Stationary signals and PCA

DS-GA 1013 / MATH-GA 2824 Mathematical Tools for Data Science

Carlos Fernandez-Granda

Prerequisites

Covariance matrices

Principal component analysis

Discrete Fourier transform

Frequency transformations in multiple dimensions

Linear translation-invariant models and convolution

Motivation: Covariance between pixels in an image

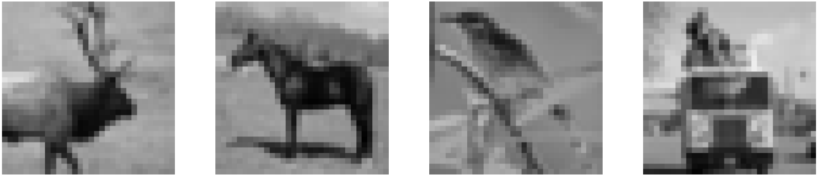

Covariance with pixel i in 16th row

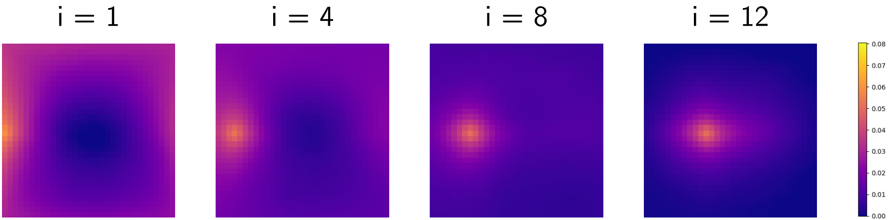

Covariance with pixel i in 16th row

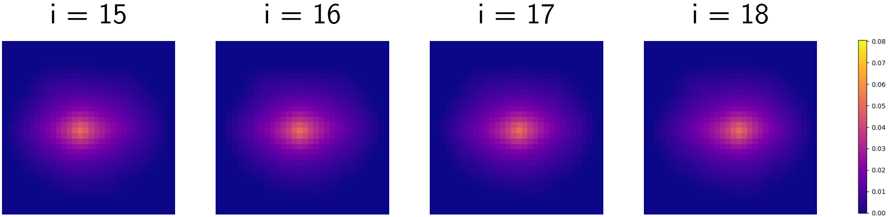

Rows of covariance matrix

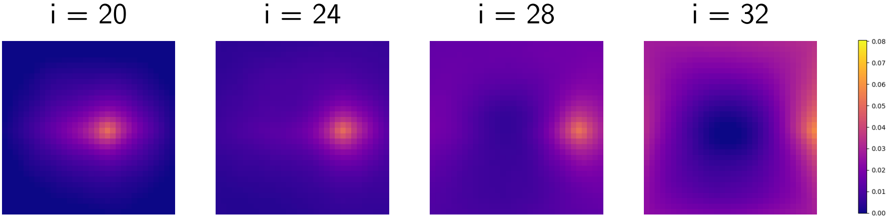

Principal directions

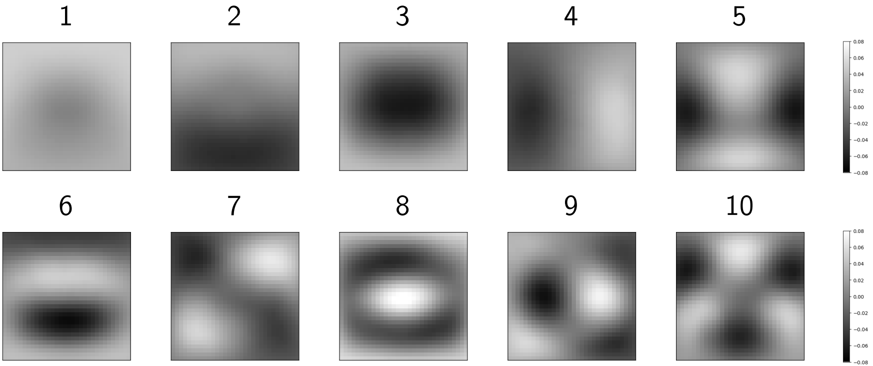

Principal directions

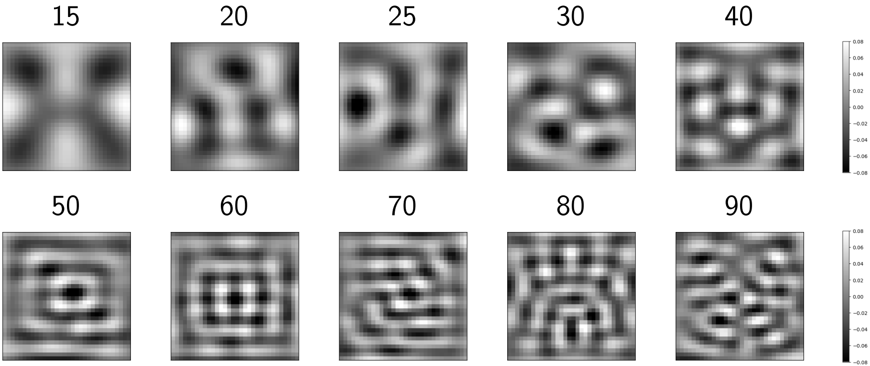

Principal directions

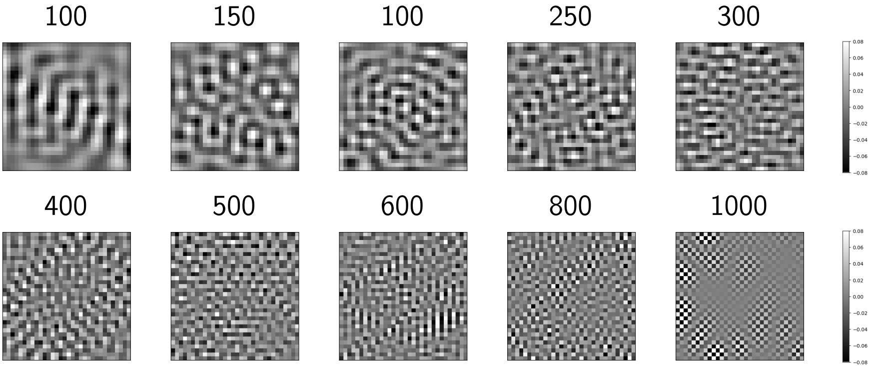

Toy model of 1D piecewise constant signals

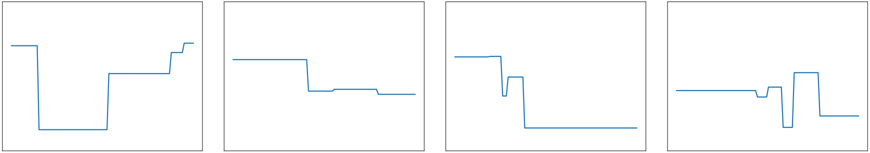

Sample covariance matrix

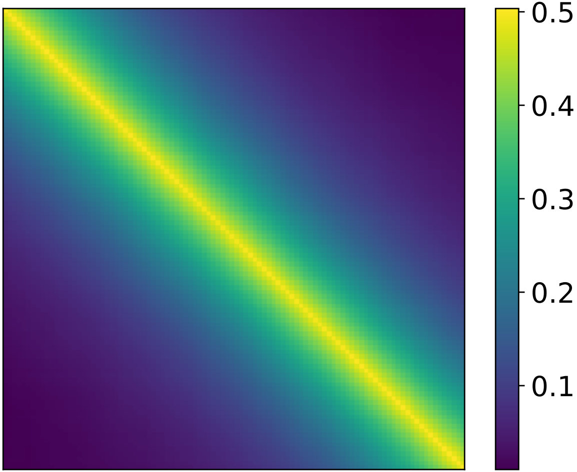

Principal directions

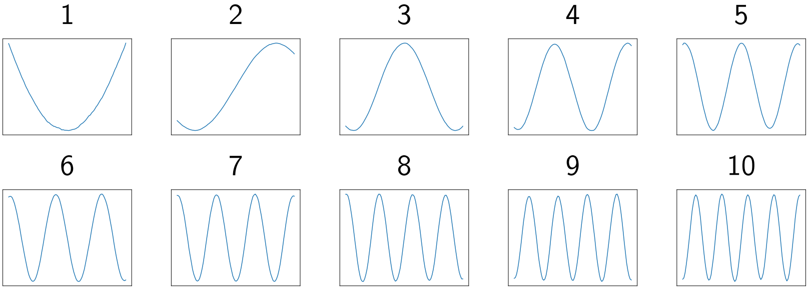

Principal directions

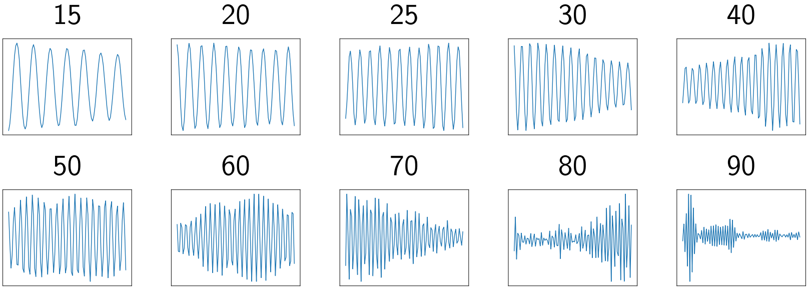

Stationary signals

$\tilde{x}$ is wide-sense or weak-sense stationary if

1. it has a constant mean

$$\mathrm{E}\left(\tilde{x}[j]\right) = \mu, \quad 1 \leq j \leq N$$

2. there is an *autocovariance* function $a_{\tilde{x}}$ such that

$$\mathrm{Cov}\left(\tilde{x}[j_1]\tilde{x}[j_2]\right) = \mathrm{ac}_{\tilde{x}}(j_2 - j_1 \bmod N), \quad 0 \leq j_1, j_2 \leq N-1$$

i.e. it has translation-invariant covariance

Autocovariance

For any j, $\text{ac}_{\tilde{x}}(j) = \text{ac}_{\tilde{x}}(-j) = \text{ac}_{\tilde{x}}(N-j)$

$$\Sigma_{\tilde{x}} = \begin{bmatrix} a_{\tilde{x}} & a_{\tilde{x}}^{\downarrow 1} & a_{\tilde{x}}^{\downarrow 2} & \cdots a_{\tilde{x}}^{\downarrow N-1} \end{bmatrix}$$
$$= \begin{bmatrix} \text{ac}_{\tilde{x}}(0) & \text{ac}_{\tilde{x}}(N-1) & \cdots & \text{ac}_{\tilde{x}}(1) \\ \text{ac}_{\tilde{x}}(1) & \text{ac}_{\tilde{x}}(0) & \cdots & \text{ac}_{\tilde{x}}(2) \\ & & \cdots & \\ \text{ac}_{\tilde{x}}(N-1) & \text{ac}_{\tilde{x}}(N-2) & \cdots & \text{ac}_{\tilde{x}}(0) \end{bmatrix}$$

where

$$a_{\tilde{x}} := \begin{bmatrix} \text{ac}_{\tilde{x}}(0) \\ \text{ac}_{\tilde{x}}(1) \\ \text{ac}_{\tilde{x}}(2) \\ \cdots \end{bmatrix}$$

Circulant matrix

Each column vector is a unit circular shift of previous column

$$\begin{bmatrix} a & d & c & b \\ b & a & d & c \\ c & b & a & d \\ d & c & b & a \end{bmatrix} = \begin{bmatrix} h & h^{\downarrow 1} & h^{\downarrow 2} & h^{\downarrow 3} \end{bmatrix}$$

Sample covariance matrix of piecewise constant signals

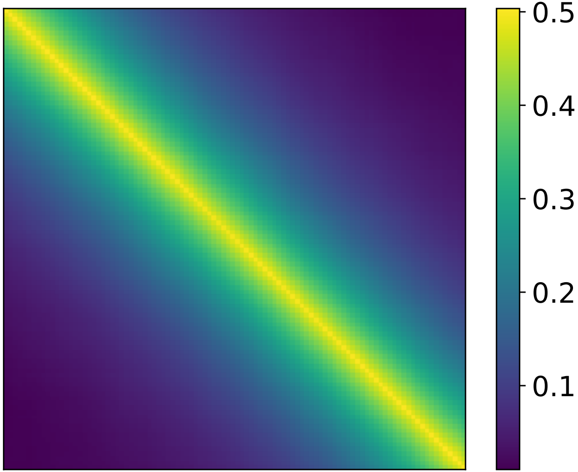

Rows of covariance matrix of image dataset (reshaped)

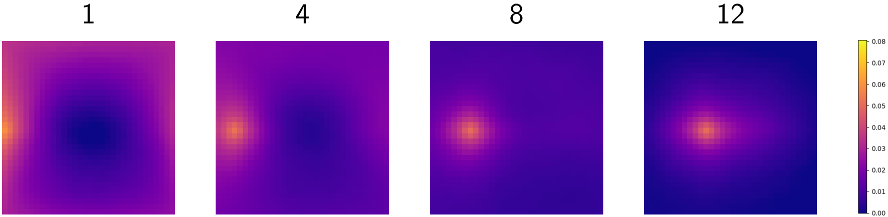

Rows of covariance matrix of image dataset (reshaped)

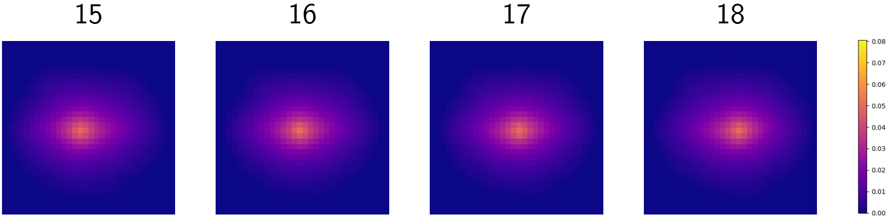

Rows of covariance matrix of image dataset (reshaped)

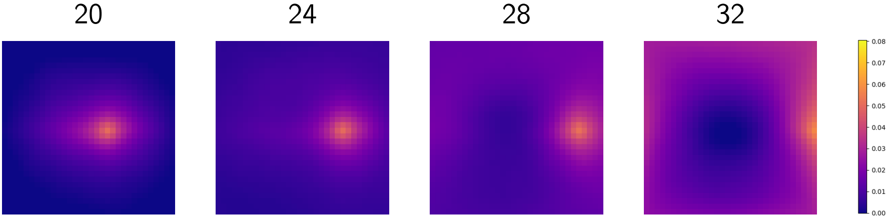

Applying a circulant matrix

Any circulant matrix $C \in \mathbb{C}^{N \times N}$ and any vector $x \in \mathbb{C}^N$

$$\begin{aligned} Cx &= \sum_{s=0}^{N-1} x\,[s]\, h^{\downarrow s} \\ &= h * x \\ &= \frac{1}{N} F_{[N]}^{*}\, \mathrm{diag}(\hat{h}) F_{[N]} x \end{aligned}$$

This is an eigendecomposition!

Eigendecomposition of circulant matrix

$$C := \frac{1}{N} F_{[N]}^{*} \operatorname{diag}(\hat{h}) F_{[N]}$$

where $F_{[N]}$ is the DFT matrix and $\hat{h}$ is the DFT of the first column

PCA of stationary signals

Let $\tilde{x}$ be wide-sense stationary with autocovariance vector $a_{\tilde{x}}$

The eigendecomposition of the covariance matrix of $\tilde{x}$ equals

$$\begin{aligned}\Sigma_{\tilde{x}} &= \begin{bmatrix} a_{\tilde{x}} & a_{\tilde{x}}^{\downarrow 1} & a_{\tilde{x}}^{\downarrow 2} & \cdots a_{\tilde{x}}^{\downarrow N-1} \end{bmatrix} \\ &= \frac{1}{N} F^{*} \operatorname{diag}(\hat{a}_{\tilde{x}}) F\end{aligned}$$

Toy model of 1D piecewise constant signals

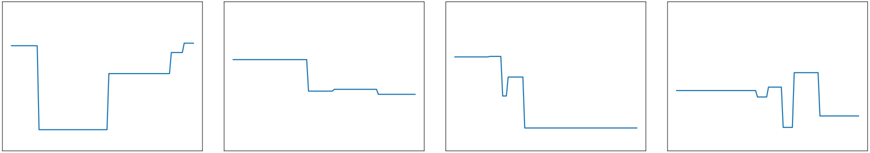

Principal directions

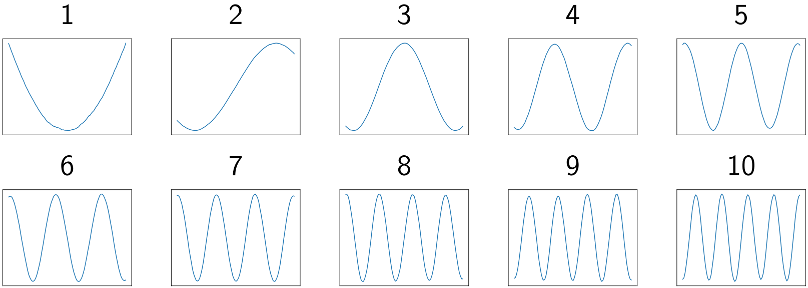

Principal directions

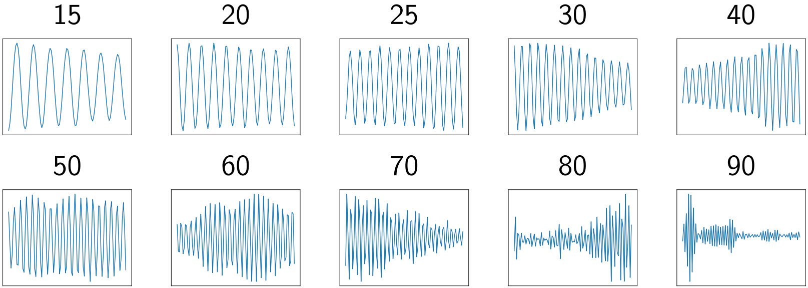

CIFAR-10 images

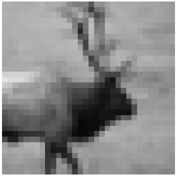 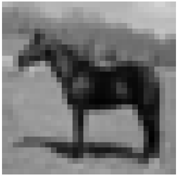 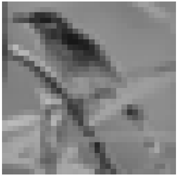 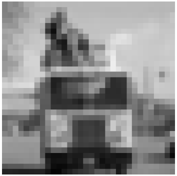

Principal directions

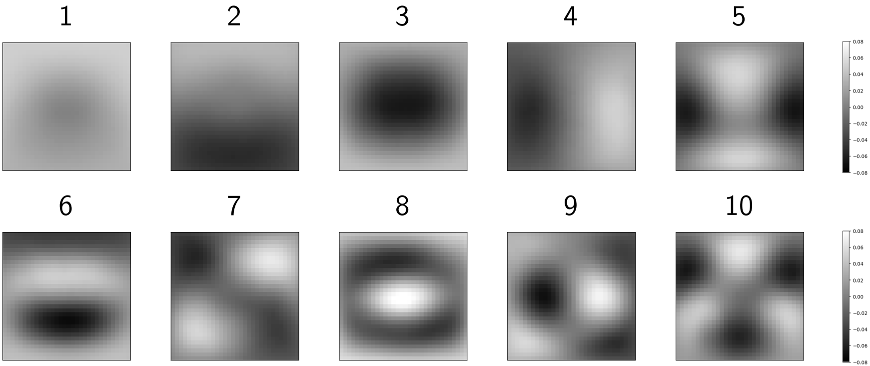

Principal directions

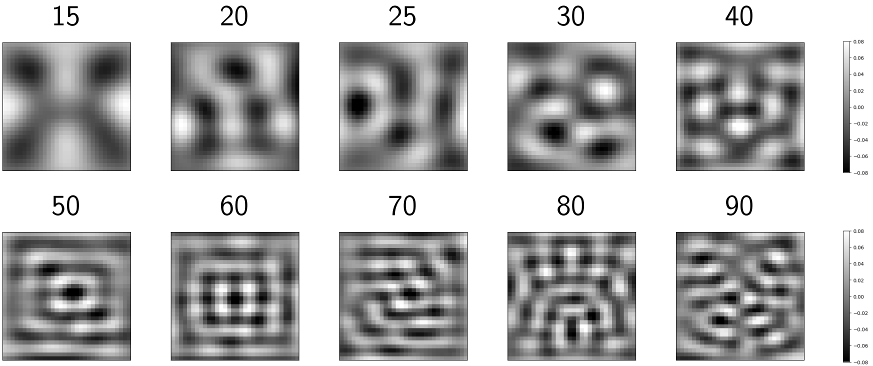

Principal directions

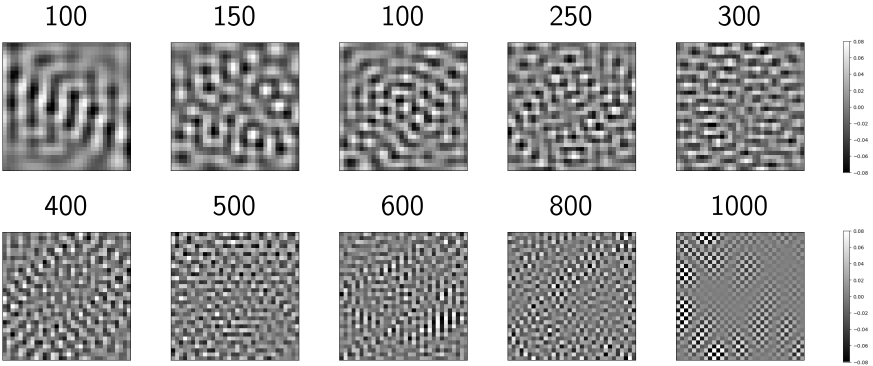

PCA of natural images

Principal directions tend to be sinusoidal

This suggests using 2D sinusoids for dimensionality reduction

JPEG compresses images using discrete cosine transform (DCT):

1. Image is divided into 8×8 patches
2. Each DCT band is quantized differently (more bits for lower frequencies)

DCT basis vectors

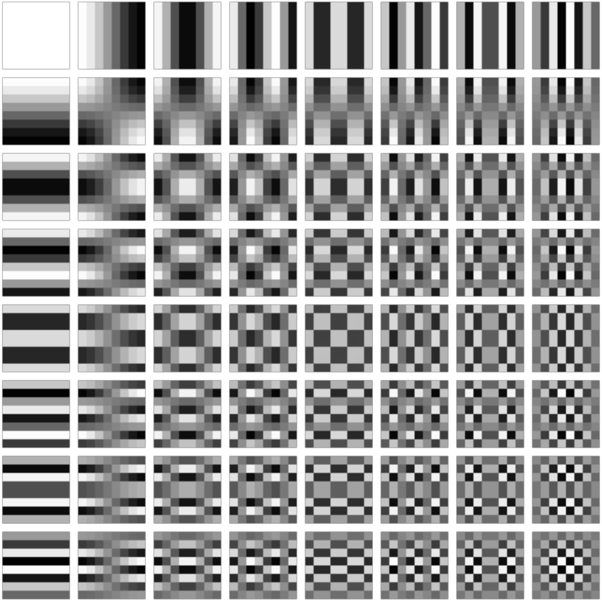

Projection of each 8x8 block onto first DCT coefficients

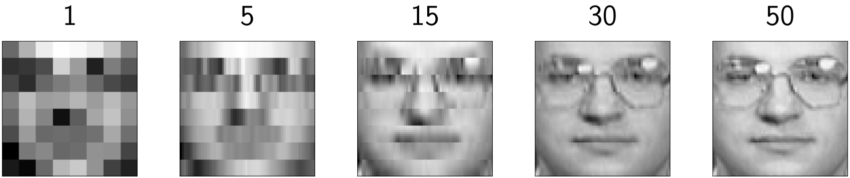

What have we learned?

Stationary signals have translation-invariant statistics

Their covariance matrix are circulant

Circulant matrices have sinusoidal eigendecompositions

PCA on stationary signals yields sinusoidal principal directions