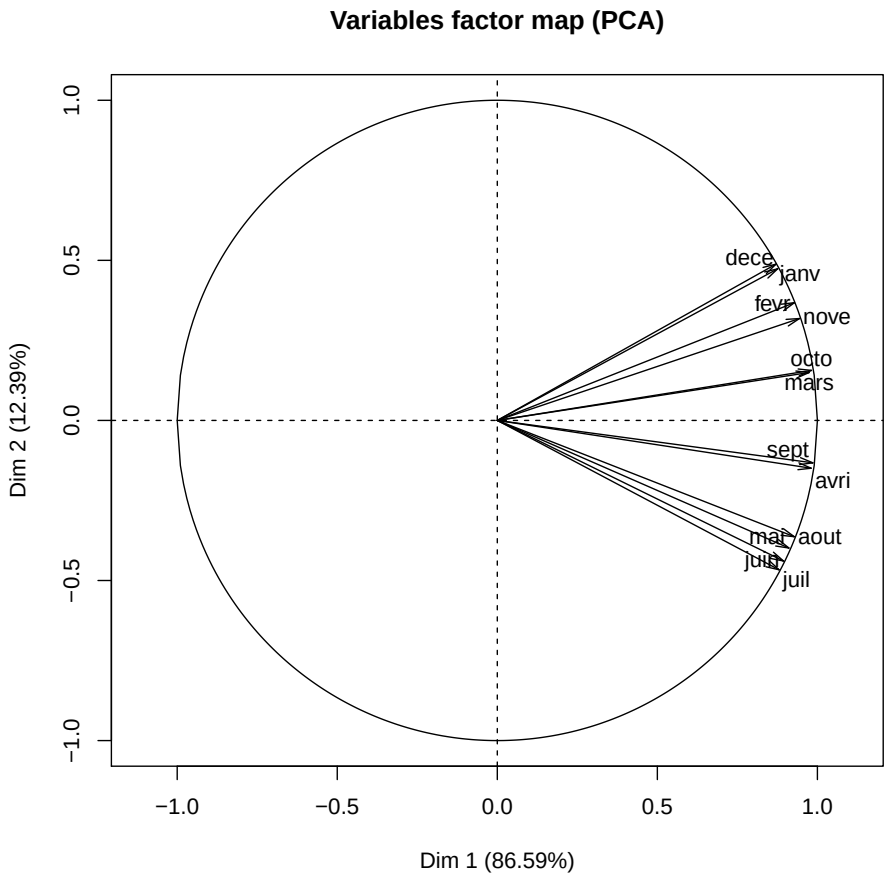

Circle of correlation (normalized variables). Here all variables are well-represented by the first 2 principal components.

Graphics for variables

Exercise

- Check that $\mathbf{U^k}$ is the k-th column of the matrix $\mathbf{U}$ of the SVD decomposition of $\mathbf{X}$ (see slide 21).
- Check that the coordinate of $\mathbf{X}^j$ onto $\mathbf{U}^k$ is $a_{j,k} = (v_k)_j\sqrt{\lambda_k}$. Explain the link between the circle of correlation and the plot of eigenvectors coordinates (slides 31 and 35).

Conclusion and further readings

- PCA is a **dimension reduction technique** which finds **uncorrelated variables, called principal variables**, that are linear combination of the original ones, which approximate the best the data in the **mean-square** sense.
- **PCA = spectral decomposition of the covariance matrix**
  - Up to isometric transformations (metric, weights)
- Several graphs can be used to interpret principal components: projection of individuals, circle of correlation (normalized case).
  - Mind that **what you visualize is only a projection.** Several tools quantify the **quality of the representation.**
    $\rightarrow$ See textbook page 29, 30.