Principal Component Analysis (PCA)

- Quadratic Forms
- Multivariate Gaussian Density
- Inner and Outer Products
- Covariance Matrix
- Isodensity Surfaces
- Principal Components Theorem
- Diagonalizing the Covariance Matrix
- KL Transform
- Reducing Dimensionality

Quadratic Forms

Let $f(\mathbf{x}) = \mathbf{x}^{\mathrm{T}}\mathbf{A}\mathbf{x}$ where $\mathbf{A} = \mathbf{A}^{\mathrm{T}}$. In two-dimensions, we have

$$\mathbf{A} = \begin{bmatrix} a & b \\ b & c \end{bmatrix} \quad \text{and} \quad \mathbf{x} = \begin{bmatrix} x & y \end{bmatrix}^{\mathrm{T}}$$

so that

$$\mathbf{A}\mathbf{x} = \begin{bmatrix} a & b \\ b & c \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} ax + by \\ bx + cy \end{bmatrix}$$

and

$$\mathbf{x}^{\mathrm{T}}\mathbf{A}\mathbf{x} = \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} ax + by \\ bx + cy \end{bmatrix} = ax^2 + 2bxy + cy^2.$$

Quadratic Forms (contd.)

When $\mathbf{A}$ is positive definite, then

$$f(\mathbf{x}) = \mathbf{x}^{\mathrm{T}}\mathbf{A}\mathbf{x}$$

is a *paraboloid* and the *isovalue contours*,

$$\mathbf{x}^{\mathrm{T}}\mathbf{A}\mathbf{x} = D$$

are *ellipses*. A matrix is positive definite iff all of its eigenvalues are positive.

Example

If $\mathbf{A} = \begin{bmatrix} 5 & -2 \\ -2 & 5 \end{bmatrix}$ then $\mathbf{x}^{\mathrm{T}}\mathbf{A}\mathbf{x}$ equals

$$5x^2 - 4xy + 5y^2.$$

The eigenvalues of $\mathbf{A}$ are 3 and 7 and the corresponding eigenvectors are $\mathbf{u} = \begin{bmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix}$ and $\mathbf{v} = \begin{bmatrix} -1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix}$. Now, let $\mathbf{A} = \mathbf{U}\mathbf{B}\mathbf{U}^{\mathrm{T}}$, where $\mathbf{U}$ equals

$$\mathbf{U} = \begin{bmatrix} 1/\sqrt{2} & -1/\sqrt{2} \\ 1/\sqrt{2} & 1/\sqrt{2} \end{bmatrix},$$

then $\mathbf{B} = \mathbf{U}^{\mathrm{T}}\mathbf{A}\mathbf{U}$, which is

$$\mathbf{B} = \begin{bmatrix} 3 & 0 \\ 0 & 7 \end{bmatrix}.$$

The corresponding quadratic form, $\mathbf{u}^{\mathrm{T}}\mathbf{B}\mathbf{u}$, is

$$3u^2 + 7v^2.$$

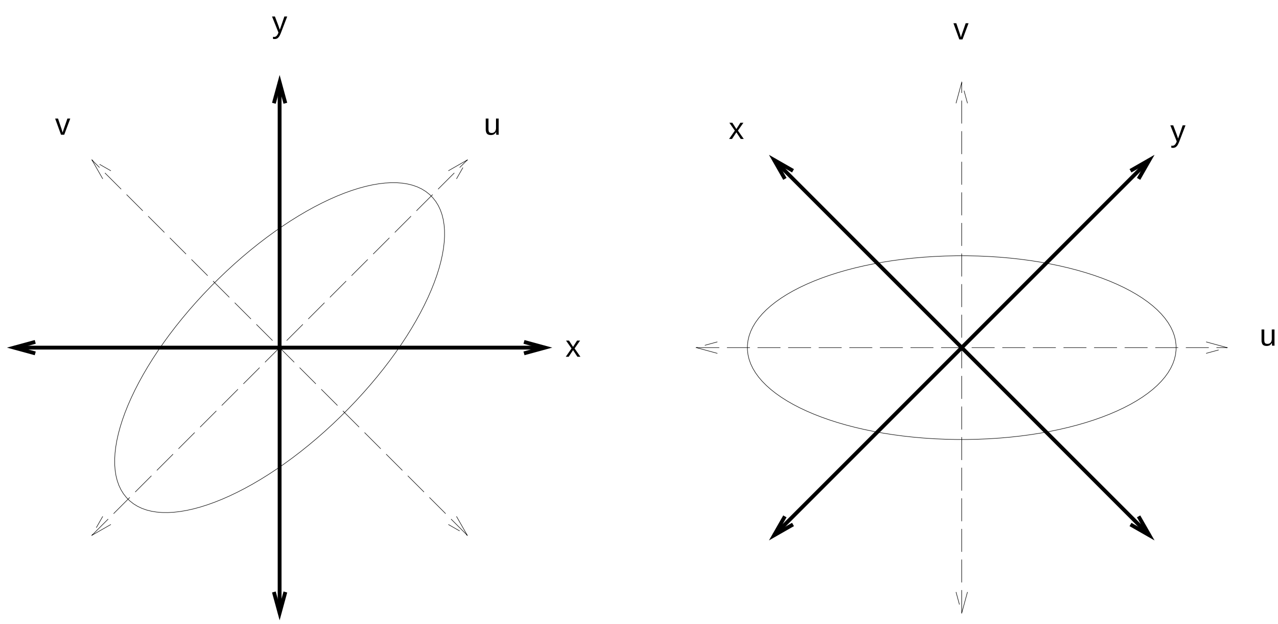

Figure 1: Left: The ellipse, $5x^2 - 4xy + 5y^2 = D$. Right: The ellipse, $3u^2 + 7v^2 = D$.

Multivariate Gaussian Density

The *multivariate Gaussian density* is defined as follows:

$$G(\mathbf{x}) = \frac{1}{(2\pi)^{\frac{K}{2}}|\mathbf{C}|^{\frac{1}{2}}} e^{-\frac{1}{2}\mathbf{x}^{\mathrm{T}}\mathbf{C}^{-1}\mathbf{x}}$$

where K is the number of dimensions and $\mathbf{C}$ is the $K \times K$ *covariance matrix*. In the *bivariate* case, $\mathbf{C}$ looks like this:

$$\mathbf{C} = \begin{bmatrix} \sigma_{xx} & \sigma_{xy} \\ \sigma_{xy} & \sigma_{yy} \end{bmatrix}.$$

Note: If $\mathbf{C}$ is symmetric and positive definite, then $\mathbf{C}^{-1}$ is also symmetric and positive definite.

Inner and Outer Products

Let $\mathbf{x} = \begin{bmatrix} 1 & 2 & 3 \end{bmatrix}^{\mathrm{T}}$. The *inner product* of $\mathbf{x}$ with itself, or $\mathbf{x}^{\mathrm{T}}\mathbf{x}$ is a scalar:

$$\begin{bmatrix} 1 & 2 & 3 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} = 1 \cdot 1 + 2 \cdot 2 + 3 \cdot 3 = 14.$$

The *outer product* of $\mathbf{x}$ with itself, or $\mathbf{x}\mathbf{x}^{\mathrm{T}}$ is a matrix:

$$\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \begin{bmatrix} 1 & 2 & 3 \end{bmatrix} = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 4 & 6 \\ 3 & 6 & 9 \end{bmatrix}.$$

Covariance Matrix

First we construct an $N \times K$ matrix, $\mathbf{X}$, where the n-th row is the n-th sample of a multivariate Gaussian r.v., $\mathbf{x} = \begin{bmatrix} x & y \end{bmatrix}^{\mathrm{T}}$. For example, when $K = 2$:

$$\mathbf{X} = \begin{bmatrix} x_1 & y_1 \\ x_2 & y_2 \\ \vdots & \vdots \\ x_N & y_N \end{bmatrix}.$$

The *sample mean* of the N samples is

$$\vec{\mu} = \frac{1}{N} \sum_{n=1}^{N} \mathbf{x}_n.$$

We will assume that $\vec{\mu} = \begin{bmatrix} 0 & 0 \end{bmatrix}^{\mathrm{T}}$. If this is false, we can always make it true by subtracting $\vec{\mu}$ from each of the samples prior to constructing $\mathbf{X}$.

Covariance Matrix (contd.)

We observe that

$$\begin{aligned}\mathbf{X}^{\mathrm{T}}\mathbf{X} &= \sum_{n=1}^{N} \mathbf{x}_n \mathbf{x}_n^{\mathrm{T}} \\ &= \begin{bmatrix} x_1x_1 & x_1y_1 \\ x_1y_1 & y_1y_1 \end{bmatrix} + \cdots + \begin{bmatrix} x_Nx_N & x_Ny_N \\ x_Ny_N & y_Ny_N \end{bmatrix}.\end{aligned}$$

The *covariance matrix* is the matrix of the expected values of the products of the x and y components of the samples:

$$\mathbf{C} = \frac{1}{N}\mathbf{X}^{\mathrm{T}}\mathbf{X} = \begin{bmatrix} \langle xx \rangle & \langle xy \rangle \\ \langle xy \rangle & \langle yy \rangle \end{bmatrix} = \begin{bmatrix} \sigma_{xx} & \sigma_{xy} \\ \sigma_{xy} & \sigma_{yy} \end{bmatrix}$$

where $\langle . \rangle$ denotes expected value.

Isodensity Surfaces

The *isodensity surfaces* of the multivariate Gaussian are the locus of those points where $G(\mathbf{x})$ has constant density:

$$G(\mathbf{x}) = \frac{1}{(2\pi)^{\frac{K}{2}}|\mathbf{C}|^{\frac{1}{2}}} e^{-\frac{1}{2}\mathbf{x}^{\mathrm{T}}\mathbf{C}^{-1}\mathbf{x}} = D$$

which can be re-arranged to yield:

$$\mathbf{x}^{\mathrm{T}}\mathbf{C}^{-1}\mathbf{x} = -2\ln\left[(2\pi)^{\frac{K}{2}}|\mathbf{C}|^{\frac{1}{2}}D\right].$$

Since $\mathbf{C}^{-1}$ is positive definite the isodensity surfaces are *ellipsoids*. The *axes* of these ellipsoids are mutually orthogonal and point in the same directions as the eigenvectors of $\mathbf{C}$. These eigenvectors are the *principal components* of the multivariate Gaussian density.

Principal Components Theorem

The principal components of a multivariate Gaussian density are given by the eigenvectors of its covariance matrix.

Proof (in two-dimensions): We observe that

$$e^{-\frac{1}{2}\mathbf{x}^{\mathrm{T}}\mathbf{C}^{-1}\mathbf{x}}$$

is maximized (or minimized) when $\mathbf{x}^{\mathrm{T}}\mathbf{C}\mathbf{x}$ is maximized (or minimized). We therefore wish to find the unit vectors $\mathbf{x}$ which maximize (or minimize):

$$\begin{aligned}
\mathbf{x}^{\mathrm{T}}\mathbf{C}\mathbf{x} &= \mathbf{x}^{\mathrm{T}}\mathbf{C}^{\frac{1}{2}}\mathbf{C}^{\frac{1}{2}}\mathbf{x} \\
&= \mathbf{x}^{\mathrm{T}}\left(\mathbf{C}^{\frac{1}{2}}\right)^{\mathrm{T}}\mathbf{C}^{\frac{1}{2}}\mathbf{x} \\
&= \left(\mathbf{C}^{\frac{1}{2}}\mathbf{x}\right)^{\mathrm{T}}\mathbf{C}^{\frac{1}{2}}\mathbf{x} \\
&= ||\mathbf{C}^{\frac{1}{2}}\mathbf{x}||^{2}
\end{aligned}$$

where $\mathbf{C}$ is symmetric and positive definite.

Principal Components Theorem (contd.)

Let $\mathbf{w}_1$ and $\mathbf{w}_2$ be eigenvectors of $\mathbf{C}$ with eigenvalues λ_1 and λ_2. Note that $\mathbf{w}_1$ and $\mathbf{w}_2$ are also eigenvectors of $\mathbf{C}^{\frac{1}{2}}$ but its eigenvalues are $\sqrt{\lambda_1}$ and $\sqrt{\lambda_2}$. Now consider a unit vector, $\mathbf{x}$, in the plane. Let θ be the relative orientation between $\mathbf{x}$ and $\mathbf{w}_1$. It follows that

$$[\mathbf{x}]_{\mathcal{W}} = \begin{bmatrix} \cos\theta \\ \sin\theta \end{bmatrix}$$

is the representation of $\mathbf{x}$ in the basis defined by $\mathbf{w}_1$ and $\mathbf{w}_2$. Consequently,

$$\begin{aligned} \left|\left| \left[\mathbf{C}^{\frac{1}{2}}\mathbf{x}\right]_{\mathcal{W}} \right|\right|^2 &= \left(\sqrt{\lambda_1}\right)^2 \cos^2\theta + \left(\sqrt{\lambda_2}\right)^2 \sin^2\theta \\ &= \lambda_1 \cos^2\theta + \lambda_2 \sin^2\theta. \end{aligned}$$

Principal Components Theorem (contd.)

Calculus tells us that $\left|\left|\left[\mathbf{C}^{\frac{1}{2}}\mathbf{x}\right]_{\mathcal{W}}\right|\right|^2$ is maximized (or minimized) when

$$\frac{d\left(\lambda_1\cos^2\theta+\lambda_2\sin^2\theta\right)}{d\theta}=0.$$

Evaluating the above derivative:

$$2\lambda_1\cos\theta\sin\theta-2\lambda_2\sin\theta\cos\theta \;=\; 2\cos\theta\sin\theta\left(\lambda_1-\lambda_2\right).$$

It follows that $\left|\left|\left[\mathbf{C}^{\frac{1}{2}}\mathbf{x}\right]_{\mathcal{W}}\right|\right|^2$ is maximized (or minimized) when $\theta=0$ (or $\theta=\pi/2$), *i.e.*, when $\mathbf{x}=\mathbf{w_1}$ (or $\mathbf{x}=\mathbf{w_2}$). Now, because $\mathcal{W}$ is orthonormal

$$\left|\left|\left[\mathbf{C}^{\frac{1}{2}}\mathbf{x}\right]_{\mathcal{W}}\right|\right|^2=||\mathbf{C}^{\frac{1}{2}}\mathbf{x}||^2$$

and because

$$||\mathbf{C}^{\frac{1}{2}}\mathbf{x}||^2=\mathbf{x}^{\mathrm{T}}\mathbf{C}\mathbf{x}$$

we conclude that $\mathbf{x}^{\mathrm{T}}\mathbf{C}\mathbf{x}$ is maximized (or minimized) when $\mathbf{x}$ is an eigenvector of $\mathbf{C}$.

Diagonalizing the Covariance Matrix

Because the covariance matrix $\mathbf{C}$ is symmetric and positive definite, it has K orthogonal eigenvectors:

$$\lambda_k \mathbf{w}_k = \mathbf{C}\mathbf{w}_k$$

where $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_K$. It can therefore be diagonalized as follows:

$$\mathbf{C} = \mathbf{W}\mathbf{D}\mathbf{W}^{\mathrm{T}}$$

where $\mathbf{W}$ is a $K \times K$ matrix of eigenvectors:

$$\mathbf{W} = \left[\, \mathbf{w}_1 \middle| \mathbf{w}_\mathbf{2} \middle| \ldots \middle| \mathbf{w}_K \,\right]$$

and $\mathbf{D}$ is a $K \times K$ diagonal matrix of eigenvalues:

$$\mathbf{D} = \mathrm{diag}(\lambda_1, \lambda_2, \ldots, \lambda_K).$$

<u>The KL Transform</u>

We can represent a sample $\mathbf{x}$ of a multivariate Gaussian r.v. with covariance matrix $\mathbf{C}$ in the basis $\mathcal{W}$ formed by $\mathbf{C}$'s eigenvectors. This change of basis is termed the *Karhunen-Loeve transform*:

$$[\mathbf{x}]_{\mathcal{W}} = \mathbf{W}^{\mathrm{T}}\mathbf{x}.$$

Because $\mathbf{C}$ is symmetric, the $\mathbf{w}_k$ are mutually orthogonal, and $\mathbf{W}^{\mathrm{T}}$ is unitary. Consequently, the KL transform (like the DFT) is a rotation in $\mathbb{R}^K$.

The KL Transform (contd.)

- **Question** Let $\mathbf{u} = [\mathbf{x}]_{\mathcal{W}}$ be the representation of $\mathbf{x}$ in the basis $\mathcal{W}$ formed by the eigenvectors of $\mathbf{C}$. What is the density of $\mathbf{u}$?

- **Answer** It is the multivariate Gaussian density with covariance matrix, $\mathbf{D}$:

$$G'(\mathbf{u}) = \frac{1}{(2\pi)^{\frac{K}{2}}|\mathbf{D}|^{\frac{1}{2}}} e^{-\frac{1}{2}\mathbf{u}^{\mathrm{T}}\mathbf{D}^{-1}\mathbf{u}}$$

where $\mathbf{D} = \mathbf{W}^{\mathrm{T}}\mathbf{C}\mathbf{W}$.

The Bivariate Case

In the bivariate case

$$\mathbf{D} = \mathbf{W}^{\mathrm{T}}\mathbf{C}\mathbf{W} = \begin{bmatrix} \sigma_{uu} & 0 \\ 0 & \sigma_{vv} \end{bmatrix}.$$

Since $\mathbf{D}$ is diagonal,

$$|\mathbf{D}| = \sigma_{uu}\sigma_{vv}$$

and $\mathbf{D}^{-1}$ has an especially simple form:

$$\mathbf{D}^{-1} = \begin{bmatrix} 1/\sigma_{uu} & 0 \\ 0 & 1/\sigma_{vv} \end{bmatrix}.$$

<u>The Bivariate Case (contd.)</u>

It follows that the multivariate Gaussian density with covariance matrix $\begin{bmatrix} \sigma_{uu} & 0 \\ 0 & \sigma_{vv} \end{bmatrix}$ is:

$$G'(u,v) = \frac{1}{2\pi\sqrt{\sigma_{uu}\sigma_{vv}}} e^{-\frac{1}{2}\left(\frac{u^2}{\sigma_{uu}} + \frac{v^2}{\sigma_{vv}}\right)}.$$

We observe that G' is *separable*:

$$G'(u,v) = \frac{1}{\sqrt{2\pi\sigma_{uu}}} e^{-\frac{u^2}{2\sigma_{uu}}} \frac{1}{\sqrt{2\pi\sigma_{vv}}} e^{-\frac{v^2}{2\sigma_{vv}}}.$$

Since the joint density function of u and v can be expressed as the product of the density function for u and the density function for v, we say that u and v are *uncorrelated.* Stated differently, knowing the value of u tells you nothing about the value of v!

Reducing Dimensionality

Since $\mathbf{W}^{\mathrm{T}}$ is unitary, its inverse is simply $\mathbf{W}$. Consequently, the KL transform can be inverted as follows:

$$\mathbf{x} = \mathbf{W}\mathbf{u}$$

which (in the general case of K dimensions) is simply:

$$\mathbf{x} = u_1\mathbf{w}_1 + u_2\mathbf{w}_2 + \cdots + u_K\mathbf{w}_K.$$

Let $\mathbf{u}'$ be a vector of length $J \leq K$ consisting of the components of $\mathbf{u}$ in the directions of eigenvectors with the J eigenvalues of largest magnitude. It is possible to recover, $\mathbf{x}'$, an approximation to $\mathbf{x}$, from $\mathbf{u}'$ as follows:

$$\mathbf{x}' = u_1\mathbf{w}_1 + u_2\mathbf{w}_2 + \cdots + u_J\mathbf{w}_J.$$