# Principal Component Analysis 

Laurenz Wiskott<br>Institute for Theoretical Biology<br>Humboldt-University Berlin<br>Invalidenstraße 43<br>D-10115 Berlin, Germany

11 March 2004

## 1 Intuition

Problem Statement Experimental data to be analyzed is often represented as a number of vectors of fixed dimensionality. A single vector could for example be a set of temperature measurements across Germany. Taking such a vector of measurements at different times results in a number of vectors that altogether constitute the data. Each vector can also be interpreted as a point in a high dimensional space. Then the data are simply a cloud of points in this space.

When analyzing such data one often encounters the problem that the dimensionality of the data points is too high to be visualized or analyzed with some particular technique. Thus the problem arises to reduce the dimensionality of the data in some optimal way.

To keep things simple we insist that the dimensionality reduction is done linearly, i.e. we are looking for a low-dimensional linear subspace of the data space, onto which the data can be projected. As a criterion for what the optimal subspace might be it seems reasonable to require that it should be possible to reconstruct the original data points from the reduced ones as well as possible. Thus if one were to project the data back from the low-dimensional space into the original high-dimensional space, the reconstructed data points should lie as close as possible to the original ones, with the mean squared distance between original and reconstructed data points being the reconstruction error. The question is, how can we find the linear subspace that minimizes this reconstruction error.

Projection and Reconstruction The task of principal component analysis (PCA) is to reduce the dimensionality of some high-dimensional data points by linearly projecting them onto a lower-dimensional space in such a way that the reconstruction error made by this projection is minimal. In order to develop an intuition for PCA we first take a closer look at what it means to project the data points and to reconstruct them. Figure 1 illustrates the process. (a) A few data points are given in a two-dimensional space and are represented by pairs of numbers $\left(x_{1}, x_{2}\right)$. (b) In order to reduce the dimensionality down to one, we have to choose a one-dimensional subspace and project the data points onto

[^0]it. (c) The points can now be represented by just one number, $y$, and we do not care that they originally came from a two-dimensional space. (d) If we want to reconstruct the original two-dimensional positions of the data points as well as possible, we have to embed the one-dimensional space in the original two-dimensional space in exactly the orientation used during the projection. However, we cannot recover the accurate 2D-position; the points remain on the one-dimensional subspace. The reconstruction error is therefore the average distance of the original 2D-positions from the one-dimensional subspace (the length of the projection arrows in (b)). For mathematical convenience one actually takes the average squared distance.

Reconstruction Error and Variance The question now is how we can find the direction of the one-dimensional subspace that minimizes the reconstruction error. For that it is interesting to inspect more closely what happens as we rotate the subspace. Figure 2 illustrates the projection onto two different subspaces. Focus just on the one point $\mathbf{x}$ and its projection $\mathbf{x}_{\|} . d$ is the distance of $\mathbf{x}$ from the origin, $r$ is the distance of $\mathbf{x}$ from $\mathbf{x}_{\|}$in the subspace, and $v$ is the distance of $\mathbf{x}_{\| \mid}$from the origin. $r$ and $v$ depend on the direction of the subspace while $d$ does not. Interestingly, since the triangles between $\mathbf{x}, \mathbf{x}_{\|}$, and the origin are right-angled, $r$ and $v$ are related by Pythagoras' theorem, i.e. $r^{2}+v^{2}=d^{2}$. We know that $r^{2}$ contributes to the reconstruction error. $v^{2}$ on the other hand contributes to the variance of the projected data within the subspace. Thus we see that the sum over the reconstruction error plus the variance of the projected data is constant and equals the variance of the original data. Therefore, minimizing the reconstruction error is equivalent to maximizing the variance of the projected data.


Figure 1: Projection of 2D data points onto a 1D subspace and their reconstruction.



Figure 2 Variance of the projected data and reconstruction error as the linear subspace is rotated.

Cavariance Matrix How can we determine the direction of maximal variance? The first we can do is to determine the variances of the individual components. If the data points (or vectors) are written as $\mathbf{x}=\left(x_{1}, x_{2}\right)^{T}$ ( $T$ indicates transpose), then the variances of the first and second component can be written as $C_{11}:=\left\langle x_{1} x_{1}\right\rangle$ and $C_{22}:=\left\langle x_{2} x_{2}\right\rangle$ (angle brackets indicate averaging over all data points). If $C_{11}$ is large compared to $C_{22}$, then the direction of maximal variance is close to $(1,0)^{T}$, while if $C_{11}$ is small, the direction of maximal variance is close to $(0,1)^{T}$. (Notice that variance doesn't have a polarity, so that one could use the inverse vector $(-1,0)^{T}$ instead of $(1,0)^{T}$ equally well for indicating the direction of maximal variance.)

But what if $C_{11}$ is of similar value as $C_{22}$, like in the example of Figure 1? Then the co-variance between the two components, $C_{12}:=\left\langle x_{1} x_{2}\right\rangle$, can give us additional information (notice that $C_{21}:=\left\langle x_{2} x_{1}\right\rangle$ is equal to $C_{12}$ ). A large positive value of $C_{12}$ indicates a strong correlation between $x_{1}$ and $x_{2}$ and that the data cloud is extended along the $(1,1)^{T}$ direction. A negative value would indicate anti-correlation and an extension along the $(-1,1)^{T}$ direction. A small value of $C_{12}$ would indicate no correlation and thus little structure of the data, i.e. no prominent direction of maximal variance. The variances and covariances are conveniently arranged in a matrix with components $C_{i j}$, which is called covariance matrix (assuming zero mean data). Figure 3 shows several data clouds and the corresponding covariance matrices.


Figure 3: Several data distributions and their covariance matrices.

Covariance Matrix and Higher Order Structure Notice that the covariance matrix only gives you information about the general extent of the data (the second order moments). It does not give you any information about the structure of the data cloud. Figure 4 shows different data distributions that all have the same covariance matrix. Thus as long as we consider only the covariance matrix, i.e. second oder moments, we can always assume a Gaussian data distribution with an ellipsoid shape, because the covariance matrix does not represent any more structure in any case.

PCA by Diagonalizing the Covariance Matrix Now that we have learned that the covariance matrix in principle contains the information about the direction of maximal variance the question arises how we can get at this information. From Figure 3 (a) and (b) we can see that there are two fundamentally different situations: in (a) the data cloud is aligned with the axes of the coordinate system and the covariance matrix is diagonal; in (b) the data cloud is oblique to the axes and the matrix is not diagonal. In the former case the direction of maximal variance is simply the axis belonging to the largest value on the diagonal of the covariance matrix. In the latter case, we cannot directly say what the direction
of maximal variance might be. Thus, since the case of a diagonal covariance matrix is so much simpler, the strategy we are going to take is to make a non-diagonal covariance matrix digonal by rotating the coordinate system accordingly. This is illustrated in Figure 5. From linear algebra we know that diagonalizing a matrix can be done by solving the corresponding eigenvalue equation. It will turn out that the eigenvectors of the covariance matrix point into the directions of maximal (and minimal) variance and that the eigenvalues are equal to the variances along these directions. Projecting the data onto the eigenvectors with largest eigenvalues is therefore the optimal linear dimensionality reduction.

## 2 Formalism

Definition of the PCA-Optimization Problem The problem of principal component analysis (PCA) can be formally states as follows.

Principal Component Analysis (PCA): Given a set $\left\{\mathbf{x}^{\mu}: \mu=1, \ldots, M\right\}$ of $I$-dimensional data points $\mathbf{x}^{\mu}=\left(x_{1}^{\mu}, x_{2}^{\mu}, \ldots, x_{I}^{\mu}\right)^{T}$ with zero mean, $\left\langle\mathbf{x}^{\mu}\right\rangle=\mathbf{0}$, find an orthogonal matrix $\mathbf{U}$ with determinant $|\mathbf{U}|=+1$ generating the transformed data points $\mathbf{x}^{\prime \mu}:=\mathbf{U}^{T} \mathbf{x}^{\mu}$ such that for any given dimensionality $P$ the data projected onto the first $P$ axes, $\mathbf{x}_{\|}^{\prime \mu}:=\left(x_{1}^{\prime \mu}, x_{2}^{\prime \mu}, \ldots, x_{P}^{\prime \mu}, 0, \ldots, 0\right)^{T}$, has the smallest reconstruction error $\left.E:=\langle | \mathbf{x}^{\prime \mu}-\left.\mathbf{x}_{\|}^{\prime \mu}\right|^{2}\right\rangle_{\mu}$ among all possible projections onto a $P$-dimensional subspace. The row vectors of matrix $\mathbf{U}$ define the new axes and are called the principal components.

Some remarks: (i) If one has non-zero-mean data, one typically removes the mean before applying PCA. Even though all the math is valid also for non-zero-mean data, the results


Figure 4: Different data distributions with identical covariance matrices.


Figure 5: Diagonalizing the covariance matrix by rotating the coordinate system.
would typically be undesired and nonintuitive. (ii) Since matrix $\mathbf{U}$ is orthogonal and has determinant value +1 , it corresponds simply to a rotation of the data $\mathbf{x}$. Thus, the 'shape' of the data cloud remains the same, just the 'perspective' changes. Note also that one can interpret the multiplication with matrix $\mathbf{U}^{T}$ either as a rotation of the data or as a rotation of the coordinate system. Either interpretation is valid. (iii) Projecting the data $\mathrm{x}^{\prime}$ onto the $P$-dimensional linear subspace spanned by the first $P$ axes is simply done by setting all components higher than $P$ to zero. This can be done, because we still have an orthonormal coordinate system. If $\mathbf{U}$ and therefore the new coordinate system were not orthogonal then the projection became a mathematically more complex operation. (iv) The reconstruction error has to be minimal for any $P$. This has the advantage that we do not need to decide on $P$ before performing PCA. Often $P$ is actually choosen based on information obtained during PCA and governed by a constraint, such as that the reconstruction error should be below a certain threshold. (v) From now on we will drop the index $\mu$ and keep in mind that averages indicated by $\langle\cdot\rangle$ are always taken over all data points indexed with $\mu$.

Mapping from High- to Low-dimensional Space and New Coordinate System Assume some data points $\mathbf{x}$ are given in an $I$-dimensional space and a linear subspace is spanned by $P$ orthonormal vectors

$$
\begin{align*}
\mathbf{v}_{p} & :=\left(v_{1 p}, v_{2 p}, \ldots, v_{I p}\right)^{T}  \tag{1}\\
\text { with } \quad \mathbf{v}_{p}^{T} \mathbf{v}_{q} & =\delta_{p q}:=\left\{\begin{array}{ll}
1 & \text { if } p=q \\
0 & \text { otherwise }
\end{array} .\right. \tag{2}
\end{align*}
$$

We will typically assume $P<I$ and speak of a high $(I)$-dimensional space and a low $(P)$ dimensional (sub)space. However, $P=I$ may be possible as a limiting case as well.

Arranging these vectors in a matrix yields

$$
\begin{align*}
\mathbf{V} & :=\left(\mathbf{v}_{1}, \mathbf{v}_{2}, \ldots, \mathbf{v}_{P}\right)  \tag{3}\\
& \stackrel{(1)}{=}\left(\begin{array}{cccc}
v_{11} & v_{12} & \ldots & v_{1 P} \\
v_{21} & v_{22} & \ldots & v_{2 P} \\
\vdots & & \ddots & \vdots \\
v_{I 1} & v_{I 2} & \ldots & v_{I P}
\end{array}\right) . \tag{4}
\end{align*}
$$

This matrix can be used to map the data points $\mathbf{x}$ into the subspace spanned by the vectors $\mathbf{v}_{p}$ yielding

$$
\begin{equation*}
\mathbf{y}:=\mathbf{V}^{T} \mathbf{x} \tag{5}
\end{equation*}
$$

If $P<I$ then the dimensionality is reduced and some information is lost; if $P=I$ all information is preserved. In any case the mapped data are now represented in a new coordinate system the axes of which are given by the vectors $\mathbf{v}_{p}$. With $P=2$ and $I=3$, for example, we have

$$
\begin{gathered}
\mathbf{y}=\binom{\mathbf{v}_{1}^{T} \mathbf{x}}{\mathbf{v}_{2}^{T} \mathbf{x}}=\binom{\mathbf{v}_{1}^{T}}{\mathbf{v}_{2}^{T}} \mathbf{x}=\mathbf{V}^{T} \mathbf{x} \\
\text { or } \quad \mathbf{y}=\binom{y_{1}}{y_{2}}=\left(\begin{array}{lll}
v_{11} & v_{21} & v_{31} \\
v_{12} & v_{22} & v_{32}
\end{array}\right)\left(\begin{array}{c}
x_{1} \\
x_{2} \\
x_{3}
\end{array}\right)=\mathbf{V}^{T} \mathbf{x}
\end{gathered}
$$

Note that $\mathbf{y}$ is $P$-dimensional while $\mathbf{x}$ is $I$-dimensional.
It is important to realize that we have done two things here: firstly, we have moved the points from the high-dimensional space onto the low-dimensional subspace (the points that were already in the subspace have not been moved, of course) and secondly, we have represented the moved points in a new coordinate system that is particularly suitable for the low-dimensional subspace. Thus, we went from the high-dimensional space and the old coordinate system to the low-dimensional subspace and a new coordinate system. Note also that points in the high-dimensional space can generally not be represented in the new coordinate system.

Transforming from New to Old Coordinate System Interestingly, since the vectors $\mathbf{v}_{p}$ are orthonormal, matrix $\mathbf{V}$ can also be used to transform the points back from the new to the old coordinate system, although, the lost dimensions cannot be recovered, of course. Thus the mapped points $\mathbf{y}$ in the new coordinate system become points $\mathbf{x}_{\|}$in the old coordinate system and are given by

$$
\begin{align*}
\mathbf{x}_{\|} & :=\mathbf{V} \mathbf{y}  \tag{6}\\
& \stackrel{(5)}{=} \mathbf{V V}^{T} \mathbf{x} . \tag{7}
\end{align*}
$$

$\mathbf{y}$ and $\mathbf{x}_{\|}$are equivalent representations, i.e. they contain the same information, just in different coordinate systems.

Combined matrix $\mathbf{V}^{T} \mathbf{V}$ Before we look at the combined matrix $\mathbf{V} \mathbf{V}^{T}$ consider $\mathbf{V}^{T} \mathbf{V}$. The latter is obviously a $P \times P$-matrix and performs a transformation from the new coordinate system to the old coordinate system (6) and back again (5) plus a mapping from the high-dimensional space onto the low-dimensional space. However, since all points in the new coordinate system lie within the low-dimensional subspace already, the mapping onto the low-dimensional space is without any effect. Thus, only the back and forth transformation between the two coordinate systems remains and that in combination is without any effect either. This means that $\mathbf{V}^{T} \mathbf{V}$ is the identity matrix, which can be verified easily

$$
\begin{align*}
\left(\mathbf{V}^{T} \mathbf{V}\right)_{p q} & =\mathbf{v}_{p}^{T} \mathbf{v}_{q} \stackrel{(2)}{=} \delta_{p q}  \tag{8}\\
\Longleftrightarrow \quad \mathbf{V}^{T} \mathbf{V} & =\mathbf{1}_{P} \tag{9}
\end{align*}
$$

with $\mathbf{1}_{P}$ indicating the identity matrix of dimensionality $P$. With $P=2$, for example, we have

$$
\mathbf{V}^{T} \mathbf{V}=\binom{\mathbf{v}_{1}^{T}}{\mathbf{v}_{2}^{T}}\left(\mathbf{v}_{1} \mathbf{v}_{2}\right)=\left(\begin{array}{ll}
\mathbf{v}_{1}^{T} \mathbf{v}_{1} & \mathbf{v}_{1}^{T} \mathbf{v}_{2} \\
\mathbf{v}_{2}^{T} \mathbf{v}_{1} & \mathbf{v}_{2}^{T} \mathbf{v}_{2}
\end{array}\right) \stackrel{(2)}{=}\left(\begin{array}{ll}
1 & 0 \\
0 & 1
\end{array}\right) .
$$

Projection from High- to Low-dimensional Space As we have seen above (7) the combined matrix $\mathbf{V V}^{T}$ maps the points $\mathbf{x}$ onto the low-dimensional subspace but in contrast to matrix $\mathbf{V}^{T}$ alone the mapped points are represented within the old coordinate system and not the new one. It turns out that this is a projection operation with the important property that it does not make a difference whether you apply it once or twice, i.e. $\mathbf{P P}=\mathbf{P}$. Let us therefore define the projection matrix

$$
\begin{equation*}
\mathbf{P}:=\mathbf{V} \mathbf{V}^{T} \tag{10}
\end{equation*}
$$

and verify that

$$
\begin{align*}
\mathbf{P P} & \stackrel{(10)}{=} \mathbf{V} \underbrace{\mathbf{V}^{T} \mathbf{V}}_{\mathbf{1}_{P}} \mathbf{V}^{T}  \tag{11}\\
& \stackrel{(9)}{=} \mathbf{V} \mathbf{V}^{T}  \tag{12}\\
& \stackrel{(10)}{=} \mathbf{P} . \tag{13}
\end{align*}
$$

A closer look at $\mathbf{P}$ shows that

$$
\begin{align*}
\mathbf{P} & :=\mathbf{V V}^{T}  \tag{14}\\
& =\left(\mathbf{v}_{1}, \ldots, \mathbf{v}_{P}\right)\left(\begin{array}{c}
\mathbf{v}_{1}^{T} \\
\vdots \\
\mathbf{v}_{P}^{T}
\end{array}\right)  \tag{15}\\
& =\sum_{p=1}^{P} \mathbf{v}_{p} \mathbf{v}_{p}^{T} \tag{16}
\end{align*}
$$

$\mathbf{P}$ is obviously an $I \times I$-matrix. If $P=I$ then projecting from the old to the new and back to the old coordinate system causes no information loss and $\mathbf{P}=\mathbf{1}_{I}$. The smaller $P$ the more information is lost and the more does $\mathbf{P}$ differ from the identity matrix. Consider, for example

$$
\begin{aligned}
& \mathbf{v}_{1}:=\frac{1}{2}(\sqrt{2},-1,1)^{T} \Rightarrow \mathbf{v}_{1} \mathbf{v}_{1}^{T}=\frac{1}{4}\left(\begin{array}{rrr}
2 & -\sqrt{2} & \sqrt{2} \\
-\sqrt{2} & 1 & -1 \\
\sqrt{2} & -1 & 1
\end{array}\right), \\
& \mathbf{v}_{2}:=\frac{1}{2}(0, \sqrt{2}, \sqrt{2})^{T} \Rightarrow \mathbf{v}_{2} \mathbf{v}_{2}^{T}=\frac{1}{4}\left(\begin{array}{rrr}
0 & 0 & 0 \\
0 & 2 & 2 \\
0 & 2 & 2
\end{array}\right), \text { and } \\
& \mathbf{v}_{3}:=\frac{1}{2}(-\sqrt{2},-1,1)^{T} \Rightarrow \mathbf{v}_{3} \mathbf{v}_{3}^{T}=\frac{1}{4}\left(\begin{array}{rrr}
2 & \sqrt{2} & -\sqrt{2} \\
\sqrt{2} & 1 & -1 \\
-\sqrt{2} & -1 & 1
\end{array}\right)
\end{aligned}
$$

for which you can easily verify that $\mathbf{P}$ (16) successively becomes the identity matrix as you take more of the $\mathbf{v}_{p} \mathbf{v}_{p}^{T}$-terms.

Variance The variance of a multi-dimensional data set is defined as the sum over the variances of its components. Since we assume zero-mean data, we have

$$
\begin{align*}
\operatorname{var}(\mathbf{x}) & :=\sum_{i=1}^{I}\left\langle x_{i}^{2}\right\rangle  \tag{17}\\
& =\left\langle\sum_{i=1}^{I} x_{i}^{2}\right\rangle  \tag{18}\\
& =\left\langle\mathbf{x}^{T} \mathbf{x}\right\rangle \tag{19}
\end{align*}
$$

This also holds for the projected data, of course, $\operatorname{var}(\mathbf{y})=\left\langle\mathbf{y}^{T} \mathbf{y}\right\rangle$.

Reconstruction Error The reconstruction error $E$ is defined as the mean square sum over the distances between the original data points $\mathbf{x}$ and the projected ones $\mathbf{x}_{\|}$. If we define the projection vectors ('Lotvektoren' in German) $\mathbf{x}_{\perp}=\mathbf{x}-\mathbf{x}_{\|}$(in contrast to the projected vectors $\mathbf{x}_{\|}$) we can write the reconstruction error as the variance of the projection vectors and find

$$
\begin{align*}
E & =\left\langle\mathbf{x}_{\perp}^{T} \mathbf{x}_{\perp}\right\rangle  \tag{20}\\
& =\left\langle\left(\mathbf{x}-\mathbf{x}_{\|}\right)^{T}\left(\mathbf{x}-\mathbf{x}_{\|}\right)\right\rangle  \tag{21}\\
& =\left\langle\left(\mathbf{x}-\mathbf{V} \mathbf{V}^{T} \mathbf{x}\right)^{T}\left(\mathbf{x}-\mathbf{V} \mathbf{V}^{T} \mathbf{x}\right)\right\rangle  \tag{22}\\
& =\langle\mathbf{x}^{T} \mathbf{x}-2 \mathbf{x}^{T} \mathbf{V} \mathbf{V}^{T} \mathbf{x}+\mathbf{x}^{T} \mathbf{V} \underbrace{\left(\mathbf{V}^{T} \mathbf{V}\right)}_{=\mathbf{1}_{P}} \mathbf{V}^{T} \mathbf{x}\rangle  \tag{23}\\
& =\left\langle\mathbf{x}^{T} \mathbf{x}\right\rangle-\langle\mathbf{x}^{T} \mathbf{V} \underbrace{\left(\mathbf{V}^{T} \mathbf{V}\right)}_{=\mathbf{1}_{P}} \mathbf{V}^{T} \mathbf{x}\rangle  \tag{24}\\
& =\left\langle\mathbf{x}^{T} \mathbf{x}\right\rangle-\left\langle\mathbf{x}_{\|}^{T} \mathbf{x}_{\|}\right\rangle  \tag{25}\\
& =\left\langle\mathbf{x}^{T} \mathbf{x}\right\rangle-\left\langle\mathbf{y}^{T} \mathbf{y}\right\rangle . \tag{26}
\end{align*}
$$

This means that the reconstruction error equals the difference between the variance of the data minus the variance of the projected data. Thus, this verifies our intuition that minimizing the reconstruction error is equivalent to maximizing the variance of the projected data.

Covariance Matrix We have already argued heuristically that the covariance matrix $\mathbf{C}_{x}$ with $C_{x i j}:=\left\langle x_{i} x_{j}\right\rangle$ plays an important role in performing PCA. It is convenient to write the covariance matrix in vector notation:

$$
\begin{equation*}
\mathbf{C}_{x}:=\left\langle\mathbf{x} \mathbf{x}^{T}\right\rangle=\frac{1}{M} \sum_{\mu} \mathbf{x}^{\mu} \mathbf{x}^{\mu T} \tag{27}
\end{equation*}
$$

It is an easy exercise to show that this definition is equivalent to the componentwise one given above. Since $\left(\mathbf{x x}^{T}\right)^{T}=\mathbf{x x}^{T}$ (remember $(\mathbf{A B})^{T}=\mathbf{B}^{T} \mathbf{A}^{T}$ ), one can also see that $\mathbf{C}_{x}$ is symmetric, i.e. $\mathbf{C}_{x}^{T}=\mathbf{C}_{x}$.

Eigenvalue Equation of the Covariance Matrix Since the covariance matrix is symmetric, its eigenvalues are real and a set of orthogonal eigenvectors always exists. In mathematical terms, for a given covariance matrix $\mathbf{C}_{x}$ we can always find a complete set of eigenvalues $\lambda_{i}$ and corresponding eigenvectors $\mathbf{u}_{i}$ such that

$$
\begin{align*}
\mathbf{C}_{x} \mathbf{u}_{i} & =\mathbf{u}_{i} \lambda_{i} \quad \text { (eigenvalue equation) }  \tag{28}\\
\lambda_{i} & \geq \lambda_{i+1} \quad \text { (eigenvalues are ordered) }  \tag{29}\\
\mathbf{u}_{i}^{T} \mathbf{u}_{j} & =\delta_{i j} \quad \text { (eigenvectors are orthonormal) } . \tag{30}
\end{align*}
$$

If we combine the eigenvectors into an orthogonal matrix $\mathbf{U}$ and the eigenvalues into a diagonal matrix $\boldsymbol{\Lambda}$,

$$
\begin{align*}
\mathbf{U} & :=\left(\mathbf{u}_{1}, \mathbf{u}_{2}, \ldots, \mathbf{u}_{I}\right)  \tag{31}\\
\mathbf{\Lambda} & :=\operatorname{diag}\left(\lambda_{1}, \lambda_{2}, \ldots, \lambda_{I}\right) \tag{32}
\end{align*}
$$

then we can rewrite (30) and (28) as

$$
\begin{align*}
\mathbf{U}^{T} \mathbf{U} & =\mathbf{1}_{I} \quad \text { (matrix } \mathbf{U} \text { is orthogonal) }  \tag{33}\\
\Longleftrightarrow \mathbf{U U}^{T} & =\mathbf{1}_{I} \quad \text { (since } \mathbf{U}^{-1}=\mathbf{U}^{T} \text { and } \mathbf{U} \text { is quadratic) },  \tag{34}\\
\mathbf{C}_{x} \mathbf{U} & =\mathbf{U} \mathbf{\Lambda} \quad(\text { eigenvalue equation })  \tag{35}\\
\stackrel{(33)}{\Longleftrightarrow} \quad \mathbf{U}^{T} \mathbf{C}_{x} \mathbf{U} & =\mathbf{U}^{T} \mathbf{U} \mathbf{\Lambda} \stackrel{(33)}{=} \mathbf{\Lambda} \tag{36}
\end{align*}
$$

Total Variance of the Data $x$ Given the eigenvector matrix $\mathbf{U}$ and the eigenvalue matrix $\boldsymbol{\Lambda}$ it is easy to compute the total variance of the data

$$
\begin{align*}
\left\langle\mathbf{x}^{T} \mathbf{x}\right\rangle & =\left\langle\operatorname{tr}\left(\mathbf{x}^{T} \mathbf{x}\right)\right\rangle \quad(\text { since } \lambda=\operatorname{tr}(\lambda) \text { for any real } \lambda)  \tag{37}\\
& =\left\langle\operatorname{tr}\left(\mathbf{x} \mathbf{x}^{T}\right)\right\rangle \quad(\text { since } \operatorname{tr}(\mathbf{A B})=\operatorname{tr}(\mathbf{A B}) \text { for any matrices } \mathbf{A}, \mathbf{B})  \tag{38}\\
& =\operatorname{tr}\left(\left\langle\mathbf{\mathbf { x } ^ { T } \rangle )} \quad(\text { since } \operatorname{tr}(\cdot) \text { and }\langle\cdot\rangle \text { commute })\right.\right.  \tag{39}\\
& \stackrel{(27)}{=} \operatorname{tr}\left(\mathbf{C}_{x}\right)  \tag{40}\\
& \stackrel{(34)}{=} \operatorname{tr}\left(\mathbf{U U}^{T} \mathbf{C}_{x}\right)  \tag{41}\\
& =\operatorname{tr}\left(\mathbf{U}^{T} \mathbf{C}_{x} \mathbf{U}\right)  \tag{42}\\
& \stackrel{(36)}{=} \operatorname{tr}(\mathbf{\Lambda})  \tag{43}\\
& \stackrel{(32)}{=} \sum_{i} \lambda_{i} . \tag{44}
\end{align*}
$$

Thus the total variance of the data is simply the sum of the eigenvalues of its covariance matrix.

Notice that on the way of this proof we have shown some very general properties. From line (37) to (40) we have shown that the variance of some multi-dimensional data equals the trace of its covariance matrix. From line (40) to (42) we have shown that the trace remains invariant under any orthogonal transformation of the coordinate system.

Diagonalizing the Covariance Matrix We can now use matrix $\mathbf{U}$ to transform the data such that the covariance matrix becomes diagonal. Define $\mathbf{x}^{\prime}:=\mathbf{U}^{T} \mathbf{x}$ and denote the new covariance matrix by $\mathbf{C}_{x}^{\prime}$. We have

$$
\begin{align*}
\mathbf{x}^{\prime} & =\mathbf{U}^{T} \mathbf{x}  \tag{45}\\
\Longrightarrow \quad \mathbf{C}_{x}^{\prime} & =\left\langle\mathbf{x}^{\prime} \mathbf{x}^{\prime T}\right\rangle  \tag{46}\\
& =\left\langle\left(\mathbf{U}^{T} \mathbf{x}\right)\left(\mathbf{U}^{T} \mathbf{x}\right)^{T}\right\rangle  \tag{47}\\
& =\mathbf{U}^{T}\left\langle\mathbf{x} \mathbf{x}^{T}\right\rangle \mathbf{U}  \tag{48}\\
& =\mathbf{U}^{T} \mathbf{C}_{x} \mathbf{U},  \tag{49}\\
& =\mathbf{\Lambda} \tag{50}
\end{align*}
$$

and find that the transformed data $\mathbf{x}^{\prime}$ have a diagonal covariance matrix. Working with $\mathbf{x}^{\prime}$ instead of $\mathbf{x}$ will simplify further analysis without loss of generality.

Variance of y for a Diagonalized Covariance Matrix Now that we have the data represented in a coordinate system in which the covariance matrix is diagonal, we can try
to answer the question, which is the $P$-dimensional subspace that minimizes the reconstruction error. Our intuition would predict that it is simply the space spanned by the first $P$ eigenvectors. To show this analytically, we take an arbitrary set of $P$ orthonormal vectors $\mathbf{v}_{p}^{\prime}$, and compute the variance of $\mathbf{y}$.

$$
\begin{align*}
\mathbf{y} & :=\mathbf{V}^{\prime T} \mathbf{x}^{\prime}  \tag{51}\\
\Longrightarrow\left\langle\mathbf{y}^{T} \mathbf{y}\right\rangle & \stackrel{(51)}{=}\left\langle\mathbf{x}^{\prime T} \mathbf{V}^{\prime} \mathbf{V}^{\prime T} \mathbf{x}^{\prime}\right\rangle  \tag{52}\\
& =\left\langle\operatorname{tr}\left(\mathbf{x}^{\prime T} \mathbf{V}^{\prime} \mathbf{V}^{\prime T} \mathbf{x}^{\prime}\right)\right\rangle \quad(\text { since } \lambda=\operatorname{tr}(\lambda) \text { for any real } \lambda)  \tag{53}\\
& =\left\langle\operatorname{tr}\left(\mathbf{V}^{\prime T} \mathbf{x}^{\prime} \mathbf{x}^{\prime T} \mathbf{V}^{\prime}\right)\right\rangle \quad(\text { since } \operatorname{tr}(\mathbf{A B C})=\operatorname{tr}(\mathbf{B C A}) \text { if defined })  \tag{54}\\
& \stackrel{(46)}{=} \operatorname{tr}\left(\mathbf{V}^{\prime T} \mathbf{C}_{x}^{\prime} \mathbf{V}^{\prime}\right) \quad(\text { since } \operatorname{tr}(\cdot) \text { and }\langle\cdot\rangle \text { commute })  \tag{55}\\
& \stackrel{(50)}{=} \operatorname{tr}\left(\mathbf{V}^{\prime T} \mathbf{\Lambda} \mathbf{V}^{\prime}\right)  \tag{56}\\
& =\sum_{i} \lambda_{i} \sum_{p}\left(v_{i p}^{\prime}\right)^{2} . \quad \text { (as one can work out on a sheet of paper) } \tag{57}
\end{align*}
$$

Constraints of Matrix $\mathbf{V}^{\prime}$ Note that, since the vectors $\mathbf{v}^{\prime}{ }_{p}$ are orthonormal, $\mathbf{V}^{\prime}$ can always be completed to an orthogonal $I \times I$-matrix by adding $I-P$ additional orthonormal vectors. Since we know that an orthogonal matrix has normalized row as well as column vectors, we see that, by taking away the $I-P$ additional column vectors, we are left with the constraints

$$
\begin{align*}
\sum_{p}\left(v_{i p}^{\prime}\right)^{2} & \leq 1 \quad \text { (row vectors of } \mathbf{V}^{\prime} \text { have norm less or equal one) }  \tag{58}\\
\sum_{i}\left(v_{i p}^{\prime}\right)^{2} & =1 \quad\left(\text { column vectors of } \mathbf{V}^{\prime}\right. \text { have norm one) }  \tag{59}\\
\Longrightarrow \quad \sum_{i p}\left(v_{i p}^{\prime}\right)^{2} & =P \quad \text { (square sum over all matrix elements equals } P) . \tag{60}
\end{align*}
$$

Notice that Constraint (60) is a direct consequence of Constraint (59) and does not need to be verified separately in the following considerations.

Finding the Optimal Subspace Since the variance (57) of $\mathbf{y}$ as well as the constraints $(58,59,60)$ of Matrix $\mathbf{V}^{\prime}$ are linear in $\left(v_{i p}^{\prime}\right)^{2}$, maximization of the variance $\left\langle\mathbf{y}^{T} \mathbf{y}\right\rangle$ is obviously achieved by putting as much 'weight' as possible on the large eigenvalues, which are the first ones. The simplest way of doing that is to set

$$
v_{i p}^{\prime}:=\delta_{i p}:= \begin{cases}1 & \text { if } i=p  \tag{61}\\ 0 & \text { otherwise }\end{cases}
$$

with the Kronecker symbol $\delta_{i p}$.
Since $I \geq P$ we can verify the constraints

$$
\begin{align*}
& \sum_{p}\left(v_{i p}^{\prime}\right)^{2} \stackrel{(61)}{=} \sum_{p} \delta_{i p}^{2}=\left\{\begin{array}{ll}
1 & \text { if } i \leq P \\
0 & \text { otherwise }
\end{array}\right\} \leq 1,  \tag{62}\\
& \sum_{i}\left(v_{i p}^{\prime}\right)^{2} \stackrel{(61)}{=} \sum_{i} \delta_{i p}^{2}=\delta_{p p}^{2}=1 \tag{63}
\end{align*}
$$

and see from (62) that there is actually as much 'weight' on the first, i.e. large, eigenvalues as Constraint (58) permits.

Interpretation of the Result What does it mean to set $v_{i p}^{\prime}:=\delta_{i p}$ ? It means that $\mathbf{V}^{\prime}$ projects the data $\mathbf{x}^{\prime}$ onto the first $P$ axis, which in fact is a projection onto the first $P$ eigenvectors of the covariance matrix $\mathbf{C}_{x}$. Thus, if we define

$$
\begin{align*}
\mathbf{V} & :=  \tag{64}\\
\stackrel{(31,61)}{=} & \left(\mathbf{U V}^{\prime}, \mathbf{u}_{2}, \ldots, \mathbf{u}_{P}\right) \tag{65}
\end{align*}
$$

we can go back to the original coordinate system and find

$$
\begin{align*}
\mathbf{y} & \stackrel{(51)}{=} \mathbf{V}^{\prime T} \mathbf{x}^{\prime}  \tag{66}\\
& \stackrel{(45)}{=} \mathbf{V}^{\prime T} \mathbf{U}^{T} \mathbf{x}  \tag{67}\\
& \stackrel{(64)}{=} \mathbf{V}^{T} \mathbf{x}, \tag{68}
\end{align*}
$$

which we know has maximal variance. Thus, if we start from the original data $\mathbf{x}$ we would set $\mathbf{v}_{p}:=\mathbf{u}_{p}$.

For the variance of $\mathbf{y}$ we find

$$
\begin{align*}
\left\langle\mathbf{y}^{T} \mathbf{y}\right\rangle & \stackrel{(57)}{=} \sum_{i=1}^{I} \lambda_{i} \sum_{p=1}^{P}\left(v_{i p}^{\prime}\right)^{2}  \tag{69}\\
& \stackrel{(61)}{=} \sum_{i=1}^{I} \lambda_{i} \sum_{p=1}^{P} \delta_{i p}^{2}  \tag{70}\\
& =\sum_{i=1}^{P} \lambda_{i} \tag{71}
\end{align*}
$$

which is the sum over the first $P$ largest eigenvalues of the covariance matrix. Likewise one can determine the reconstruction error as

$$
\begin{align*}
& E \stackrel{(26)}{=}\left\langle\mathbf{x}^{T} \mathbf{x}\right\rangle-\left\langle\mathbf{y}^{T} \mathbf{y}\right\rangle  \tag{72}\\
& \stackrel{(44,71)}{=} \sum_{i=1}^{I} \lambda_{i}-\sum_{i=1}^{P} \lambda_{i}  \tag{73}\\
&= \sum_{i=P+1}^{I} \lambda_{i} . \tag{74}
\end{align*}
$$

Notice that this is just one optimal set of weights. We have seen above that the projected data can be rotated arbitrarily without changing its variance and therefore without changing its reconstruction error. This is equivalent to a rotation of the projection vectors $\mathbf{v}_{p}$ within the space spanned by the first eigenvectors.

Whitening or Sphering Sometimes it is desirable to transform a data set such that it has variance one in all directions. Such a normalization operation is called whitening or sphering. The latter term is quite intuitive, because a spherical data distribution has the same variance in all directions. Intuitively speaking sphering requires to stretch and compress the data distribution along the axes of the principal components such they have variance one. Technically speaking one first transforms the data into a coordinate system
where the covariance matrix is diagonal, then one performs the stretching along the axes, and then one transforms the data back into the original coordinate system. Principal component analysis obviously gives all the required information. The eigenvectors of the covariance matrix provide the axes of the new coordinate system and the eigenvalues $\lambda_{i}$ indicate the variances and therefore how much one has to stretch the data. If the original variance is $\lambda_{i}$ then one obviously has to stretch by a factor of $1 / \sqrt{\lambda_{i}}$ to get variance one. Thus, sphering is achieved by

$$
\begin{equation*}
\overline{\mathbf{x}}:=\mathbf{U} \operatorname{diag}\left(\frac{1}{\sqrt{\lambda_{1}}}, \frac{1}{\sqrt{\lambda_{2}}}, \ldots, \frac{1}{\sqrt{\lambda_{I}}}\right) \mathbf{U}^{T} \mathbf{x} \tag{75}
\end{equation*}
$$

It is easy to verify that the sphered data $\overline{\mathbf{x}}$ has variance one in all directions and a unit covariance matrix.


[^0]:    (c) 2004 Laurenz Wiskott

