Principal Component Analysis

Gerald Farin, Dianne Hansford

December 13, 2005

1 Introduction

Suppose we are given a set of 2D points $\mathbf{p}_1, \ldots, \mathbf{p}_L$. It is often important to bound the points by a bounding box. The "standard" bounding box is shown in the left of Fig. 1. This bounding box does not hug the points very closely. If we could find a box in more general position, it could give a much tighter fit to the points; such a box is in the right of the same figure.

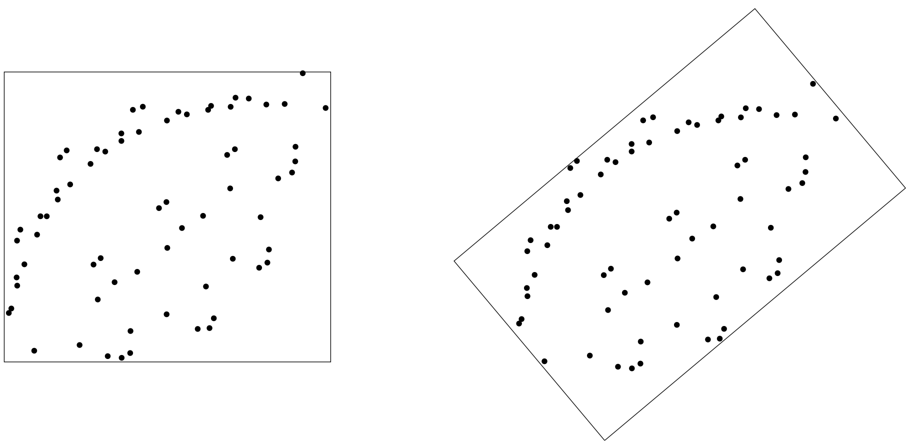

Figure 1: The same set of points with standard bounding box (left), and rotated bounding box (right).

Finding the correct orientation for the box is the main problem. It is solved using Principal Component Analysis (PCA). Before we can explain that method, a few thoughts on quadratic forms.

2 Quadratic Forms

Let a function f have 2D vectors $\mathbf{v}$ as its arguments and let it be defined in terms of a symmetric matrix $C = A^{\mathrm{T}}A$:

$$f(\mathbf{v}) = [v_1, v_2]C \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = \mathbf{v}^{\mathrm{T}}C\mathbf{v}. \tag{1}$$

Such functions are called quadratic forms. The graph of a quadratic form is a point set $[v_1, v_2, f(v_1, v_2)]^{\mathrm{T}}$, forming a surface.

Next, we ask for which vectors $\mathbf{v}$ will f be constant, i.e., for which $\mathbf{v}$ do we have $f(\mathbf{v}) = c$ for some constant c. These vectors form an *ellipse:* it is defined by

$$c = [v_1, v_2]C \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = \mathbf{v}^{\mathrm{T}}C\mathbf{v}. \tag{2}$$

As we vary the constant c, the resulting ellipses all have the same shape and they are scaled versions of each other. An example is given in Fig. 2.

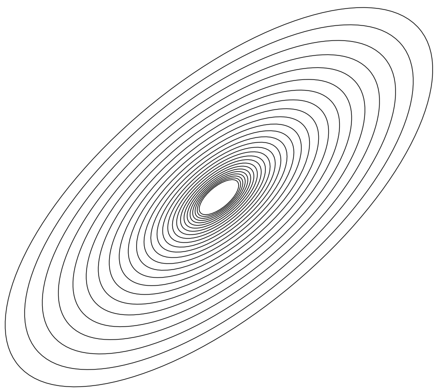

Figure 2: A set of ellipses.

The axes of all these ellipses are given by the (orthogonal) eigenvectors $\mathbf{r}_1$ and $\mathbf{r}_2$ of C. The corresponding eigenvalues are positive numbers λ_1, λ_2 (recall $C = A^{\mathrm{T}}A$ for some A) and they determine the relative lengths of the axes. We assume $\lambda_1 \geq \lambda_2$.

In order to further investigate the shape of (2), perform a rotation which aligns $\mathbf{r}_1$ and $\mathbf{r}_2$ with the axes $\mathbf{e}_1$ and $\mathbf{e}_2$. Such a rotation changes C to a diagonal matrix, with diagonal entries λ_1, λ_2.

Let us now restrict the vectors $\mathbf{v}$ to be of unit length, i.e., they are of the form $\mathbf{v} = [\cos t, \sin t]^{\mathrm{T}}$ for real numbers t. With other words, these vectors trace out the unit circle. The restriction of f to the unit circle is

$$f(\cos t, \sin t) = \lambda_1 \cos^2 t + \lambda_2 \sin^2 t.$$

You may visualize this as a curve which "floats" over the unit circle.

We now ask for which value of t this expression is minimal or maximal. Since both $\cos^2 t$ and $\sin^2 t$ vary between 0 and 1, we see that the minimum is obtained for $t = \pi/2$, since then the function value is λ_2. Similarly, the maximum is attained for $t = 0$ with function value λ_1.

Thus the min and max of f, when restricted to the unit circle, are attained at the endpoints of $C's$ two eigenvectors.

3 Principal Component Analysis

We now return to the problem of finding a good bounding box for a 2D point set.

Let $\mathbf{u}_1, \ldots, \mathbf{u}_L$ a set of 2D vectors such that $\sum_i \mathbf{u}_i = \mathbf{0}$. This condition simply means that the origin is the centroid of the vectors. Let $\mathbf{d}$ be a unit vector. Then the projection of a vector $\mathbf{u}_i$ onto $\mathbf{d}$[1] is a vector with length $\mathbf{d}^{\mathrm{T}}\mathbf{u}$. Let $l(\mathbf{d})$ be the square sum of all these lengths:

$$l(\mathbf{d}) = [\mathbf{d}^{\mathrm{T}}\mathbf{u}_1]^2 + \ldots + [\mathbf{d}^{\mathrm{T}}\mathbf{u}_L]^2.$$

The meaning of $l(\mathbf{d})$ is as follows. Imagine rotating $\mathbf{d}$ around the origin. For each position of $\mathbf{d}$, we compute the value $l(\mathbf{d})$. To be more concrete, consider the set of vectors in Fig. 3. For the shown line (right in the figure), the value of $l(\mathbf{d})$ is large; for some $\mathbf{d}$ orthogonal to it, $l(\mathbf{d})$ will be small. Thus the value of $l(\mathbf{d})$ indicates directions along which the point set does or does not line up.

We rewrite $l(\mathbf{d})$ as

$$l(\mathbf{d}) = \|\mathbf{d}^{\mathrm{T}}[\mathbf{u}_1, \ldots, \mathbf{u}_L]\|^2$$

and, using the abbreviation $U = [\mathbf{u}_1, \ldots, \mathbf{u}_L]$,

$$l(\mathbf{d}) = (\mathbf{d}^{\mathrm{T}}U)(\mathbf{d}^{\mathrm{T}}U) = \mathbf{d}^{\mathrm{T}}UU^{\mathrm{T}}\mathbf{d}. \tag{3}$$

We further abbreviate $C = UU^{\mathrm{T}}$ and note that C is a symmetric 2×2 matrix. Hence (3) describes a quadratic form just as we discussed in Section 2.

For which $\mathbf{d}$ is $l(\mathbf{d})$ maximal? The answer is simple: for $\mathbf{d}$ being the eigenvector corresponding to C's largest eigenvalue! Similarly, $l(\mathbf{d})$ is minimal for $\mathbf{d}$ being the

[1] More precisely, onto a line through the origin containing $\mathbf{d}$.

eigenvector corresponding to C's smallest eigenvalue. This follows directly from Section 2.

For a geometric application, we return to the rotated bounding box problem of Fig. 1. We find the centroid of the points and translate the whole point set such that the centroid is moved to the origin. We may now interpret our translated points as vectors $\mathbf{u}_i$ above. The axes of the rotated bounding box are the eigenvectors of C. The edge lengths of the box are λ_1 and λ_2.

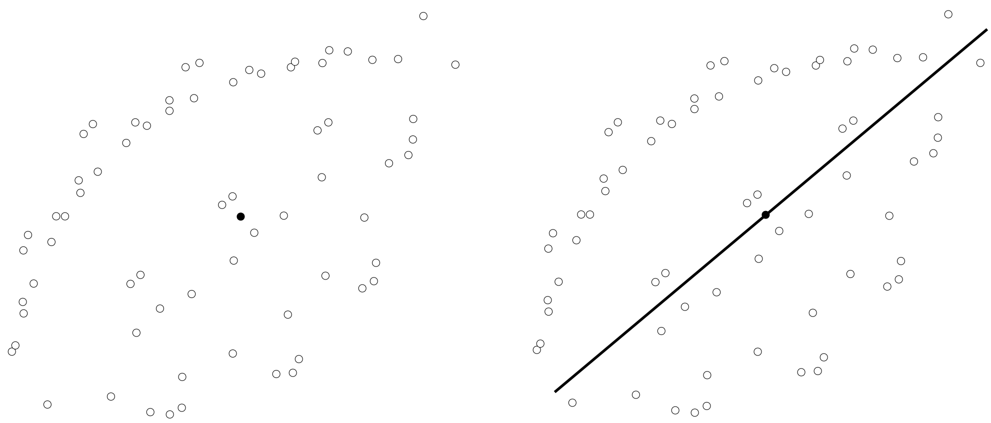

Figure 3: Left: a set of points with their centroid (solid). Right: with the line maximizing $l(\mathbf{d})$.

If $\lambda_1 = \lambda_2$, then there is no preferred direction in the point set: our ellipse is a circle.

The generalization to dimensions higher than 2 should be obvious.