# Practical Approaches to Principal Component Analysis in the Presence of Missing Values 

Alexander Ilin<br>Tapani Raiko<br>ALEXANDER.ILIN @ TKK.FI<br>TAPANI.RAIKO@TKK.FI<br>Adaptive Informatics Research Center<br>Aalto University School of Science and Technology<br>P.O. Box 15400, FI-00076 Aalto, Finland

Editor: Tommi Jaakkola


#### Abstract

Principal component analysis (PCA) is a classical data analysis technique that finds linear transformations of data that retain the maximal amount of variance. We study a case where some of the data values are missing, and show that this problem has many features which are usually associated with nonlinear models, such as overfitting and bad locally optimal solutions. A probabilistic formulation of PCA provides a good foundation for handling missing values, and we provide formulas for doing that. In case of high dimensional and very sparse data, overfitting becomes a severe problem and traditional algorithms for PCA are very slow. We introduce a novel fast algorithm and extend it to variational Bayesian learning. Different versions of PCA are compared in artificial experiments, demonstrating the effects of regularization and modeling of posterior variance. The scalability of the proposed algorithm is demonstrated by applying it to the Netflix problem.


Keywords: principal component analysis, missing values, overfitting, regularization, variational Bayes

## 1. Introduction

Principal component analysis (PCA) is a data analysis technique that can be traced back to Pearson (1901). It can be used to compress data sets of high dimensional vectors into lower dimensional ones. This is useful, for instance, in visualization and feature extraction. PCA has been extensively covered in the literature (e.g., Jolliffe, 2002; Bishop, 2006; Diamantaras and Kung, 1996; Haykin, 1989; Cichocki and Amari, 2002). PCA can be derived from a number of starting points and optimization criteria. The most important of these are minimization of the mean-square error in data compression, finding mutually orthogonal directions in the data having maximal variances, and decorrelation of the data using orthogonal transformations.

In the data compression formulation, PCA finds a smaller-dimensional linear representation of data vectors such that the original data could be reconstructed from the compressed representation with the minimum square error. Assume that we have $n d \times 1$ data vectors $\mathbf{y}_{1}, \mathbf{y}_{2}, \ldots, \mathbf{y}_{n}$ that are modeled as

$$
\begin{equation*}
\mathbf{y}_{j} \approx \mathbf{W} \mathbf{x}_{j}+\mathbf{m}, \tag{1}
\end{equation*}
$$

where $\mathbf{W}$ is a $d \times c$ matrix, $\mathbf{x}_{j}$ are $c \times 1$ vectors of principal components, and $\mathbf{m}$ is a $d \times 1$ bias vector. We assume that $c \leq d \leq n$. Principal subspace methods (e.g., Cichocki and Amari, 2002;

Diamantaras and Kung, 1996) find $\mathbf{W}, \mathbf{x}_{j}$ and $\mathbf{m}$ such that the reconstruction error

$$
\begin{equation*}
C=\sum_{j=1}^{n}\left\|\mathbf{y}_{j}-\mathbf{W} \mathbf{x}_{j}-\mathbf{m}\right\|^{2} \tag{2}
\end{equation*}
$$

is minimized. In matrix notation, the data vectors and principal components can be compiled into $d \times n$ and $c \times n$ matrices $\mathbf{Y}=\left[\begin{array}{llll}\mathbf{y}_{1} & \mathbf{y}_{2} & \ldots & \mathbf{y}_{n}\end{array}\right]$ and $\mathbf{X}=\left[\begin{array}{llll}\mathbf{x}_{1} & \mathbf{x}_{2} & \ldots & \mathbf{x}_{n}\end{array}\right]$, and $y_{i j}, w_{i k}$ and $x_{k j}$ denote the elements of the matrices $\mathbf{Y}, \mathbf{W}$ and $\mathbf{X}$, respectively. The bias matrix $\mathbf{M}$ contains $n$ copies of the bias vector $\mathbf{m}$ as its columns. Principal subspace methods find $\mathbf{W}$ and $\mathbf{X}$ such that $\mathbf{Y} \approx \mathbf{W X}+\mathbf{M}$ and the minimized cost function is the sum of the squared elements (or Frobenius norm) of matrix $\mathbf{Y}-\mathbf{W X}-\mathbf{M}$ :

$$
\begin{equation*}
C=\|\mathbf{Y}-\mathbf{W} \mathbf{X}-\mathbf{M}\|_{\mathrm{F}}^{2} . \tag{3}
\end{equation*}
$$

Without any further constraints, there exist infinitely many ways to perform a decomposition that minimizes (2) or equivalently (3). This can be seen by noting that any rotation or scaling of $\mathbf{W}$ can be compensated by rotating or scaling $\mathbf{X}$ accordingly, leaving the product $\mathbf{W X}$ the same. However, the subspace spanned by the column vectors of the matrix $\mathbf{W}$, called the principal subspace, is unique.

PCA finds a specific representation of the principal subspace. It is traditionally defined using the requirement that the column vectors of $\mathbf{W}$ are mutually orthogonal, have unit length and, furthermore, for each $k=1, \ldots, c$, the first $k$ vectors form the $k$-dimensional principal subspace. This makes the solution practically unique (except for changing the sign, and in the special case of having equal eigenvalues, e.g., Diamantaras and Kung, 1996; Jolliffe, 2002; Haykin, 1989). In this article, we use the term PCA for methods which seek representations (1) by minimizing the error (2). Thus we assume that once the principal subspace is found, it can be transformed into the PCA solution. This is indeed true for the case of fully observed vectors $\mathbf{y}_{j}$ but can be more difficult in the case with missing values, as we discuss in Section 4.

Let us now consider the same problem when the data matrix $\mathbf{Y}$ has missing entries. In this paper, we make the typical assumption that values are missing at random (MAR) (Little and Rubin, 1987), that is, given the observed data, the missingness does not depend on the unobserved data or latent variables. An example where the assumption does not hold is when out-of-scale measurements are marked missing. In the following example, the data matrix contains $N=9$ observed values and 6 missing values (marked with a sign $\times$ ):

$$
\mathbf{Y}=\left[\begin{array}{ccccc}
y_{11} & y_{12} & y_{13} & y_{14} & \times \\
y_{21} & y_{22} & \times & \times & y_{25} \\
\times & \times & y_{33} & y_{34} & \times
\end{array}\right]
$$

A natural extension of PCA for the case with missing values would be to find a representation such that $\mathbf{Y} \approx \mathbf{W} \mathbf{X}+\mathbf{M}$ for the observed values. The rest of the matrix $\mathbf{W} \mathbf{X}+\mathbf{M}$ can be taken as the reconstruction of missing values.

Although the PCA problem in the presence of missing values seems to be as easy as classical PCA, there are some important distinctions: 1) There is no analytical solution available since even the estimation of the data covariance matrix is nontrivial. Therefore, iterative learning procedures must be exploited. 2) The optimized cost function typically has multiple local minima and thus finding the optimal solution is more difficult. 3) There is no analytical solution even for the bias


Figure 1: Two examples of PCA with missing values. In the left, the missing value mechanism is missing at random (MAR), while on the right, it is missing not at random (MNAR). The black dots represent fully observed samples, while the blue crosses represent observations where only the horizontal value has been observed. The red circles represent the reconstructions of the missing values produced by VBPCA. The dashed red ellipses represent the estimated data covariance. On the left, the estimated covariance is quite close to the data covariance (black ellipse).
term $\mathbf{m}$ in (1), which is not generally equal to the row-wise mean of the data matrix, as in classical PCA. 4) Standard PCA approaches can easily lead to overfitting, thus regularization is often required. 5) The algorithms may require heavy computations, especially for large-scale problems. 6) The concept of the PCA basis in the principal subspace is not easily generalized in the presence of missing values. 7) The choice of the dimensionality of the principal subspace is generally more difficult than in classical PCA. Thus, the PCA problem has many features which are usually associated with nonlinear models. This paper discusses some of these important issues providing illustrative examples and presenting ways to overcome possible difficulties.

All the methods studied in this work assume MAR, so they cannot be expected to work in the missing not at random (MNAR) case. Fig. 1 gives a first example of PCA with missing values. The data contain 200 samples generated from a Gaussian distribution represented by the black ellipse. In the first case (left subfigure), the vertical value is missing when the horizontal value is negative. This setting is MAR since the missingness depends only on the observed data. In the second case (right subfigure), the vertical value is missing when the vertical value itself is negative. Now the missingness depends on the missing data, and the missing value mechanism is thus MNAR. As expected, the variational Bayesian PCA (VBPCA, see Section 3.3) works well in the MAR case and gives reasonable reconstructions for the missing values. Also, the vertical part of the bias term $\mathbf{m}$ is estimated to be 1.13 which is much closer to the true value 1.10 than the row-wise mean 1.50 of the observed values in the data.

PCA in the presence of missing values can be relevant for many data sets which appear in practical applications. Sometimes, data sets contain relatively few missing observations and the
problem is to adjust standard PCA algorithms to handle partially observed instances. The choice of the algorithm is not very crucial in such a simple case, as most of the approaches would provide similar solutions. However, in other applications, the available observations can be very sparse and the modeling task is often to reconstruct the missing part from the observed data only. Examples include collaborative filtering problems which is the task of predicting preferences by using other people's preferences (e.g., Hofmann, 2004), and historical data reconstruction in climatic records (e.g., Kaplan et al., 1997).

Historically, the missing value problem in PCA was first studied by Dear (1959) who used only one component and just one imputation iteration (see below). It was based on the minimum meansquare error formulation of PCA introduced by Young (1941). Christoffersson (1970) also used a one-component model for reconstructing missing values. Wiberg (1976) first suggested to directly minimize the mean-square error of the observed part of the data. An algorithm by de Ligny et al. (1981) already worked with up to half of the values missing. The missing values problem using a multivariate normal distribution has been studied even earlier than using PCA, for instance, by Anderson (1957). More historical references can be found in the book by Jolliffe (2002).

More recently, PCA with missing values was studied by Grung and Manne (1998). They proposed using either the imputation or the faster alternating $\mathbf{W}-\mathbf{X}$ algorithm (see below). They discussed the overfitting problem and suggested to delete troublesome rows or columns from data. Tipping and Bishop (1999) introduced the probabilistic formulation of PCA (PPCA). Although they mentioned shortly missing data, they did not provide the formulas to be used for incomplete data. Bishop (1999) introduced variational Bayesian PCA (VBPCA) for choosing the number of components in PCA. Raiko and Valpola (2001) reconstructed missing values with VBPCA to compare some nonlinear models to it. Oba et al. (2003) applied VBPCA for missing value estimation in gene expression profile data, and mentioned that it performs much better than the existing methods for missing value estimation.

The present article reviews possible approaches to the problem with the emphasis on probabilistic models and Bayesian methods. The rest of the article is organized as follows. In Section 2, we present classical algorithms for PCA based on minimizing the reconstruction error similar to (2) and the method based on estimation of the covariance matrix (imputation algorithm). We review how the algorithms are normally used for fully observed data and explain how they can be adapted to the case with missing values. We explain possible difficulties of the standard methods including bad locally optimal solutions, numerical problems and overfitting. Simple examples are given to illustrate these problems. In the same section, we discuss the properties of the imputation algorithm, such as implicit remedies against overfitting and overlearning and its expectation-maximization (EM) interpretation.

Section 3 presents probabilistic models for PCA which provide a good foundation for handling missing values. First, we introduce formulas for the probabilistic PCA model in the case with missing values. Then, we describe Bayesian regularization for handling problems that arise when data are sparse. We provide formulas for performing maximum a posteriori estimation and for variational Bayesian inference.

In Section 4, we propose an extension of the PCA-basis notion to the case of incomplete data. We also show how to find the PCA basis in a principal subspace estimated by different methods. Section 5 provides formulas for computing missing value reconstructions and the variance of the predicted values. We briefly discuss possible ways to select the right number of principal components to obtain reliable reconstructions.

Section 6 discusses large-scale problems and computational complexity. In case of high dimensional and very sparse data, overfitting becomes a severe problem and traditional algorithms for PCA are very slow. We introduce a novel fast optimization algorithm and show how to use it in Bayesian models.

In Section 7, different versions of PCA are compared in artificial experiments, demonstrating the effects of regularization and modeling of posterior variance. We demonstrate the importance of tricks such as fixing hyperparameters for the early learning in variational Bayesian methods. The scalability of the proposed algorithm is demonstrated by applying it to the Netflix problem. Finally, we conclude in Section 8.

### 1.1 Notation

The following notation is used throughout the article. Bold capital letters denote matrices, bold lower-case letters denote vectors, and normal lower-case letters denote scalars. The basic model equation is $\mathbf{y}_{j} \approx \mathbf{W} \mathbf{x}_{j}+\mathbf{m}$, where column vectors $\mathbf{y}_{j}$ are the data cases, $\mathbf{W}$ is the matrix that maps the principal components $\mathbf{x}_{j}$ to the data, and $\mathbf{m}$ is the bias vector. Indices $i=1, \ldots, d$ and $j=1, \ldots, n$ go over the rows and columns of $\mathbf{Y}$, respectively. The index of a principal component is denoted by $k=1, \ldots, c$. Notation $i j$ is used for the index of $y_{i j}$, the $i j$-th element of matrix $\mathbf{Y} . O$ is the set of indices $i j$ corresponding to observed values $y_{i j}, O_{i}$ is the set of indices $j$ (similarly $O_{j}$ is the set of indices $i$ ) for which $y_{i j}$ is observed. $\left|O_{i}\right|$ is the number of elements in $O_{i}$ and $N=|O|$ is the number of observed data elements. Notation $\mathbf{A}_{i:}$ and $\mathbf{A}_{: j}$ is used for the $i$-th row and $j$-th column of a matrix $\mathbf{A}$, respectively. Both $\mathbf{A}_{i:}$ and $\mathbf{A}_{: j}$ are column vectors. The $i$-th row of $\mathbf{W}$ is denoted by $\mathbf{w}_{i}=\mathbf{W}_{i \text { : }}$, and the $j$-th column of $\mathbf{X}$ is $\mathbf{x}_{j}=\mathbf{X}_{: j}$. Table 1 lists the symbols used in the paper.

## 2. Least Squares Techniques

In this section, we present classical algorithms for PCA based on minimizing the reconstruction error similar to (2) and the method based on estimation of the covariance matrix.

### 2.1 The Cost Function

The minimum mean-square error compression of data is the formulation for PCA (Young, 1941) that can be generalized to the case with missing values in a very straightforward manner. The cost function (2) is adapted such that the sum is taken over only those indices $i$ and $j$ for which the data entry $y_{i j}$ is observed (Wiberg, 1976):

$$
\begin{gather*}
C=\sum_{i j \in O}\left(y_{i j}-\hat{y}_{i j}\right)^{2},  \tag{4}\\
\hat{y}_{i j}=\mathbf{w}_{i}^{\mathrm{T}} \mathbf{x}_{j}+m_{i}=\sum_{k=1}^{c} w_{i k} x_{k j}+m_{i} . \tag{5}
\end{gather*}
$$

In this section, we consider algorithms optimizing cost function (4), which we call the least squares (LS) approach to PCA of incomplete data.

The cost function (2) in the fully observed case has practically unique (up to an arbitrary rotation in the principal subspace) global minimum w.r.t. model parameters $\mathbf{W}, \mathbf{X}, \mathbf{m}$ provided that all eigenvalues of the sample covariance matrix are distinct. Srebro and Jaakkola (2003) showed that

| $d$ | dimensionality of the data vectors (indexed by $i$ ) |
| :---: | :---: |
| $n$ | number of data vectors (indexed by $j$ ) |
| c | number of principal components (indexed by $k$ ) |
| $\mathbf{Y}=\left\{y_{i j}\right\}$ | $d \times n$ data matrix |
| $\mathbf{y}_{j}$ | $d$-dimensional data vector ( $j$-th column of $\mathbf{Y}$ ) |
| $\hat{y}_{i j}$ | reconstruction of the data element $y_{i j}$ |
| $\mathbf{m}=\left\{m_{i}\right\}$ | $d$-dimensional bias vector |
| M | $=[\mathbf{m m} \ldots \mathbf{m}], d \times n$ bias matrix |
| $\bar{m}_{i}$ | posterior mean of $m_{i}$ (scalar) |
| $\widetilde{m}_{i}$ | posterior variance of $m_{i}$ (scalar) |
| $\mathbf{W}=\left\{w_{i k}\right\}$ | $d \times c$ matrix for the mapping from principal components to the data |
| $\mathbf{W}_{: k}$ | $k$-th column of matrix $\mathbf{W}$ ( $d$-dimensional vector) |
| $\mathbf{w}_{i}$ | $i$-th row of $\mathbf{W}$ ( $c$-dimensional vector) |
| $\overline{\mathbf{w}}_{i}$ | posterior mean of $\mathbf{w}_{i}$ (c-dimensional vector) |
| $\overline{\mathbf{W}}$ | $=\left[\overline{\mathbf{w}}_{1} \overline{\mathbf{w}}_{2} \ldots \overline{\mathbf{w}}_{d}\right]^{\mathrm{T}}, d \times c$ matrix of posterior means of $\mathbf{w}_{i}$ |
| $\Sigma_{\mathbf{w}_{i}}$ | posterior covariance of $\mathbf{w}_{i}(c \times c$ matrix $)$ |
| $\widetilde{w}_{i k}$ | posterior variance of $w_{i k}$ (scalar) |
| $\mathbf{X}=\left\{x_{k j}\right\}$ | $c \times n$ matrix of principal components |
| $\mathbf{x}_{j}$ | $c$-dimensional principal component vector ( $j$-th column of $\mathbf{X}$ ) |
| $\overline{\mathbf{x}}_{j}$ | posterior mean of $\mathbf{x}_{j}$ ( $c$-dimensional vector) |
| $\overline{\mathbf{X}}$ | $=\left[\overline{\mathbf{x}}_{1} \overline{\mathbf{x}}_{2} \ldots \overline{\mathbf{x}}_{n}\right], c \times n$ matrix of posterior means of $\mathbf{x}_{j}$ |
| $\Sigma_{\mathbf{x}}$ | posterior covariance of $\mathbf{x}_{j}$ (same $c \times c$ matrix for each $j$ ) |
| $\Sigma_{\mathbf{x}_{j}}$ | posterior covariance of $\mathbf{x}_{j}$ (separate $c \times c$ matrix for each $j$ ) |
| $\widetilde{x}_{k j}$ | posterior variance of $x_{k j}$ (scalar) |
| $C$ | cost function to be minimized |
| O | set of indices $i, j$ for which $y_{i j}$ is observed |
| $O_{i}$ | set of indices $j$ for which $y_{i j}$ is observed |
| $O_{j}$ | set of indices $i$ for which $y_{i j}$ is observed |
| $\gamma$ | learning rate (positive scalar) |
| $\varepsilon_{j}$ | $d$-dimensional noise vector ( $j$-th out of $n$ vectors) |
| $v_{y}$ | noise variance (scalar) |
| $\mathcal{N}(\mathbf{x} ; \mu, \Sigma)$ | Gaussian pdf over variable $\mathbf{x}$ with mean $\mu$ and covariance $\Sigma$ |
| $v_{m}$ | prior variance for the elements of $\mathbf{m}$ (scalar) |
| $v_{w, k}$ | prior variance for the elements of $\mathbf{W}_{: k}$ (scalar for each $k$ ) |
| $\xi$ | $=\left(v_{y}, v_{w, k}, v_{m}\right)$ hyperparameters |
| $\theta$ | parameters, typically $\theta=(\mathbf{W}, \mathbf{X}, \mathbf{m})$ |

Table 1: Symbols used in the paper. See also Section 1.1 about notation.
an extension to the weighted case, where each observation has a weight varying from 0 to 1 , is a difficult non-convex problem where the cost function can have multiple local minima. We demonstrate this phenomenon for a simple data set with missing values, that is with weights either 0 (missing value) or 1 (observed value). In our example, model (1) with one principal component and fixed $\mathbf{m}=0$ is fitted to the data

$$
\mathbf{Y}=\left[\begin{array}{cc}
0.8 & 0.8 \\
1 & \times \\
\times & 1
\end{array}\right] \approx\left[\begin{array}{l}
w_{1} \\
w_{2} \\
w_{3}
\end{array}\right]\left[\begin{array}{ll}
x_{1} & x_{2}
\end{array}\right] .
$$

To discard the scaling ambiguity of the PCA solution, the column vector $\mathbf{W}$ can be restricted to have unit length. Thus, the minimized error can be represented as a function of two parameters which define a unit length vector in the three-dimensional space. The cost function in this simple example has three local minima, as shown in Fig. 2: The global minimum $\mathbf{W}= \pm\left[\begin{array}{lll}0.49 & 0.62 & 0.62\end{array}\right]^{\mathrm{T}}$ corresponds to zero cost while the other two minima (close to $\pm\left[\begin{array}{lll}0 & 1 & 0\end{array}\right]^{\mathrm{T}}$ and $\pm\left[\begin{array}{lll}0 & 0 & 1\end{array}\right]^{\mathrm{T}}$ ) provide a non-zero error. Each of the two sub-optimal solutions reconstructs perfectly only three out of four observed values in $\mathbf{Y}$.


Figure 2: Example of local minima for the cost function (4). Data are three-dimensional ( $d=3$ ), the model has one principal component $(c=1)$ and therefore the PCA solution can be defined by a unit length vector $\mathbf{W}$. Left: The cost function plotted on a surface of unit length vectors. Right: The same plot using the Euler vector representation: The matrix $\mathbf{W}$ is constructed from $\alpha$ and $\beta$ as $\mathbf{W}=e^{\mathbf{A}}\left[\begin{array}{lll}1 & 0\end{array}\right]^{\mathrm{T}}$ with $\mathbf{A}$ a $3 \times 3$ matrix with four nonzero elements $a_{12}=-a_{21}=\alpha$ and $a_{13}=-a_{31}=\beta$. Coordinates in brackets correspond to coordinates in Euclidian space in the left hand side plot.

The cost function (4) can be minimized using any optimization procedure. Two possible approaches are presented in the following.

### 2.2 Alternating W-X Algorithm

Complete data. It is possible to optimize the cost function (3), and therefore to find the principal subspace, by updating $\mathbf{W}$ and $\mathbf{X}$ alternately. When either of these matrices is fixed, the other one
can be obtained from an ordinary least squares problem. We will further refer to this approach as the alternating algorithm.

The algorithm alternates between the updates

$$
\begin{align*}
\mathbf{X} & =\left(\mathbf{W}^{\mathrm{T}} \mathbf{W}\right)^{-1} \mathbf{W}^{\mathrm{T}} \mathbf{Y},  \tag{6}\\
\mathbf{W} & =\mathbf{Y X}^{\mathrm{T}}\left(\mathbf{X X}^{\mathrm{T}}\right)^{-1} \tag{7}
\end{align*}
$$

This iteration can be efficient when only a few principal components are needed, that is $c \ll d$ (Roweis, 1998). The bias vector $\mathbf{m}$ is the row-wise mean of the data and the above equations assume that it has been subtracted from the data as a preprocessing step.

Incomplete data. Grung and Manne (1998) studied the alternating algorithm in the case of missing values. In order to get the accurate least squares solution, we include the bias term into the estimation procedure, yielding the update rules

$$
\begin{align*}
\mathbf{x}_{j} & =\left(\sum_{i \in O_{j}} \mathbf{w}_{i} \mathbf{w}_{i}^{\mathrm{T}}\right)^{-1} \sum_{i \in O_{j}} \mathbf{w}_{i}\left(y_{i j}-m_{i}\right), \quad j=1, \ldots, n,  \tag{8}\\
m_{i} & =\frac{1}{\left|O_{i}\right|} \sum_{j \in O_{i}}\left[y_{i j}-\mathbf{w}_{i}^{\mathrm{T}} \mathbf{x}_{j}\right],  \tag{9}\\
\mathbf{w}_{i} & =\left(\sum_{j \in O_{i}} \mathbf{x}_{j} \mathbf{x}_{j}^{\mathrm{T}}\right)^{-1} \sum_{j \in O_{i}} \mathbf{x}_{j}\left(y_{i j}-m_{i}\right), \quad i=1, \ldots, d, \tag{10}
\end{align*}
$$

where $m_{i}$ is the $i$-th element of $\mathbf{m}$.

### 2.3 Gradient Descent Algorithm

Complete data. The basis of neural network implementation of PCA learning rules is a gradient descent optimization procedure. Such algorithms work online, processing only one input vector $\mathbf{y}_{j}$ at once. The learning rules implement stochastic gradient descent and the algorithms will eventually converge to a basis in the principal subspace. Each data vector may have to be processed several times for convergence. The same can also be implemented in a batch procedure.

Using gradient descent for minimization of (4) w.r.t. $\mathbf{W}$ yields the update rule

$$
\begin{equation*}
\mathbf{W} \leftarrow \mathbf{W}+\gamma(\mathbf{Y}-\mathbf{W} \mathbf{X}) \mathbf{X}^{\mathrm{T}}, \tag{11}
\end{equation*}
$$

where $\gamma>0$ is called the learning rate. Minimization w.r.t. matrix $\mathbf{X}$ can be performed using the least squares solution in (6). Some neural algorithms use learning rules which either explicitly orthogonalize $\mathbf{W}$ or which yield an orthogonal $\mathbf{W}$ at the convergence. Then the update of $\mathbf{X}$ can be simplified to $\mathbf{X}=\mathbf{W}^{\mathrm{T}} \mathbf{Y}$, which together with (11) is the batch version of the Oja learning algorithm ( Oja , 1983; Diamantaras and Kung, 1996). The bias $\mathbf{m}$ is again removed from the data as a preprocessing step.

Incomplete data. In the presence of missing values, the gradient descent update rule for $\mathbf{W}$ is

$$
\mathbf{W} \leftarrow \mathbf{W}-\gamma \frac{\partial C}{\partial \mathbf{W}}, \quad \text { with } \quad \frac{\partial C}{\partial w_{i k}}=-2 \sum_{j \in O_{i}}\left(y_{i j}-\hat{y}_{i j}\right) x_{k j}
$$

where $\hat{y}_{i j}$ is given in (5). Matrix $\mathbf{X}$ could be updated using (8) but a gradient-based update can also be used:

$$
\mathbf{X} \leftarrow \mathbf{X}-\gamma \frac{\partial C}{\partial \mathbf{X}}, \quad \text { with } \quad \frac{\partial C}{\partial x_{k j}}=-2 \sum_{i \in O_{j}}\left(y_{i j}-\hat{y}_{i j}\right) w_{i k} .
$$

The bias term $\mathbf{m}$ can be updated using (9). As we discuss in Section 6.1, gradient-based learning can be computationally more efficient than the alternating algorithm because there is no need to compute matrices $\left(\sum_{i \in O_{j}} \mathbf{w}_{i} \mathbf{w}_{i}^{\mathrm{T}}\right)^{-1}$ and $\left(\sum_{j \in O_{i}} \mathbf{x}_{j} \mathbf{x}_{j}^{\mathrm{T}}\right)^{-1}$ in (8) and (10).

### 2.4 Numerical Problems and Overfitting

Least squares PCA methods can work well for data sets with few missing values but they may not be applicable to sparse data sets because of severe problems with overfitting. Suppose that for some $j$ the number $\left|O_{j}\right|$ of observed measurements $y_{i j}$ is smaller than the number of principal components $c$. Then, the corresponding least square problem is ill posed: the matrix $\sum_{i \in O_{j}} \mathbf{w}_{i} \mathbf{w}_{i}^{\mathrm{T}}$ is rank deficient and Equation (8) cannot be used. Moreover, even if $\left|O_{j}\right|$ is greater than $c$, matrix $\sum_{i \in O_{j}} \mathbf{w}_{i} \mathbf{w}_{i}^{\mathrm{T}}$ may be badly conditioned and the corresponding $\mathbf{x}_{j}$ can become infinitely large. This means that the parameters would be overfitted to explain well a few observed values but the generalization ability of the model would be poor (e.g., reconstruction of missing data would be very inaccurate). This is exactly what happens in the vicinity of the two local minima in the example in Fig. 2.

The same problem can happen when the rows of $\mathbf{W}$ are estimated using, for example, (10): matrix $\sum_{j \in O_{i}} \mathbf{x}_{j} \mathbf{x}_{j}^{\mathrm{T}}$ can be rank deficient or badly conditioned. This is more probable when some rows of $\mathbf{Y}$ contain very few measurements. However, these problems can generally appear for any data set and for any algorithm optimizing the cost function (4).

Overfitting problems can be avoided by using proper regularization. A common way to prevent unbounded growth of model parameters is to add terms which penalize large parameter values into the cost function. The amount of regularization can be determined, for instance, by cross-validation. Another possibility is to use probabilistic methods (see Section 3) in which the extra penalty terms come naturally from a probabilistic model.

### 2.5 Method Using Singular Value Decomposition

Complete data. Perhaps the most popular approach to PCA is based on singular value decomposition (SVD) of the data matrix or (equivalently) eigen-decomposition of the sample covariance matrix. SVD of the data matrix is given by:

$$
\mathbf{Y}=\mathbf{U} \Sigma \mathbf{V}^{\mathrm{T}}
$$

where $\mathbf{U}$ is a $d \times d$ orthogonal matrix, $\mathbf{V}$ is an $n \times n$ orthogonal matrix and $\Sigma$ is a $d \times n$ pseudodiagonal matrix (diagonal if $d=n$ ) with the singular values on the main diagonal (e.g., Haykin, 1989). The PCA solution is obtained by selecting the $c$ largest singular values from $\Sigma$, by forming $\mathbf{W}$ from the corresponding $c$ columns of $\mathbf{U}$, and $\mathbf{X}$ from the corresponding $c$ rows of $\Sigma \mathbf{V}^{\mathrm{T}}$ (Jolliffe, 2002). The bias $\mathbf{m}$ is again removed from the data as a preprocessing step.

PCA can equivalently be defined using the eigen-decomposition of the $d \times d$ covariance matrix $\mathbf{C}$ of the column vectors of the data matrix $\mathbf{Y}$ :

$$
\mathbf{C}=\frac{1}{n} \mathbf{Y} \mathbf{Y}^{\mathrm{T}}=\mathbf{U D U}^{\mathrm{T}} .
$$

The diagonal matrix $\mathbf{D}$ contains the eigenvalues of $\mathbf{C}$, and the columns of the matrix $\mathbf{U}$ contain the unit-length eigenvectors of $\mathbf{C}$ in the same order (Cichocki and Amari, 2002; Diamantaras and Kung, 1996; Jolliffe, 2002; Haykin, 1989). Again, the columns of $\mathbf{U}$ corresponding to the largest eigenvalues are taken as $\mathbf{W}$, and $\mathbf{X}$ is computed as $\mathbf{W}^{\mathrm{T}} \mathbf{Y}$. This approach can be more efficient for cases where $d \ll n$, since it avoids the computation of the $n \times n$ matrix $\mathbf{V}$.

Incomplete data. The same approach cannot be directly used in the presence of missing values. Estimating the covariance matrix $\mathbf{C}$ becomes difficult. Let us consider a simple (but incorrect) estimate where we just leave out terms with missing values from the average for each element of $\mathbf{C}$. For the data matrix below, we get

$$
\mathbf{Y}=\left[\begin{array}{ccccc}
-1 & +1 & 0 & 0 & \times \\
-1 & +1 & \times & \times & 0 \\
\times & \times & -1 & +1 & \times
\end{array}\right], \quad \mathbf{C}=\frac{1}{n} \mathbf{Y} \mathbf{Y}^{\mathrm{T}}=\left[\begin{array}{ccc}
0.5 & 1 & 0 \\
1 & 0.667 & \times \\
0 & \times & 1
\end{array}\right] .
$$

There are at least two problems. First, the estimated covariance 1 between the first and second components is larger than their estimated variances 0.5 and 0.667 . This clearly leads to the situation where the covariance matrix is not positive (semi)definite and some of its eigenvalues are negative. Secondly, the covariance between the second and the third component could not be estimated at all. For these reasons, this is a viable option only if the number of missing values is not significant. There are many algorithms for finding a proper estimate of the covariance matrix (e.g., Ghahramani and Jordan, 1994; Boscardin and Zhang, 2004) but they are computationally intensive iterative algorithms.

A simple alternative is an iterative procedure which performs PCA more directly. It alternates between imputing the missing values in $\mathbf{Y}$ and applying standard PCA to the infilled (complete) data matrix $\mathbf{Y}_{c}$ (e.g., Jolliffe, 2002). Initially, the missing values can be replaced, for example, by the row-wise means of $\mathbf{Y}$. The covariance matrix of the complete data $\mathbf{Y}_{c}$ can be estimated without the problems mentioned above, and $\mathbf{W}$ can be computed using its eigen-decomposition. The bias term $\mathbf{m}$ can be updated as the row-wise mean of $\mathbf{Y}_{c}$. Next, the principal components $\mathbf{X}$ are calculated:

$$
\mathbf{X}=\mathbf{W}^{\mathrm{T}}\left(\mathbf{Y}_{c}-\mathbf{M}\right),
$$

the reconstructions can be used as a better estimate for the missing values:

$$
\mathbf{Y}_{c}= \begin{cases}\mathbf{Y} & \text { for observed values } \\ \mathbf{W} \mathbf{X}+\mathbf{M} & \text { for missing values }\end{cases}
$$

and PCA can be applied again to the updated data matrix $\mathbf{Y}_{c}$. This process can be iterated until convergence. We will further refer to this approach as the imputation algorithm. In Appendix A, we show that the imputation algorithm also minimizes the cost function (4) and that it implements the EM steps for a simple probabilistic model. This approach requires the use of the complete data matrix, and therefore it is computationally very expensive for large-scale problems.

Although the imputation algorithm belongs to the class of least-squares PCA algorithms (i.e., it does not use explicit regularization), it still provides some remedy against overfitting. Performing SVD of a complete data matrix is equivalent to minimizing a cost function which is a sum of two terms: the reconstruction error for the observed values (which is the true minimized error) and the reconstruction error for the infilled missing data (which is a penalty term forcing the new reconstructions of missing data to be close to the infilled values). Initialization of the reconstructions
with the row-wise means of the data matrix introduces a bias in favor of the most "conservative" solutions. In practice, there can be infinitely many solutions that provide the same reconstruction error for the observed part but that have different reconstructions for the missing part. In such cases, the effect of the initialization can last even if the algorithm is iterated indefinitely.

Our experiments show that the imputation algorithm also has resistance against overlearning. By overlearning we mean situations when the training error decreases but the generalization ability of the model gets worse. Badly overfitted solutions correspond to regions in the parameter space where a small decrease of the training error can cause a large increase of the test error. However, the penalty term (which keeps the reconstructions of missing data close to the infilled values) makes the steps shorter and learning can practically stop once the algorithm enters regions where the training error changes very little. Thus, this regularization effect is due to early stopping.

## 3. Probabilistic Models for PCA

In this section, we presents probabilistic models for PCA which provide a good foundation for handling missing values.

### 3.1 Probabilistic PCA

The probabilistic formulation of PCA offers a number of benefits, including well-founded regularization, model comparison, interpretation of results, and extendability. Probabilistic PCA (Tipping and Bishop, 1999) explicitly includes the noise term in a generative model

$$
\begin{equation*}
\mathbf{y}_{j}=\mathbf{W} \mathbf{x}_{j}+\mathbf{m}+\varepsilon_{j} . \tag{12}
\end{equation*}
$$

Both the principal components $\mathbf{x}_{j}$ and the noise $\varepsilon_{j}$ are assumed normally distributed:

$$
\begin{align*}
& p\left(\mathbf{x}_{j}\right)=\mathfrak{N}\left(\mathbf{x}_{j} ; 0, \mathbf{I}\right),  \tag{13}\\
& p\left(\boldsymbol{\varepsilon}_{j}\right)=\mathcal{N}\left(\boldsymbol{\varepsilon}_{j} ; 0, v_{\mathbf{y}} \mathbf{I}\right), \tag{14}
\end{align*}
$$

where $\mathcal{N}(\mathbf{x} ; \mu, \Sigma)$ denotes the normal probability density function (pdf) over variable $\mathbf{x}$ with mean $\mu$ and covariance $\Sigma$. The parameters of the model include $\mathbf{W}, \mathbf{m}$ and $v_{y}$. The model can be identified by finding the maximum likelihood (ML) estimate for the model parameters using the EM algorithm (Dempster et al., 1977).

Complete data. The E-step estimates the conditional distribution of the hidden variables $\mathbf{X}$ given the data and the current values of the model parameters:

$$
p\left(\mathbf{X} \mid \mathbf{Y}, \mathbf{W}, v_{y}\right)=\prod_{j=1}^{n} \mathcal{N}\left(\mathbf{x}_{j} ; \overline{\mathbf{x}}_{j}, \Sigma_{\mathbf{x}_{j}}\right),
$$

where the covariance matrix $\Sigma_{\mathbf{x}_{j}}$ is same for all $\mathbf{x}_{j}$ :

$$
\Sigma_{\mathbf{x}_{j}}=\Sigma_{\mathbf{x}}=v_{y}\left(v_{y} \mathbf{I}+\mathbf{W}^{\mathrm{T}} \mathbf{W}\right)^{-1}, \quad j=1, \ldots n
$$

and the means $\overline{\mathbf{x}}_{j}$ can be summarized in the columns of matrix

$$
\begin{equation*}
\overline{\mathbf{X}}=\frac{1}{v_{y}} \Sigma_{\mathbf{x}} \mathbf{W}^{\mathrm{T}} \mathbf{Y} . \tag{15}
\end{equation*}
$$

The M-step re-estimates the model parameters as

$$
\begin{align*}
\mathbf{W} & =\mathbf{Y} \overline{\mathbf{X}}^{\mathrm{T}}\left(\overline{\left.\mathbf{X} \mathbf{X}^{\mathrm{T}}+n \Sigma_{\mathbf{x}}\right)^{-1},}\right.  \tag{16}\\
v_{y} & =\frac{1}{n d} \sum_{i=1}^{d} \sum_{j=1}^{n}\left(y_{i j}-\mathbf{w}_{i}^{\mathrm{T}} \overline{\mathbf{x}}_{j}\right)^{2}+\frac{1}{d} \operatorname{tr}\left(\mathbf{W} \Sigma_{\mathbf{x}} \mathbf{W}^{\mathrm{T}}\right) .
\end{align*}
$$

Note that $v_{y}$ is updated as the mean-square error plus the extra term which accounts for the uncertainty in $\mathbf{X}$.

Tipping and Bishop (1999) showed that $\mathbf{W}$ converges to the principal subspace. In the zeronoise limit, that is for $v_{y} \rightarrow 0$, the iterations (15)-(16) reduce to the least squares projections in (6)-(7) (Roweis, 1998). Note also that the ML estimator for $\mathbf{m}$ is given by the row-wise mean of the data (Tipping and Bishop, 1999). Therefore, it can be computed once and removed from the data in the preprocessing step.

Incomplete data. The generalization to the case with missing values happens as follows. First we write an element-wise version of the probabilistic model in (12)-(14):

$$
\begin{aligned}
y_{i j} & =\mathbf{w}_{i}^{\mathrm{T}} \mathbf{x}_{j}+m_{i}+\varepsilon_{i j}, \quad \forall i j \in O \\
p\left(\mathbf{x}_{j}\right) & =\mathcal{N}\left(\mathbf{x}_{j} ; 0, \mathbf{I}\right), \\
p\left(\varepsilon_{i j}\right) & =\mathcal{N}\left(\varepsilon_{i j} ; 0, v_{y}\right) .
\end{aligned}
$$

We treat $\mathbf{w}_{i}, m_{i}$, and $v_{y}$ as model parameters and $\mathbf{x}_{j}$ as latent variables and apply the standard EM algorithm to arrive at the following update rules:

$$
\begin{align*}
\Sigma_{\mathbf{x}_{j}} & =v_{y}\left(v_{y} \mathbf{I}+\sum_{i \in O_{j}} \mathbf{w}_{i} \mathbf{w}_{i}^{\mathrm{T}}\right)^{-1},  \tag{17}\\
\overline{\mathbf{x}}_{j} & =\frac{1}{v_{y}} \Sigma_{\mathbf{x}_{j}} \sum_{i \in O_{j}} \mathbf{w}_{i}\left(y_{i j}-m_{i}\right), \quad j=1, \ldots, n, \\
m_{i} & =\frac{1}{\left|O_{i}\right|} \sum_{j \in O_{i}}\left[y_{i j}-\mathbf{w}_{i}^{\mathrm{T}} \overline{\mathbf{x}}_{j}\right], \\
\mathbf{w}_{i} & =\left(\sum_{j \in O_{i}}\left[\overline{\mathbf{x}}_{j} \overline{\mathbf{x}}_{j}^{\mathrm{T}}+\Sigma_{\mathbf{x}_{j}}\right]\right)^{-1} \sum_{j \in O_{i}} \overline{\mathbf{x}}_{j}\left(y_{i j}-m_{i}\right), \quad i=1, \ldots, d,  \tag{18}\\
v_{y} & =\frac{1}{N} \sum_{i j \in O}\left[\left(y_{i j}-\mathbf{w}_{i}^{\mathrm{T}} \overline{\mathbf{x}}_{j}-m_{i}\right)^{2}+\mathbf{w}_{i}^{\mathrm{T}} \Sigma_{\mathbf{x}_{j}} \mathbf{w}_{i}\right] .
\end{align*}
$$

There are several distinctions compared to the fully observed data: 1) The optimal $\mathbf{m}$ depends on other parameters and therefore it has to be updated in the iterative procedure. 2) The covariance $\Sigma_{\mathbf{x}_{j}}$ is different for each $\mathbf{x}_{j}$. 3) Each row of $\mathbf{W}$ is recomputed based only on those columns of $\overline{\mathbf{X}}$ which contribute to the reconstruction of the observed values in the corresponding row of the data matrix. The same applies to the computation of the columns of $\overline{\mathbf{X}}$. Thus, the computations required for incomplete data are generally heavier.

### 3.2 Examples of Overfitting with Probabilistic PCA

PPCA provides some remedy against overfitting, in contrast to the least squares approach. First, a nonzero noise level $v_{y}$ regularizes badly conditioned matrices $\sum_{i \in O_{j}} \mathbf{w}_{i} \mathbf{w}_{i}^{\mathrm{T}}$ in (17). Second, even
if the noise level is small, badly conditioned matrices $\sum_{i \in O_{j}} \mathbf{w}_{i} \mathbf{w}_{i}^{\mathrm{T}}$ result in large values both in the means $\overline{\mathbf{x}}_{j}$ and in the covariance matrices $\Sigma_{\mathbf{x}_{j}}$. This diminishes the effect of large values in $\overline{\mathbf{x}}_{j}$ when $\mathbf{W}$ is re-estimated using (18).

However, PPCA can overfit, for example, if the estimated number of principal components is unreasonably large. In the example in Fig. 3, PPCA with one ( $c=1$ ) principal component is applied to two-dimensional $(d=2)$ data and each data vector is only partly observed (i.e., either $y_{1 j}$ or $y_{2 j}$ is known for all $j \mathrm{~s}$ ). The observations are represented by triangles placed on the two axes in Fig. 3. The data points reconstructed with a trained PPCA model lie on a line, thus the model invents correlations that are not observed in the data. In fact, there are infinitely many PPCA models which are equally good and the final solution (and hence the reconstructions) depends on initialization. The situation would not change significantly if there were a few fully observed data vectors. Then, the trained model would be defined by those few samples and it would not generalize well for new data.

A similar overfitting example is presented in Fig. 4. There, three-dimensional data are described well by a model with $c=1$ principal component. The first two rows of $\mathbf{Y}$ are fully observed while there are only two measurements in the third row (Fig. 4a). The reconstructions of the third row of $\mathbf{Y}$ are very inaccurate because the model relies only on the two available observations.

### 3.3 Variational Bayesian PCA (VBPCA)

A common way to cope with the overfitting problem is to penalize parameter values which correspond to more complex explanations of the data. A natural regularization in PCA is using penalization of large values in matrices $\mathbf{W}$ and $\mathbf{X}$. In the Bayesian formulation, this is equivalent to introducing a prior over the model parameters. For example, the PPCA model in (12)-(14) can be complemented with Gaussian priors over the elements of vector $\mathbf{m}$ and matrix $\mathbf{W}$ :

$$
\begin{align*}
& p(\mathbf{m})=\mathcal{N}\left(\mathbf{m} ; 0, v_{m} \mathbf{I}\right)  \tag{19}\\
& p(\mathbf{W})=\prod_{k=1}^{c} \mathcal{N}\left(\mathbf{W}_{: k} ; 0, v_{w, k} \mathbf{I}\right) . \tag{20}
\end{align*}
$$

Here, we use a zero mean prior for $\mathbf{m}$ for the sake of simplicity. Including a mean hyperparameter $\mu$, that is $p(\mathbf{m})=\prod_{i} \mathcal{N}\left(m_{i} ; \mu, v_{m}\right)$, can be useful in practice.

The model (20) uses a shared prior for all elements in the same column of $\mathbf{W}$, parameterized with $v_{w, k}$. This is done to allow automatic selection of the right number of components needed for PCA. The hyperparameters $v_{m}, v_{w, k}$ can also be updated during learning (e.g., using the evidence framework or variational approximations). If the evidence of the relevance of the $k$-th principal component for reliable data modeling is weak, the corresponding $v_{w, k}$ should tend to zero. This is called automatic relevance determination (e.g., Bishop, 2006). The prior in (19)-(20) also introduces a bias towards the PCA basis within the principal subspace (Luttinen and Ilin, 2010).

Let us assume that we perform ML estimation of the hyperparameters $\boldsymbol{\xi}=\left(v_{y}, v_{w, k}, v_{m}\right)$ in the probabilistic model defined in (12)-(14) and (19)-(20). This can be done using the EM-algorithm if one treats the model parameters $\theta=(\mathbf{W}, \mathbf{X}, \mathbf{m})$ as hidden random variables. Implementation of the EM algorithm would require the computation of the posterior of the hidden variables $p(\theta \mid \mathbf{Y}, \xi)$ on the E-step. Unfortunately, the true posterior $p(\theta \mid \mathbf{Y}, \boldsymbol{\xi})$ does not have an analytic form and one possible solution is to approximate it with a simpler pdf $q(\theta)$. This is justified by the variational view of the EM algorithm (Neal and Hinton, 1999).


Figure 3: A simple example of overfitting with PPCA: data are two-dimensional and each sample contains only one observed value. The measurements are marked with blue and red triangles at the two axes. The reconstructions provided by a trained PPCA model are shown with crosses.


Figure 4: Example of inaccurate reconstruction with PPCA. Left: The "true" values in the three rows of $\mathbf{Y}$ are shown with blue lines. The first two rows are fully observed, the only two observations in the third row are marked with circles. The PPCA reconstruction is shown with red dots. Right: The blue dots is the scatter plot of the third row of $\mathbf{Y}$ (x-axis) against the principal component estimated by PPCA (y-axis). The dots corresponding to the two observations are marked with circles. The red crosses is the reconstruction of the third row (x-axis) against the estimated principal component (y-axis). The estimated correlation are due to the two fully observed data vectors.

Using the variational approach (Wallace, 1990; Hinton and van Camp, 1993), the E-step is modified to update the approximation $q(\theta)$ so as to minimize the cost function

$$
\begin{equation*}
C(q(\theta), \xi)=\int q(\theta) \log \frac{q(\theta)}{p(\mathbf{Y}, \theta \mid \xi)} d \theta=\int q(\theta) \log \frac{q(\theta)}{p(\theta \mid \mathbf{Y}, \xi)} d \theta-\log p(\mathbf{Y} \mid \xi) \tag{21}
\end{equation*}
$$

On the M-step, the approximation $q(\theta)$ is used as it was the actual posterior $p(\theta \mid \mathbf{Y}, \xi)$ in order to increase likelihood $p(\mathbf{Y} \mid \xi)$. This can be seen as minimization of (21) w.r.t. $\xi$ (Neal and Hinton, 1999).

The first term in (21) is the Kullback-Leibler divergence between the true posterior and its approximation. Since it is always non-negative, the cost function provides a lower bound of the log-likelihood:

$$
\begin{equation*}
\log p(\mathbf{Y} \mid \xi) \geq-C(q(\theta), \xi) \tag{22}
\end{equation*}
$$

This property can be used to compare solutions corresponding to different local minima of (21). More details on variational methods can be found, for example, in the book by Bishop (2006).

The variational Bayesian (VB) approach to PCA can be implemented by minimization of (21) w.r.t. the approximating $\operatorname{pdf} q(\theta)$ and $\xi .{ }^{1}$ The complexity of the cost function (21) depends on the form of the approximating pdf $q(\theta)$. A computationally convenient form for the PCA model is

$$
\begin{equation*}
q(\theta)=\prod_{i=1}^{d} q\left(m_{i}\right) \prod_{i=1}^{d} q\left(\mathbf{w}_{i}\right) \prod_{j=1}^{n} q\left(\mathbf{x}_{j}\right) . \tag{23}
\end{equation*}
$$

Then, the cost function can be minimized alternately w.r.t. one factor $q\left(\theta_{i}\right)$ in (23) while keeping the other ones fixed. Because we use conjugate priors in (19)-(20), it is possible to find optimal pdfs $q\left(\theta_{i}\right)$ on each step: The optimal $q\left(m_{i}\right), q\left(\mathbf{w}_{i}\right)$ and $q\left(\mathbf{x}_{j}\right)$ are Gaussian and their update boils down to re-estimation of the corresponding means and covariance matrices.

In Appendix C, we present the update rules for the resulting algorithm, which we call VBPCA. There, we assume incomplete data sets, the update rules for the fully observed data can be found in the paper by Bishop (1999). Note the resemblance of the learning rules to the EM algorithm applied to PPCA. The mean parameters $\overline{\mathbf{W}}, \overline{\mathbf{X}}, \overline{\mathbf{m}}$ of the approximating pdfs can be used as estimates of the corresponding model parameters. The covariance matrices $\Sigma_{\mathbf{w}_{i}}, \Sigma_{\mathbf{x}_{j}}$ and variances $\widetilde{m}_{i}$ reflect the uncertainty about the corresponding quantities.

The advantages of the VB approach can be summarized in the following.

- VB learning is sensitive to posterior probability mass rather than posterior probability density, which makes it more resistant against overfitting compared to point estimation (e.g., PPCA or MAPPCA presented in Section 6.4). Typically, overfitted solutions correspond to high peaks in the posterior density while robust solutions contain much of the posterior probability mass in their neighborhood.
- The method provides information about uncertainty for the unknown quantities (in the estimated posterior covariance matrices). This can be useful to detect unreliable results similar to the one presented in Fig. 4. For example, the uncertainty of the reconstructions of missing values can be estimated as we show in Section 5.

1. In the fully Bayesian treatment, the hyperparameters $\xi$ are also assigned priors and they are treated equally to the rest of the parameters $\theta$. We omit the effect of the prior for the hyperparameters to simplify the equations.

- The VB cost function can be used for comparison between different solutions based on (22): A smaller cost yields a greater lower bound of the solution evidence. This is a useful feature for PCA of incomplete data as, in this problem, even simpler models can provide cost functions with multiple local minima. A greater lower bound does not guarantee a better solution, but still it seems to be a reliable measure in practice, see Figure 11 for an example.


## 4. Rotation in the Principal Subspace to the PCA Basis

A useful property of classical PCA is that the principal components are ordered by the amount of data variance they explain, which allows an intuitive interpretation of the results. Such a representation is quite trivial for complete data sets but the extensions to the case of incomplete data is not straightforward. In this section, we show how to find such a basis in the principal subspace estimated by different algorithms for incomplete data. We refer to this basis as the PCA basis in the principal subspace.

One can define the PCA basis for complete data as the solution which minimizes (2) and which satisfies the following conditions: 1) the principal components (in the rows of $\mathbf{X}$ ) are zero mean, mutually uncorrelated and scaled to unit variance:

$$
\begin{gather*}
\frac{1}{n} \sum_{j=1}^{n} \mathbf{x}_{j}=0  \tag{24}\\
\frac{1}{n} \sum_{j=1}^{n} \mathbf{x}_{j} \mathbf{x}_{j}^{\mathrm{T}}=\frac{1}{n} \mathbf{X} \mathbf{X}^{\mathrm{T}}=\mathbf{I} \tag{25}
\end{gather*}
$$

and 2) matrix $\mathbf{W}$ has mutually orthogonal columns which are ordered according to their norms:

$$
\begin{align*}
& \mathbf{W}^{\mathrm{T}} \mathbf{W}=\operatorname{diag}(\mathbf{s}),  \tag{26}\\
& s_{k} \geq s_{l}, \quad k<l \tag{27}
\end{align*}
$$

where $\operatorname{diag}(\mathbf{s})$ denotes a diagonal matrix with diagonal elements $s_{k}=\left\|\mathbf{W}_{: k}\right\|^{2}$. Note that the normalized columns of $\mathbf{W}$ and their squared norms $s_{k}$ are the eigenvectors and eigenvalues of the data covariance matrix.

We propose to use the same conditions (24)-(27) to define the PCA basis in the principal subspace estimated for incomplete data. Note the important difference to the complete data case: The PCA basis for complete data is unique (assuming distinct eigenvalues of the data covariance matrix) and it does not depend on the dimensionality for the assumed principal subspace. This cannot be guaranteed for incomplete data case: The principal subspace estimated with the same algorithm using fewer components may differ from the leading directions of the PCA basis found in the subspace with more components.

In the following, we show how to transform a solution found with different algorithms to the PCA basis such that conditions (24)-(27) are satisfied.

### 4.1 Least Squares Approaches

The first step to transform a solution $\{\mathbf{W}, \mathbf{m}, \mathbf{X}\}$ found by a least squares algorithm (as presented in Sections 2.2-2.3) is to center the rows in $\mathbf{X}$ to zero mean:

$$
\begin{gather*}
\mathbf{x}_{j} \leftarrow \mathbf{x}_{j}-\mu,  \tag{28}\\
\mathbf{m}_{\mathrm{pca}}=\mathbf{m}+\mathbf{W} \mu, \tag{29}
\end{gather*}
$$

with $\mu=\frac{1}{n} \sum_{j=1}^{n} \mathbf{x}_{j}$. This ensures (24). The next step is to transform linearly $\mathbf{W}$ and $\mathbf{X}$ :

$$
\begin{align*}
\mathbf{W}_{\mathrm{pca}} & =\mathbf{W} \mathbf{U} \mathbf{D}_{x}^{1 / 2} \mathbf{V},  \tag{30}\\
\mathbf{X}_{\mathrm{pca}} & =\mathbf{V}^{\mathrm{T}} \mathbf{D}_{x}^{-1 / 2} \mathbf{U}^{\mathrm{T}} \mathbf{X}, \tag{31}
\end{align*}
$$

where matrices $\mathbf{U}, \mathbf{D}_{x}$ are computed by eigen-decomposition

$$
\begin{equation*}
\frac{1}{n} \mathbf{X} \mathbf{X}^{\mathrm{T}}=\mathbf{U} \mathbf{D}_{x} \mathbf{U}^{\mathrm{T}} \tag{32}
\end{equation*}
$$

and $\mathbf{V}$ is calculated by eigen-decomposition

$$
\begin{equation*}
\mathbf{D}_{x}^{1 / 2} \mathbf{U}^{\mathrm{T}} \mathbf{W}^{\mathrm{T}} \mathbf{W} \mathbf{U} \mathbf{D}_{x}^{1 / 2}=\mathbf{V} \mathbf{D}_{w} \mathbf{V}^{\mathrm{T}} . \tag{33}
\end{equation*}
$$

It is easy to show that the transformed solution satisfies (25)-(26):

$$
\begin{aligned}
\frac{1}{n} \mathbf{X}_{\mathrm{pca}} \mathbf{X}_{\mathrm{pca}}^{\mathrm{T}} & =\frac{1}{n} \mathbf{V}^{\mathrm{T}} \mathbf{D}_{x}^{-1 / 2} \mathbf{U}^{\mathrm{T}} \mathbf{X} \mathbf{X}^{\mathrm{T}} \mathbf{U} \mathbf{D}_{x}^{-1 / 2} \mathbf{V} \\
& =\mathbf{V}^{\mathrm{T}} \mathbf{D}_{x}^{-1 / 2} \mathbf{U}^{\mathrm{T}} \mathbf{U} \mathbf{D}_{x} \mathbf{U}^{\mathrm{T}} \mathbf{U} \mathbf{D}_{x}^{-1 / 2} \mathbf{V}=\mathbf{V}^{\mathrm{T}} \mathbf{D}_{x}^{-1 / 2} \mathbf{D}_{x} \mathbf{D}_{x}^{-1 / 2} \mathbf{V}=\mathbf{V}^{\mathrm{T}} \mathbf{V}=\mathbf{I}, \\
\mathbf{W}_{\mathrm{pca}}^{\mathrm{T}} \mathbf{W}_{\mathrm{pca}} & =\mathbf{V}^{\mathrm{T}} \mathbf{D}_{x}^{1 / 2} \mathbf{U}^{\mathrm{T}} \mathbf{W}^{\mathrm{T}} \mathbf{W} \mathbf{U D}_{x}^{1 / 2} \mathbf{V}=\mathbf{V}^{\mathrm{T}} \mathbf{V} \mathbf{D}_{w} \mathbf{V}^{\mathrm{T}} \mathbf{V}=\mathbf{D}_{w}
\end{aligned}
$$

### 4.2 Probabilistic PCA

We show in Appendix B that the following conditions hold at the convergence of PPCA:

$$
\begin{gather*}
\frac{1}{n} \sum_{j=1}^{n} \overline{\mathbf{x}}_{j}=0,  \tag{34}\\
\Sigma_{*}=\frac{1}{n} \sum_{j=1}^{n}\left[\overline{\mathbf{x}}_{j} \overline{\mathbf{x}}_{j}^{\mathrm{T}}+\Sigma_{\mathbf{x}_{j}}\right]=\mathbf{I}, \tag{35}
\end{gather*}
$$

which can be seen as the analogue of (24)-(25). Note that $\Sigma_{*}$ becomes the sample covariance matrix of the principal components in the noiseless limit, when $v_{y} \rightarrow 0$. The extra requirements (26)-(27) can be used to define the (practically) unique PCA basis for a PPCA solution. They resolve the rotational ambiguity caused by the fact that any orthogonal transformation of the principal subspace does not change $\Sigma_{*}$ in (35).

One can easily perform a transformation of a PPCA solution such that (34)-(35) and (26)-(27) hold. This is done by first normalizing $\overline{\mathbf{x}}_{j}$ to zero mean and updating the bias term similarly to (28)(29). Then, a rotation of the subspace is performed similarly to (30)-(33) with the exception that
$\frac{1}{n} \mathbf{X} \mathbf{X}^{\mathrm{T}}$ is replaced with $\Sigma_{*}$. Finally, one can rotate the posterior covariance matrices of the estimated principal components:

$$
\begin{equation*}
\Sigma_{\mathbf{x}_{j} \mathrm{pca}}=\mathbf{V}^{\mathrm{T}} \mathbf{D}_{x}^{-1 / 2} \mathbf{U}^{\mathrm{T}} \Sigma_{\mathbf{x}_{j}} \mathbf{U} \mathbf{D}_{x}^{-1 / 2} \mathbf{V} \tag{36}
\end{equation*}
$$

which follows from (31).
Although the optimal PPCA solution satisfy at least conditions (34)-(35), convergence of the learning algorithm to the optimal solution can be very slow. Therefore, performing the described transformations during learning can speed up learning. In the experiments, we use the transformations after each iteration.

### 4.3 VBPCA

It can be shown (Luttinen and Ilin, 2010) that VBPCA described in Section 3.3 converges to the solution which satisfies (34)-(35) and the condition analogous to (26):

$$
\begin{equation*}
\sum_{i=1}^{d}\left[{\overline{\mathbf{w}_{i} \mathbf{w}_{i}}}^{\mathrm{T}}+\Sigma_{\mathbf{w}_{i}}\right]=\operatorname{diag}(\mathbf{s}), \tag{37}
\end{equation*}
$$

where $\overline{\mathbf{w}_{i}}, \Sigma_{\mathbf{w}_{i}}$ are the posterior means and covariance matrices for the rows of matrix $\mathbf{W}$ (see Appendix $C$ for notation details). Thus, one can simply use the ordering requirement (27) to define a meaningful basis for a VBPCA solution. Performing a transformation that ensures (34)-(35) and (37) during learning can speed up convergence (Luttinen and Ilin, 2010). The transformation can be performed similarly to (30)-(33) with the exception that $\mathbf{X} \mathbf{X}^{\mathrm{T}}$ is replaced with $\sum_{j=1}^{n}\left[\overline{\mathbf{x}}_{j} \overline{\mathbf{x}}_{j}^{\mathrm{T}}+\Sigma_{\mathbf{x}_{j}}\right]$ in (32) and $\mathbf{W}^{\mathrm{T}} \mathbf{W}$ is replaced with $\sum_{i=1}^{d}\left[{\overline{\mathbf{w}_{i}} \mathbf{w}_{i}^{\mathrm{T}}}^{\mathrm{T}}+\Sigma_{\mathbf{w}_{i}}\right]$ in (33). Finally, one can rotate the posterior covariance matrices using (36) and

$$
\Sigma_{\mathbf{w}_{i}, \mathrm{pca}}=\mathbf{V}^{\mathrm{T}} \mathbf{D}_{x}^{1 / 2} \mathbf{U}^{\mathrm{T}} \Sigma_{\mathbf{w}_{i}} \mathbf{U} \mathbf{D}_{x}^{1 / 2} \mathbf{V} .
$$

In the experiments, we use the transformations after each iteration.

## 5. Reconstruction of Missing Values and Selection of Model Rank

PCA is a popular approach to the problem of missing value reconstruction. The estimated principal components capture the correlations between different variables, which allows for reconstruction of missing values from the observed ones. The models discussed in this article compute reconstructions using (5), in which $\mathbf{w}_{i}, m_{i}, \mathbf{x}_{j}$ are replaced with the respective posterior means $\overline{\mathbf{w}}_{i}, \bar{m}_{i}, \overline{\mathbf{x}}_{j}$ when applicable. Additionally, one can estimate the uncertainty of the predictions for PPCA and VBPCA by computing its variance

$$
\tilde{y}_{i j}=\widetilde{m}_{i}+\overline{\mathbf{w}}_{i}^{\mathrm{T}} \Sigma_{\mathbf{x}_{j}} \overline{\mathbf{w}}_{i}+\overline{\mathbf{x}}_{j}^{\mathrm{T}} \Sigma_{\mathbf{w}_{i}} \overline{\mathbf{x}}_{j}+\operatorname{tr}\left(\Sigma_{\mathbf{x}_{j}} \Sigma_{\mathbf{w}_{i}}\right),
$$

where $\widetilde{m}_{i}, \Sigma_{\mathbf{x}_{j}}, \Sigma_{\mathbf{w}_{i}}$ represent the posterior uncertainties of the respective parameters and $\Sigma_{\mathbf{w}_{i}}$ is zero for PPCA. This is a useful feature of the probabilistic methods compared to the least-squares approach.

The quality of reconstruction depends on the number of estimated principal components and therefore selection of the model rank is an important problem. In this section, we briefly outline possible solutions to this problem leaving out the detailed discussion.

The most straightforward way to define the rank of the model is to use cross-validation. The data set is divided into training and validation parts, models with different ranks are fitted to the training set and the the rank is selected based on the performance on the held-out validation set.

Another way is to use probabilistic methods for model selection. For example, Minka (2001) developed a means to detect the dimensionality of the principal subspace for complete data using Laplace's method and Bayesian information criterion (BIC) applied to the PPCA model. He also discussed other alternatives but his approach showed the best performance. However, the extension of Laplace's method to the case of incomplete data is not straightforward. An additional difficulty is the possibility that the posterior distribution over the model parameters has multiple modes corresponding to different principal subspaces. VBPCA can select the optimal number of components automatically by setting some columns of $\mathbf{W}$ to zero. However, the VB cost function may have many local minima which correspond to different model ranks and exploring different local solutions can be a tedious procedure. Hoff (2008) estimated the dimensionality of the principal subspace in the complete data using a sampling procedure.

## 6. Speeding Up Learning For Sparse High-Dimensional Data Sets

Obtaining a reasonably good solution in appropriate time is a very important issue for high-dimensional problems in which learning may take several days, as in the example considered in Section 7.4. The algorithms presented in the previous sections scale differently to problems with large dimensions and large degree of sparsity of observed values.

The computational complexities of different algorithms are summarized in Tables 2 and 3. For example, the alternating optimization algorithm for PPCA requires computation and inversion of matrices $\sum_{j \in O_{i}}\left[\mathbf{x}_{j} \mathbf{x}_{j}^{\mathrm{T}}+\Sigma_{\mathbf{x}_{j}}\right]$ and $\left(v_{y} \mathbf{I}+\sum_{i \in O_{j}} \mathbf{w}_{i} \mathbf{w}_{i}^{\mathrm{T}}\right)$, which are generally unique for each row of $\mathbf{W}$ and each column of $\mathbf{X}$, respectively. The corresponding computational complexity is $O\left(N c^{2}+n c^{3}\right)$ per iteration, ${ }^{2}$ which is quite a bit heavier than $O(n d c)$ per iteration (Roweis, 1998) required for complete data.

The computational complexity $O(N c+n c)$ of the gradient method is small compared to previously reported results. For instance, Srebro and Jaakkola (2003) reported the complexity $O\left(n d c^{2}\right)$ for a more general model which was a huge improvement over the complexity $O\left(n^{3} c^{3}\right)$ of the algorithms proposed even earlier. Using the numbers from the Netflix problem considered in Section $7.4\left(N \approx 10^{8}, n \approx 480189, d=17770, c=30\right)$, we get rough estimates of the differences in practice: Imputation algorithm $n d^{2} \approx 10^{14}$, (Srebro and Jaakkola, 2003) $n d c^{2} \approx 10^{12}$, alternating $\mathbf{W}-\mathbf{X} N c^{2}+n c^{3} \approx 10^{11}$, and gradient $N c+n c \approx 10^{9.5}$.

Once the model has been learned and fixed, doing inference on a new sample is fast. Using the alternating algorithm, it requires only one E-step with $O\left(N_{0} c^{2}+c^{3}\right)$, where $N_{0} \leq d$ is the number of observed values in the new sample, while imputation and gradient algorithms require a number of iterations with $O(d c)$ and $O\left(N_{0} c\right)$ each.

Efficient memory usage is an important issue for sparse data sets. It is preferable that the amount of used memory scales linearly with the number of observed values regardless of the data matrix dimensionalities. In Table 2, we provide the number of parameters estimated by different models, which gives the lower bound for the amount of the required memory. Only the imputation algorithm
2. The computational complexity of the alternating $\mathbf{W}-\mathbf{X}$ scheme can be reduced if the matrices requiring inversion are computed once for all rows (columns) of $\mathbf{Y}$ that have observed values at exactly the same columns (rows respectively).

|  | Section | Number of <br> parameters | Imputation <br> with SVD | Alternating <br> W-X | Gradient <br> descent |
| :---: | :---: | :---: | :---: | :---: | :---: |
| LS | 2 | $(d+n) c$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |
| MAPPCA | 6.4 | $(d+n) c$ |  | $\checkmark$ | $\checkmark$ |
| PPCA | 3.1 | $(d+n) c+n c^{2}$ |  | $\checkmark$ |  |
| VBPCA | 3.3 | $(d+n)\left(c+c^{2}\right)$ |  | $\checkmark$ |  |
| PPCAd | 6.3 | $(d+2 n) c$ |  |  | $\checkmark$ |
| VBPCAd | 6.3 | $(2 d+2 n) c$ |  |  | $\checkmark$ |

Table 2: Memory requirements and applicability of optimization methods to different models. We mark only the optimization schemes which can easily be implemented using the formulas provided in this paper. Some of the methods mentioned here will be presented in the upcoming sections.

| Imputation <br> with SVD | Alternating <br> $\mathbf{W}-\mathbf{X}$ | Gradient <br> descent |
| :---: | :---: | :---: |
| $O\left(n d^{2}\right)$ | $O\left(N c^{2}+n c^{3}\right)$ | $O(N c+n c)$ |

Table 3: Summary of computational complexities (per iteration) of different optimization methods, assuming naïve computation of products and inverses of matrices and ignoring the computation of SVD in the imputation algorithm.
requires memory for the complete data matrix, and even that requirement can be diminished at a cost of greater time complexity.

### 6.1 Gradient-Based Learning

The gradient-based optimization scheme, extending the Oja learning algorithm to regularized PCA models, can be very efficient for large-scale problems and sparse data sets. The computational complexity of one iteration scales very well with dimensionalities, which can lead to faster learning in practice. The gradient-based approach can also be advantageous compared to the alternating optimization scheme as the latter discards the joint effect of parameters $\mathbf{W}$ and $\mathbf{X}$ on the changes in the cost function. This results in slow convergence without proper speed-up procedures (e.g., Honkela et al., 2003).

We also propose to use a speed-up to the gradient descent algorithm. In Newton's method for optimization, the gradient is multiplied by the inverse of the Hessian matrix. Newton's method is known to be fast-converging, but using the full Hessian is computationally costly in high-dimensional problems ( $d \gg 1$ ). We propose a simplified approach which uses only the diagonal part of the Hessian matrix and includes a control parameter $\alpha$ that allows the learning algorithm to vary from the standard gradient descent $(\alpha=0)$ to the diagonal Newton's method $(\alpha=1)$. The final learning rules then take the form

$$
\theta_{i} \leftarrow \theta_{i}-\gamma\left(\frac{\partial^{2} C}{\partial \theta_{i}^{2}}\right)^{-\alpha} \frac{\partial C}{\partial \theta_{i}} .
$$

For example, the cost function (4) has the second-order derivatives

$$
\frac{\partial^{2} C}{\partial w_{i k}^{2}}=\sum_{j \in O_{i}} x_{k j}^{2}, \quad \frac{\partial^{2} C}{\partial x_{k j}^{2}}=\sum_{i \in O_{j}} w_{i k}^{2},
$$

which can be computed efficiently for sparse data. We showed the efficiency of the proposed optimization scheme in the application to large-scale PCA problems in our conference papers (Raiko et al., 2008, 2007a).

### 6.2 Online Learning

In cases where the number of data samples is very large, $n \gg d$, it is wasteful to go through the whole data set before updating $\mathbf{W}$ (and $\mathbf{m}$ ). Using online learning ( Oja , 1983), one sample $j$ is processed at a time as follows: The principal components $\mathbf{x}_{j}$ are inferred using (8) (or the E-step in general), and $\mathbf{W}$ is updated using the gradient approach. When sweeping through the data, updates concerning the latter samples benefit from using the $\mathbf{W}$ that has already been updated using the earlier ones. This reduces the required number of iterations. Downsides of this approach include that determining the step size, or using speed-ups such as the one in Section 6.1 or the conjugate gradient, becomes difficult. Also, for enabling parallel implementation, one should process at least a small batch of samples at a time before updating $\mathbf{W}$ (e.g., Salakhutdinov et al., 2007). Online learning was not considered in this paper.

### 6.3 Factorial Variational Approximations

The overfitting problem can be more severe for high-dimensional sparse data sets. The situations discussed in Sections 2.4 and 3.2 are more probable when data points are sparsely distributed in high dimensions. Thus, methods penalizing large values of model parameters and taking into account their posterior uncertainty are especially relevant here.

The PPCA and VBPCA approaches, which take into account the posterior uncertainty at least in $\mathbf{X}$, are hardly applicable to large-scale problems. First, the computational burden of one iteration of the alternating $\mathbf{W}-\mathbf{X}$ scheme is very large (see Table 3) and application of the gradient-based optimization is cumbersome because of the need to compute $c \times c$ covariance matrices. Second, the required memory is at least $n c^{2}$ elements for storing posterior correlations only in $\mathbf{X}$. This becomes infeasible for many large-scale problems even for a decent number of principal components.

A possible solution is to take into account only some posterior correlations and to use variational techniques for learning. For example, the posterior approximation $q(\theta)$ in VBPCA can be made fully factorial leading to

$$
\begin{equation*}
q(\theta)=\prod_{i=1}^{d} q\left(m_{i}\right) \prod_{i=1}^{d} \prod_{k=1}^{c} q\left(w_{i k}\right) \prod_{c=1}^{k} \prod_{j=1}^{n} q\left(x_{k j}\right) . \tag{38}
\end{equation*}
$$

instead of (23). Such posterior approximation was used, for example, by Raiko et al. (2007b) for PCA in the presence of missing values. The implementation proposed there was based on the imputation algorithm and thus not easily scalable to high-dimensional sparse problems.

The fully factorial approximation (38) reduces significantly the number of variational parameters. They now include the mean parameters $\bar{m}_{i}, \bar{w}_{i k}$ and $\bar{x}_{k j}$ and the variance parameters $\widetilde{m}_{i}, \widetilde{w}_{i k}$, $\widetilde{x}_{k j}$ in addition to hyperparameters $v_{y}, v_{w, k}, v_{m}$, which can be point-estimated. The corresponding
cost function is given in Appendix D. It can be minimized in different ways, for example, by using the gradient-based optimization scheme explained in Section 6.1. The complete algorithm works by alternating four update steps: $\left\{\widetilde{w}_{i k}, \forall i, k\right\},\left\{\widetilde{x}_{k j}, \forall k, j\right\},\left\{\bar{w}_{i k}, \bar{x}_{k j}, \forall i, k, j\right\}$, and $\left\{v_{y}, v_{w, k}, v_{m}, \forall k\right\}$. The required derivatives are reported in Appendix D. We will further refer to this algorithm as VBPCAd.

The idea of fully factorial approximation can also be used for reducing the complexity of PPCA. The posterior of the hidden states $p\left(\mathbf{X} \mid \mathbf{Y}, \mathbf{W}, \mathbf{m}, v_{y}\right)$ can be approximated to be fully factorial on the E-step. The approximating pdf (38) can be fitted to the true posterior by minimizing the cost function

$$
C_{\mathrm{PPCA}}=\int q(\mathbf{X}) \log \frac{q(\mathbf{X})}{p\left(\mathbf{Y}, \mathbf{X} \mid \mathbf{W}, \mathbf{m}, v_{y}\right)} d \mathbf{X},
$$

which is motivated by the variational view of the EM algorithm. The resulting update rules resemble the ones of VBPCAd with the exceptions outlined in Fig. 5. We refer to this approach as PPCAd.

Note that the approximation used in PPCAd does not restrict the generality of the PPCA model when it is applied to complete data. The variational approximation introduces a bias in favor of solutions for which the form of the posterior approximation agrees with the form of the true posterior (e.g., Ilin and Valpola, 2005). The posterior covariance matrix $\Sigma_{\mathbf{x}_{j}}$, which is the same for each $j$ in the fully observed case, is diagonal if and only if the columns of $\mathbf{W}$ are mutually orthogonal. Therefore, restricting the posterior covariance to be diagonal guarantees that (26) holds at convergence. Requirements (24)-(25) are fulfilled because of the assumed prior model for $\mathbf{X}$. Thus, PPCAd for complete data should converge to the PCA basis. PPCAd applied to incomplete data is biased in favor of solutions in which the true posterior covariance matrices are closer to being diagonal. Note also that the idea of speeding up learning by using transformations as described in Section 4.3 for VBPCA, cannot be used for VBPCAd, because the transformations would break up the diagonality of the posterior covariance matrices.

There is an interesting connection between the speed-up proposed in Section 6.1 and the fully diagonal posterior approximation. The second order derivatives needed for the speed-up coincide with the inverse of the posterior variance of each parameter, and thus their computation is practically free.

### 6.4 Ignoring Posterior Uncertainty (MAPPCA)

A further simplification is to use maximum a posteriori (MAP) estimation for the model parameters. The minimized cost function is then minus log-posterior of the parameters. Assuming uniform prior for hyperparameters $v_{y}, v_{w, k}, v_{m}$, it is proportional to

$$
\begin{align*}
C_{\mathrm{MAP}}=\frac{1}{v_{y}} \sum_{i j \in O}\left(y_{i j}-\mathbf{w}_{i}^{\mathrm{T}} \mathbf{x}_{j}-m_{i}\right)^{2}+ & N \log 2 \pi v_{y}+\frac{1}{v_{m}} \sum_{i=1}^{d} m_{i}^{2}+d \log 2 \pi v_{m} \\
& +\sum_{k=1}^{c}\left[\frac{1}{v_{w, k}} \sum_{i=1}^{d} w_{i k}^{2}+d \log 2 \pi v_{w, k}+\sum_{j=1}^{n} x_{k j}^{2}+n \log 2 \pi\right] . \tag{39}
\end{align*}
$$

and it should be minimized w.r.t. the unknown quantities $\mathbf{W}, \mathbf{X}, v_{y}, v_{w, k}, v_{m}$. We refer to this approach as MAPPCA.

Minimization of (39) can be done by any suitable optimization procedure. Derivatives required for gradient-based optimization are given in Appendix E. However, some difficulties should be
avoided. First, using an improper prior for hyperparameters leads to a situation where the posterior pdf is infinitely large when $v_{w, k} \rightarrow 0, \sum_{i=1}^{d} w_{i k}^{2} \rightarrow 0$. This can be overcome by penalizing too small values of the variance hyperparameters. For example, a hyperparameter $v_{z}$, which is a variance parameter of the prior for a set of zero-mean variables $\left\{z_{i}, i=1, \ldots, M\right\}$, can be updated as $v_{z}=$ $\frac{2 \beta+\sum_{i=1}^{M}\left\langle z_{i}^{2}\right\rangle}{2 \alpha+M}$, where $\left\langle z_{i}^{2}\right\rangle$ denotes the expectation of $z_{i}^{2}$ over the posterior, and $\alpha$ and $\beta$ are some small values (we used $\alpha=10^{-3}$ and $\beta=10^{-3}$ ). This update rule corresponds to broad priors restricting $v_{z}$ to be positive and it guarantees that $v_{z} \geq \frac{\beta}{\alpha+M / 2}$. The maximum likelihood (ML) estimate $v_{z}=\frac{1}{M} \sum_{i=1}^{M}\left\langle z_{i}^{2}\right\rangle$ is achieved when $\alpha \rightarrow 0, \beta \rightarrow 0 .{ }^{3}$

Another issue is the non-identifiability of scaling between $\mathbf{W}$ and $\mathbf{X}$. A practical way to avoid this problem is to fix the scale by normalizing the rows of $\mathbf{X}$ to zero mean and unit variance (after each update) and compensating by scaling the columns of $\mathbf{W}$ accordingly. This guarantees the fulfilment of (24)-(25), which are the conditions justified by the prior model.

### 6.5 Summary of the PCA Variants

The differences between the six described variants of the PCA models are summarized in Table 4 and Fig. 5. The approaches differ on the posterior models and on the use of prior for $\mathbf{W}$ and $\mathbf{m}$. All the variants can be derived as special cases of VBPCA by making simplifying assumptions, which is shown in Fig. 5.

## 7. Experiments

We study the properties of the discussed algorithms first using artificial data and then presenting a case study using the Netflix competition data (Netflix, 2007). In the experiments with artificial data, we generated data matrices $\mathbf{Y}$ according to model (12)-(14) with fixed dimensionalities $d, c$ and $n$. The mixing matrix $\mathbf{W}$ was generated by taking a random orthogonal matrix and scaling its columns by $1,2, \ldots, c$. The bias term $\mathbf{m}$ was selected randomly from a Gaussian distribution with variance 10 . For one of the data sets ( $10-5-\mathrm{mlp}$ ), we nonlinearly transformed the resulting data. The observed values were selected randomly in $\mathbf{Y}$ and the rest of the data matrix was used to compute the test reconstruction error.

We used four types of data sets in the experiments:

- Data Set $10-5$ - $g$ : The data set is for the case where $d \ll n$ and Gaussian components. We generated $n=1000$ samples from a 10 -dimensional Gaussian distribution with $c=5$ principal directions. The standard deviation of the noise in (14) was $\sqrt{v_{y}}=0.5$.
- Data Set $10-5-m l p$ : The data lie within a nonlinear manifold and $d \ll n$. The data were generated by a random nonlinear transformation of five-dimensional Gaussian data to the 10 -dimensional space and adding relatively small Gaussian noise. The transformation was performed by randomly initialized multi-layer perceptron networks. The number of samples was $n=1000$.
- Data Set 10-4-mog: The data set is a linear combination (12) of non-Gaussian $\mathbf{x}_{j}$ and $d \ll n$. We used $c=4$ independently distributed non-Gaussian components (elements of $\mathbf{x}_{j}$ ). The

3. In the formulas reported in the article, we always present the ML update rules, although more robust update rules were used in practice.

|  | prior for $\mathbf{W}, \mathbf{m}$ | posterior for $\mathbf{W}, \mathbf{m}$ | posterior for $\mathbf{X}$ |
| :--- | :---: | :---: | :---: |
| LS | uniform | point | point |
| PPCA | uniform | point | full |
| PPCAd | uniform | point | diagonal |
| VBPCA | yes | full | full |
| VBPCAd | yes | diagonal | diagonal |
| MAPPCA | yes | point | point |

Table 4: Variants of probabilistic approaches to PCA.


Figure 5: Relationships between the PCA variants are easily described by deriving each of them as a special case of VBPCA. The corresponding update rules can also be derived from the VBPCA update rules using the given restrictions.


Figure 6: The distributions of the four components used to generate the 10-4-mog data set.
\(\left.\begin{array}{c|ccccl}Data Set \& d \& n \& c \& Gaussianity \& Square root of variances <br>

\hline 10-5-\mathrm{g} \& 10 \& 1000 \& 5 \& yes \& {\left[$$
\begin{array}{lll}5 & 4 & 3\end{array}
$$ 10.5 ··· 0.5\right.}\end{array}\right]\)| no |
| :---: |

Table 5: Characteristics of four types of data sets used in artificial experiments.
distributions of the four components are presented in Fig. 6. The standard deviation of the noise in (14) was $\sqrt{v_{y}}=0.5$. The number of samples was $n=1000$.

- Data Set 100-10-g: The data set is for the case where $d \approx n$ and small noise. The data were generated from a 100 -dimensional Gaussian distribution with $c=10$ principal directions. The standard deviation of the noise in (14) was $\sqrt{v_{y}}=0.5$. The number of samples was $n=100$.

The ratio of missing values was varied for all data sets. We used data sets with relatively small dimensionalities in order to test the performance of many alternative algorithms in a reasonable time. The summary of the data sets is shown in Table 5.

Different PCA algorithms considered in this article were used to estimate the same number $c$ of principal components that was used for generating the data (see Table 5). The algorithms were initialized randomly except for the imputation algorithm which has a standard procedure of infilling missing values with the row-wise means of $\mathbf{Y}$.

The performance of the algorithms was assessed by the speed of their convergence and the accuracy in the task of missing value reconstruction. We computed three quantities: the average training root mean square (RMS) reconstruction error $\sqrt{\frac{1}{N} \sum_{i j \in O}\left(y_{i j}-\hat{y}_{i j}\right)^{2}}$, the average test RMS reconstruction error and the average time to convergence. Averaging was done across 30 different data sets of the same type and the same ratio of missing values. The algorithms were run until the deviation of the training error was less than 0.0001 during 100 iterations. This stopping criterion was used to estimate the time to convergence.

### 7.1 Performance of Least Squares Approaches and MAPPCA When $d \ll n$

In this section, we present the performance of three techniques:

- LS: the least squares approach which optimizes the cost function (4) using the gradient-based procedure with the speed-up explained in Section 6.1
- LS (es): the same algorithm which also used early stopping: $20 \%$ of data was reserved as validation set and learning was stopped when the reconstruction error for the validation set started to increase
- LS (imp): the imputation algorithm presented in Section 2.5
- MAPPCA: method presented in Section 6.4. We used the gradient-based optimization procedure with the speed-up explained in Section 6.1.

The estimated performance of the four algorithms is presented in Fig. 7. There, we also show the results of VBPCA which turned out to be the most accurate method. We summarize the observed results in the following.

The gradient-based LS approach can be used for data sets with a relatively small amount of missing data. For such data sets, it provides reasonable accuracy, converges fast and scales well to high-dimensional data. However, its performance deteriorates fast with the increase of the ratio of missing values. The algorithm can get stuck at regions where learning proceeds very slowly and overfitting problems become very frequent. When overfitting happens, some parameter values can grow very large to explain perfectly only part of data. This results in very bad generalization and large test reconstruction errors (see discussion in Section 2.1). Thus, this method can badly overfit


Figure 7: Left: Average test RMS errors (colored bars) and training RMS errors (white bars) obtained from 30 realizations of data sets of three types. Right: Average time to convergence, in seconds, estimated in the same way. Note the different scale on the r.h.s. plots for VBPCA.
even in relatively simple problems. We observed similar performance for the alternating optimization (see Section 2.2) of the sum-squared cost function: The alternating optimization procedure provided similar reconstruction errors but it seemed to get stuck in regions with slow convergence more often.

The simplest strategy of early stopping did not seem to improve the performance of the LS algorithms. On the contrary, the obtained accuracy was the worst among the considered approaches.

The imputation algorithm can also be a good choice for data sets with a relatively small amount of missing data. The algorithm converges fast, provides good accuracy and it has less problems with overfitting compared to explicit optimization of the sum-squared error. However, the complexity of the required computations grows fast with the increase of the dimensionalities, and therefore the algorithm is hardly applicable to large-scale data sets.

MAPPCA is a fast algorithm which provides reasonable accuracy for data sets with a relatively small amount of missing data. In many cases, it slightly outperforms the LS approaches in terms of the reconstruction accuracy. It is also less prone to serious overfitting problems compared to explicit optimization of (4). Again, similar performance was observed for MAPPCA optimized with the alternating optimization procedure explained in Section 2.2.

### 7.2 Performance of More Advanced Probabilistic Approaches When $d \ll n$

In this section, we present the performance of the more advanced probabilistic approaches:

- PPCA presented in Section 3.1
- PPCAd which is PPCA with fully factorial variational approximation explained in Section 6.3
- VBPCA presented in Section 3.3
- VBPCAd which is VBPCA with fully factorial variational approximation explained in Section 6.3.

We used the gradient-based optimization procedure with the speed-up explained in Section 6.1 for training PPCAd and VBPCAd.

Fig. 8 presents the estimated performance of the algorithms for three data sets with $d \ll n$. We also show the results of MAPPCA for comparison. We summarize the results in the following.

Approaches using variational approximations, especially VBPCA and VBPCAd, may suffer from the underfitting problem. The algorithms may converge to a sub-optimal solution (corresponding to a local minimum of the VB cost function) in which some potentially useful principal components are not used. The algorithms seem to find such sub-optimal solutions more often for sparse data. Special care has to be taken to avoid such local minima. For example, in the presented experiments we fixed priors for $\mathbf{W}$ to be very broad at the beginning of learning and this helped improve the performance of VBPCA and VBPCAd significantly. See also the experiment showing the effect of the broad priors in Section 7.3.

PPCAd is a very efficient algorithm for data sets with $d \ll n$. Its reconstruction accuracy is comparable to the best results (obtained with VBPCA) but the algorithm converges fast and scales well to high-dimensional problems. The difference of the accuracy of PPCAd compared to the best results of VBPCA becomes more noticeable when the ratio of missing values or the number of estimated principle components increase.


Figure 8: Left: Average test RMS errors (colored bars) and training RMS errors (white bars) obtained from 30 realizations of data sets of three types. Right: Average time to convergence, in seconds, estimated in the same way. Note the different scale on the r.h.s. plots for VBPCA.

VBPCAd provides similar performance to PPCAd for data sets with $d \ll n$ but it is more prone to the underfitting problem mentioned earlier.

VBPCA and PPCA are the most accurate algorithms in these experiments. However, they are also very slow because of the required computations of the posterior covariance matrices. Thus, the methods are computationally very expensive to use in high-dimensional problems but if they are feasible, they are worth trying. The difference in the accuracy of VBPCA and PPCA is not noticeable in these experiments. This is likely due to the fact that the number $n$ of samples was very large and the effect of priors was very small. VBPCA is also more prone to the underfitting problem compared to PPCA.

### 7.3 Performance in Difficult Sparse Problems

In this section, we test the performance of the considered algorithms on a more challenging data set $100-10-\mathrm{g}$. The data set is quite sparse (up to $75 \%$ of values are missing) and the number $n$ of samples is very small compared to the data dimensionality $d$. Sparse data sets with $d \approx n$ can appear in some practical applications. For example, in collaborative filtering the number of users can be comparable to the number of ranked items and the number of missing values can be large. Here we used Gaussian data with $n=d=100, c=10$ principal directions and with relatively small noise.

Fig. 9 presents the performance of the considered algorithms for these data. The LS algorithms and MAPPCA do not provide good performance because of overfitting problems. The probabilistic approaches perform much better and the best results are obtained with VBPCA. The advantage of VBPCA becomes more significant for data sets with many missing values. The experiments also show a larger effect of using priors in VBPCA and VBPCAd (in comparison with PPCA and PPCAd) for data sets with a relatively small number of samples.

One of the reasons for the superior performance of the VB approach is its ability to select automatically the optimal rank of the model by cutting off useless principal components. Model selection is useful even in such artificial tests (when the number of principal components used for data generation is known) because the number of components which are useful for reliable reconstruction of missing data can turn out to be smaller than the model rank used for generating the data. Thus, the reported accuracy of the algorithms alternative to VBPCA and VBPCAd might be improved if model selection was explicitly performed for them (e.g., by cross-validation).

In Section 7.2, we pointed out that the downside of the VB model selection is the existence of multiple sub-optimal solutions corresponding to different local minima of the VB cost function. Solutions in which the effect of some components is set to zero are potentially attractive for the method. In order to avoid sub-optimal solutions, we fixed the priors for $\mathbf{W}$ to be very broad at the beginning of learning for VBPCA and VBPCAd in our experiments. We show the effect of the number of iterations with fixed broad priors on the resulting accuracy in Fig. 10.

A useful feature of the VB approach is that the value of the VB cost function at a local minimum can be used to estimate the accuracy of the corresponding model. This follows from the fact that the VB cost gives the lower bound of the likelihood, as shown in (22). In Fig. 11, we demonstrate the correlations between the accuracy of a trained VBPCA model and the corresponding value of the cost function. In that experiment, we initialized VBPCA in two different ways and applied it to 30 realizations of the $100-10-\mathrm{g}$ data set with $80 \%$ of missing data. The negative result was that different initializations led to different VBPCA solutions for most of the 30 realizations. The positive result


Figure 9: Left: Average test rms errors (colored bars) and training rms errors (white bars) obtained from 30 realizations of data set $100-10-\mathrm{g}$. Right: Average time to convergence, in seconds, estimated in the same way.
was that the solutions with the smaller cost function values generally provided better reconstruction accuracy. Similar results were obtained for VBPCAd (not shown here).

Thus, in order to find a globally good VB solution, one could, in practice, run the learning algorithm many times. The algorithm should be run until full convergence, otherwise the cost function values cannot be used for model comparison. Running VBPCA can be computationally heavy because it is the slowest algorithm among the considered ones. Therefore, VBPCAd is an attractive alternative for large-scale data sets.

### 7.4 Experiments with the Netflix Data

We have tested different PCA approaches in the task of collaborative filtering. The Netflix (2007) problem is a collaborative filtering task that has received a lot of attention recently. It consists of movie ratings given by $n=480189$ customers to $d=17770$ movies. There are $N=100480507$ ratings from 1 to 5 given, and the task is to predict 2817131 other ratings among the same group of customers and movies. 1408395 of the ratings are reserved for validation (or probing). Note that $98.8 \%$ of the values are thus missing.


Figure 10: The underfitting effect for VBPCA when the priors are updated too early: Box plots of test (left) and training (right) RMS errors depending on the number of iterations with fixed broad priors for $\mathbf{W}$ (x-axes) at the beginning of learning. VBPCA was applied to 30 data sets of type $100-10-\mathrm{g}$ with $75 \%$ of missing values.


Figure 11: Test RMS error (y-axis) against cost function values (x-axis) obtained with VBPCA for 30 realizations of the data set $100-10-\mathrm{g}$ with $80 \%$ of missing values. Two initializations were used for each realization: random (the corresponding local minima are shown with triangles) and with 100 iterations of PPCAd (the corresponding local minima are shown with circles). The dotted lines connect two solutions obtained for the same realization.

The PCA approach is one of the most popular techniques considered by the Netflix contestants (e.g., Funk, 2006; Bell et al., 2007; Salakhutdinov et al., 2007; Paterek, 2007; Lim and Teh, 2007). In our recent conference papers (Raiko et al., 2007a, 2008), we tried to estimate $c=15$ principal components from the data using unregularized (LS) PCA, MAPPCA, and VBPCAd approaches. We reproduce the results for the LS approaches in Fig. 12. The results clearly show that overfitting is a serious problem for this sparse data set and the LS approach fails. It is also evident that the alternating optimization algorithm is very slow and it can hardly be used for such high-dimensional data. The imputation algorithm is also infeasible in this problem (Raiko et al., 2008). The proposed remedy is to use probabilistic approaches.

Fig. 13 presents the results obtained with MAPPCA and VBPCAd using a varying number of principal components. MAPPCA starts to overfit after about five hours of computation, that is, the RMS error on the test data (lower curves) starts to degrade even if the RMS error on the training data is still diminishing. This does not happen with VBPCAd as the validation error seems to decrease monotonically. The experiments also confirm that VBPCAd is computationally scalable to very large problems. We also tried to run VBPCA, but the straightforward Matlab implementation turned out to be too slow to produce meaningful results. The experiments were run on a dual cpu AMD Opteron SE 2220 using Matlab and the implementations did not use parallel computing.

Our best RMS error for the probing Netflix data was 0.9055 and it was achieved with VBPCAd with 50 components. The same model re-trained using both training and probing data provided an RMS error 0.8990 on the quiz set. The only pre/post-processing that we used was bounding the reconstructions $\hat{y}_{i j}$ such that $1 \leq \hat{y}_{i j} \leq 5$. This is a good result compared to conceptually similar models developed by other contestants. For example, Bell et al. (2007) reported an RMS error of 0.9135 with basic factor analysis with 40 components. Lim and Teh (2007) studied a model which basically implements VBPCA and they reported an RMS error of 0.9141 for a model with 30 components. Salakhutdinov and Mnih (2008b) reported the probing RMS error 0.9280 for an SVD method which essentially implements the LS approach. Their constrained probabilistic matrix factorization model provided a probing RMS error 0.9016 but it used extra informations and more sophisticated priors. The accuracy of our VBPCAd model is comparable to the accuracy of an MCMC approach applied to a similar generative model (Salakhutdinov and Mnih, 2008a). The RMS error achieved with a model with 60 components trained by MCMC was 0.8989 for the quiz set.

Besides reconstruction of missing values, probabilistic PCA methods provide extra information which can be useful in practice. For example, one can predict the uncertainty of the provided reconstructions, as discussed in Section 5. Fig. 14 shows good correlation between the probing RMS errors and the uncertainty of the reconstructions computed by VBPCAd. It is clear that the predicted uncertainty is underestimated, which is largely because of the approximations used in modeling the posterior distributions.

The estimated uncertainty can provide additional information when PCA is used as a preprocesing step for more sophisticated methods. For example, Bell and Koren (2007) used PCA results for constructing a user-oriented neighborhood model for collaborative filtering. Because each column of matrix $\mathbf{X}$ computed by PCA can be considered as a collection of some features associated with a particular user, the similarity between $n$ users can be estimated based on the similarity between the columns of $\mathbf{X}$. The information about the uncertainty of the feature estimates could be used to build a robust similarity measure: Users who have rated few movies (and whose features


Figure 12: Learning curves for LS algorithms in the Netflix problem. The RMS error (y-axis) is plotted against the processor time in hours. The upper curves show the training error while the lower curves show the test error (for the probing data provided by Netflix). Note that the time scale is logarithmic. The probing error provided by the alternating algorithm is not shown because it is too large compared to the other two curves. The two gradient-based approaches were implemented using parallel computing (this implementation was about twice as fast as the implementation on a single CPU).


Figure 13: Learning curves for regularized methods in the Netflix problem. The RMS error (y-axis) is plotted against the processor time in hours. The upper curves show the training error while the lower curves show the test error (for the probing data provided by Netflix). Note that the time scale is logarithmic.


Figure 14: Uncertainty predicted by VBPCAd for the Netflix probing data. The RMS error for the probing data (y-axis) is plotted against the square root of the reconstruction variance estimated by VBPCAd with 50 components (x-axis). To show the dependency more clearly, the estimates of the probing RMS errors are obtained by averaging over a batch of ratings with similar predicted uncertainties.
are estimated with high uncertainty) should not affect the predictions of their neighbors as much as users who have rated many movies (and whose features are estimated more accurately).

Finally, we note that the accuracy achieved by the best teams in the Netflix competition is significantly better than the numbers reported here. However, that does not mean that PCA (and its implementations discussed in this paper) is of no use in collaborative filtering tasks. On the contrary, PCA or its close relatives were used by the leading teams (e.g., Bell et al., 2007) as an important element of the final blend of models. The key to success in that competition was in combining a vast collection of different models rather than in perfecting a single approach (e.g., The Ensemble, 2009).

## 8. Discussion

We have reviewed the problem of PCA learning in the presence of missing values and discussed various approaches to it. We demonstrated that the simplicity of PCA is lost when introducing missing values. Firstly, the estimation of the bias term and the covariance matrix of the data becomes difficult and thus the solution by eigen-decomposition cannot be used directly. Secondly, the convergence to a unique solution cannot be guaranteed even for the simplest PCA models.

Missing values also make the problem of overfitting more relevant to PCA, so there is a need for some form of regularization. In regularized solutions, reconstructions of data vectors are no longer

| conditions | small-scale problems | large-scale problems |
| :---: | :---: | :---: |
| few missing values | imputation | gradient-based LS |
| $d \ll n$ | PPCA | PPCAd |
| many missing values or $d \approx n$ | VBPCA | VBPCAd |

Table 6: Recommendations on the use of different PCA algorithms.
the projections of the data to the principal subspace. Regularization can be done elegantly using probabilistic models.

Inference and learning in the relevant probabilistic models are generally more complex compared to the complete data case. For example, the posterior covariance of principal components $\Sigma_{\mathbf{x}_{j}}$ is no longer same for each sample. Thus, a diagonal posterior covariance $\Sigma_{\mathbf{x}_{j}}$ is no longer sufficient for finding the optimal PCA solution. Regularized models typically have more local minima of the optimized cost function. In particular, there can be local minima corresponding to zero values of hyperparameters.

Another type of treatment for random variables is to use MCMC methods. Their benefits compared to VBPCA includes not having to assume posterior independence like in (23) and being able to vary the number of components (Hoff, 2008), but the downsides include worse interpretability, higher computational complexity, and the need to store samples in case the model is to be applied to new incoming data. Salakhutdinov and Mnih (2008a) used Gibbs sampling initialized with the MAPPCA method for the Netflix problem. The resulting accuracy was similar to ours with similar number of components ( 0.8989 with 60 components against our 0.8990 for VBPCAd with 50 components).

Large-scale PCA problems require fast converging learning algorithms which allow for efficient implementation. Additionally, sparse data sets require that the amount of computer memory would scale linearly with the number of observed values regardless of the data matrix dimensionalities. In this paper, efficient solutions are proposed based on fully factorial variational approximations and approximate Newton's iteration for the relevant optimization procedure. An alternative learning algorithm based on natural gradient was discussed in our conference paper (Raiko et al., 2008).

The choice of the right PCA algorithm depends on a particular data analysis problem. We gave some recommendations in Sections 7.1-7.3 and we summarize our recommendations in Table 6. The LS approaches can be applied to data sets with relatively few missing data: The imputation algorithm can be a good choice for smaller-scale problems and the gradient-based LS algorithms would be good for large data sets. When the number $n$ of samples is relatively large and the data matrix is relatively densely populated, PPCA is a proper choice (or PPCAd for large-scale problems). VB algorithms can be most efficient when there are very few samples ( $d \approx n$ ) or for very sparse data sets. Again, VBPCAd would be a better alternative for large-scale problems.

A Matlab toolbox which contains implementations of all the considered PCA techniques is available online at http://www.cis.hut.fi/projects/bayes/. Some of the implemented algorithms scale well to large-scale sparse data sets, which is achieved by low level implementation of core computations and by support of parallel computing. For example, the experiments with the Netflix data reported in Fig. 13 were obtained using the provided Matlab code.

## Acknowledgments

The authors would like to thank Juha Karhunen, Erkki Oja, Alexey Kaplan, and Jaakko Luttinen for useful discussions. This work was supported by the Academy of Finland (projects 'Unsupervised machine learning in latent variable models', 'Novel machine learning techniques for studying climate variability and its impacts', 'Auditory approaches to automatic speech recognition' and 'Bayesian Machine Learning Algorithms for Discovering Relations with Latent Variable Models'), and by the IST Program of the European Community, under the PASCAL2 Network of Excellence. This publication only reflects the authors' views. We would like to thank Netflix (2007) for providing the data.

## Appendix A. Notes on the Imputation Algorithm

The imputation algorithm can be seen as the application of the EM algorithm (Dempster et al., 1977) to the model

$$
\mathbf{y}_{j}=\mathbf{W} \mathbf{x}_{j}+\mathbf{m}+\varepsilon_{j}
$$

with isotropic Gaussian noise $p\left(\varepsilon_{j}\right)=\mathcal{N}\left(\varepsilon_{j} ; 0, v_{y} \mathbf{I}\right)$ and model parameters $\mathbf{W}, \mathbf{x}_{j}, \mathbf{m}, v_{y}$. Note that $\mathbf{x}_{j}$ belong to the model parameters and missing values in $\mathbf{y}_{j}$ are treated as hidden variables. This is contrary to the probabilistic PCA model (see Section 3.1) where $\mathbf{x}_{j}$ are treated as hidden variables. Note also that with this interpretation, there is no well defined inference procedure for a new data case outside of the training set.

Following the view of the EM algorithm presented by Neal and Hinton (1999), we can write down the cost function

$$
\begin{align*}
& C_{\text {imp }}\left(\theta, v_{y}, q\left(Y_{\text {mis }}\right)\right)=\int q\left(Y_{\text {mis }}\right) \log \frac{q\left(Y_{\text {mis }}\right)}{p\left(\mathbf{Y} \mid \theta, v_{y}\right)} d Y_{\text {mis }} \\
= & \int q\left(Y_{\text {mis }}\right) \log \frac{q\left(Y_{\text {mis }}\right)}{p\left(Y_{\text {mis }} \mid \theta, v_{y}\right)} d Y_{\text {mis }}-\log p\left(Y_{\text {obs }} \mid \theta, v_{y}\right) \tag{40}
\end{align*}
$$

which should be minimized w.r.t. model parameters $\theta=\left\{\mathbf{W}, \mathbf{m}, \mathbf{x}_{j}, \forall j\right\}, v_{y}$ and the pdf over the missing data $q\left(Y_{\mathrm{imp}}\right)$. We denoted the missing data by $Y_{\mathrm{mis}}=\left\{y_{i j}, i j \notin O\right\}$ and the observed data by $Y_{\mathrm{obs}}=\left\{y_{i j}, i j \in O\right\}$.

Minimization of (40) w.r.t. $q\left(Y_{\text {mis }}\right)$ for fixed $\theta$ and $v_{y}$ yields that $q\left(Y_{\text {mis }}\right)=p\left(Y_{\text {mis }} \mid \theta, v_{y}\right)$ because the first term in (40) is simply the Kullback-Leibler divergence between the two pdfs. It is straightforward that

$$
p\left(Y_{\mathrm{mis}} \mid \theta, v_{y}\right)=\prod_{i j \in O} \mathcal{N}\left(y_{i j} ; \hat{y}_{i j}(\theta), v_{y}\right)
$$

where $\hat{y}_{i j}(\theta)$ are the reconstructed missing values using (5) with the current estimates $\theta$. Thus, infilling missing data with the reconstructions $\hat{y}_{i j}(\theta)$ is the step of the imputation algorithm which corresponds to the E-step of the EM-algorithm.

Collecting the terms in (40) which depend only on $\theta$ and $v_{y}$ gives

$$
\begin{gather*}
-\int q\left(Y_{\text {mis }}\right) \log p\left(Y_{\mathrm{obs}}, Y_{\mathrm{mis}} \mid \theta, v_{y}\right) d Y_{\text {mis }} \\
=-\frac{d n}{2} \log 2 \pi v_{y}-\frac{1}{2 v_{y}} \sum_{i j \in O}\left(y_{i j}-\hat{y}_{i j}(\theta)\right)^{2}-\frac{1}{2 v_{y}} \sum_{i j \notin O}\left[\left(\bar{y}_{i j}-\hat{y}_{i j}(\theta)\right)^{2}+v_{y}\right] \\
=-\frac{d n}{2} \log 2 \pi v_{y}-\frac{1}{2 v_{y}} \sum_{i j}\left(y_{i j}-\hat{y}_{i j}(\theta)\right)^{2}-\frac{d n-N}{2} . \tag{41}
\end{gather*}
$$

It is easy to see that minimization of (41) w.r.t. $\theta$ is equivalent to minimizing the $\mathrm{LS} \operatorname{cost}(2)$ for the infilled data matrix. This is done in the imputation algorithm by SVD of the complete data matrix. Thus, the M-step of the EM-algorithm corresponds to performing SVD in the imputation algorithm. Minimization of (41) w.r.t. $v_{y}$ does not affect the steps of the imputation algorithm.

The imputation algorithm implicitly minimizes the cost function (4), thus it belongs to the class of LS PCA algorithms. It can be shown (Neal and Hinton, 1999) that the minimum of (40) coincides with the minimum of minus log-likelihood

$$
\begin{equation*}
-\log p\left(Y_{\mathrm{obs}} \mid \theta, v_{y}\right)=\frac{N}{2} \log 2 \pi v_{y}+\frac{1}{2 v_{y}} \sum_{i j \in O}\left(y_{i j}-\hat{y}_{i j}(\theta)\right)^{2} . \tag{42}
\end{equation*}
$$

Replacing $v_{y}$ in (42) with its maximum likelihood estimate yields that the function being minimized is

$$
\frac{N}{2}\left[\log \left(\frac{2 \pi}{N} \sum_{i j \in O}\left(y_{i j}-\hat{y}_{i j}(\theta)\right)^{2}\right)+1\right]
$$

from which it follows that the imputation algorithm implicitly optimizes the sum-squared error.

## Appendix B. Conditions Fulfilled at the Convergence of Probabilistic PCA

The variational view of the EM algorithm allows for an interpretation of the PPCA learning algorithm in which the cost function

$$
\begin{equation*}
C=\int q(\mathbf{X}) \log \frac{q(\mathbf{X})}{p(\mathbf{X})} d \mathbf{X}-\int q(\mathbf{X}) \log p\left(\mathbf{Y} \mid \mathbf{W}, \mathbf{X}, \mathbf{m}, v_{y}\right) d \mathbf{X} \tag{43}
\end{equation*}
$$

is minimized w.r.t. to the model parameters $\mathbf{W}, \mathbf{m}, v_{y}$ and the $\operatorname{pdf} q(\mathbf{X})$. The E-step fixes the model parameters and minimizes the cost function w.r.t. the distribution $q(\mathbf{X})$ while the M -step minimizes $C$ w.r.t. $\mathbf{W}, \mathbf{m}$ and $v_{y}$ assuming that $q(\mathbf{X})$ is fixed.

Let us consider simple transformations of $\{\mathbf{W}, \mathbf{m}, q(\mathbf{X})\}$ which do not change the second term in (43). For example, subtraction of a constant $\mathbf{b}$ from the columns of $\mathbf{X}$ can be compensated by changing the bias term $\mathbf{m}$ correspondingly:

$$
\begin{align*}
q(\mathbf{X}) & \leftarrow \prod_{j=1}^{n} \mathcal{N}\left(\mathbf{x}_{j} ; \overline{\mathbf{x}}_{j}-\mu, \Sigma_{\mathbf{x}_{j}}\right),  \tag{44}\\
\mathbf{m} & \leftarrow \mathbf{m}+\mathbf{W} \mu . \tag{45}
\end{align*}
$$

because $\mathbf{y}=\mathbf{W} \mathbf{x}+\mathbf{m}=\mathbf{W}(\mathbf{x}-\mu)+(\mathbf{W} \mu+\mathbf{m})$. Similarly, rotation of the columns of $\mathbf{X}$ can be compensated by changing $\mathbf{W}$ :

$$
\begin{align*}
q(\mathbf{X}) & \leftarrow \prod_{j=1}^{n} \mathcal{N}\left(\mathbf{x}_{j} ; \mathbf{A} \overline{\mathbf{x}}_{j}, \mathbf{A} \Sigma_{\mathbf{x}_{j}} \mathbf{A}^{\mathrm{T}}\right),  \tag{46}\\
\mathbf{W} & \leftarrow \mathbf{W A}^{-1} \tag{47}
\end{align*}
$$

because $\mathbf{y}=\mathbf{W x}+\mathbf{m}=\left(\mathbf{W A}^{-1}\right)(\mathbf{A x})+\mathbf{m}$.
These transformations generally affect the first term in (43) which is the Kullback-Leibler divergence between $q(\mathbf{X})$ and $p(\mathbf{X})$ :

$$
\begin{equation*}
D=\sum_{j=1}^{n} \int q\left(\mathbf{x}_{j}\right) \log \frac{q\left(\mathbf{x}_{j}\right)}{p\left(\mathbf{x}_{j}\right)} d \mathbf{x}_{j}=\frac{1}{2} \sum_{j=1}^{n}\left[\operatorname{tr}\left(\Sigma_{\mathbf{x}_{j}}\right)+\overline{\mathbf{x}}_{j}^{\mathrm{T}} \overline{\mathbf{x}}_{j}-\log \operatorname{det} \Sigma_{\mathbf{x}_{j}}\right] . \tag{48}
\end{equation*}
$$

and at the convergence this term cannot be made smaller by transformations of the form (44)-(45) or (46)-(47).

Transformation (44)-(45) changes (48) to

$$
D=\frac{1}{2} \sum_{j=1}^{n}\left[\operatorname{tr}\left(\Sigma_{\mathbf{x}_{j}}\right)+\left(\overline{\mathbf{x}}_{j}-\mu\right)^{\mathrm{T}}\left(\overline{\mathbf{x}}_{j}-\mu\right)-\log \operatorname{det} \Sigma_{\mathbf{x}_{j}}\right] .
$$

Now taking the derivative w.r.t. $\mu$ and equating it to zero gives the optimal $\mu=\frac{1}{n} \sum_{j=1}^{n} \overline{\mathbf{x}}_{j}$. At the convergence, the optimal $\mu$ should be zero and therefore (34) should hold.

Similarly, transformation (46)-(47) gives

$$
\begin{aligned}
D & =\frac{1}{2} \sum_{j=1}^{n}\left[\operatorname{tr}\left(\mathbf{A} \Sigma_{\mathbf{x}_{j}} \mathbf{A}^{\mathrm{T}}\right)+\overline{\mathbf{x}}_{j}^{\mathrm{T}} \mathbf{A}^{\mathrm{T}} \mathbf{A} \overline{\mathbf{x}}_{j}-\log \operatorname{det}\left(\mathbf{A} \Sigma_{\mathbf{x}_{j}} \mathbf{A}^{\mathrm{T}}\right)\right] \\
& =\frac{n}{2} \operatorname{tr}\left(\mathbf{A} \Sigma_{*} \mathbf{A}^{\mathrm{T}}\right)-\frac{1}{2} \sum_{j=1}^{n} \log \operatorname{det}\left(\mathbf{A} \Sigma_{\mathbf{x}_{j}} \mathbf{A}^{\mathrm{T}}\right),
\end{aligned}
$$

where $\Sigma_{*}=\frac{1}{n} \sum_{j=1}^{n}\left[\overline{\mathbf{x}}_{j} \overline{\mathbf{x}}_{j}^{\mathrm{T}}+\Sigma_{\mathbf{x}_{j}}\right]$. Taking the derivative w.r.t. A gives

$$
\mathbf{A} \Sigma_{*}-\left(\mathbf{A}^{\mathrm{T}}\right)^{-1}=0
$$

and therefore the optimal $\mathbf{A}$ satisfies $\mathbf{A} \Sigma_{*} \mathbf{A}^{\mathrm{T}}=\mathbf{I}$. Transformation (46)-(47) yields that $\Sigma_{*} \leftarrow \mathbf{A} \boldsymbol{\Sigma}_{*} \mathbf{A}^{\mathrm{T}}$ and therefore at the convergence $\Sigma_{*}=\mathbf{I}$, which is condition (35).

## Appendix C. Variational Bayesian PCA (VBPCA)

The following form of the posterior approximation is used:

$$
q(\mathbf{W}, \mathbf{X}, \mathbf{m})=\prod_{i=1}^{d} \mathcal{N}\left(m_{i} ; \bar{m}_{i}, \widetilde{m}_{i}\right) \prod_{i=1}^{d} \mathcal{N}\left(\mathbf{w}_{i} ; \overline{\mathbf{w}}_{i}, \Sigma_{\mathbf{w}_{i}}\right) \prod_{j=1}^{n} \mathcal{N}\left(\mathbf{x}_{j} ; \overline{\mathbf{x}}_{j}, \Sigma_{\mathbf{x}_{j}}\right) .
$$

The update of the principal components:

$$
\begin{aligned}
\Sigma_{\mathbf{x}_{j}} & =v_{y}\left(v_{y} \mathbf{I}+\sum_{i \in O_{j}}\left[\overline{\mathbf{w}}_{i} \overline{\mathbf{w}}_{i}^{\mathrm{T}}+\Sigma_{\mathbf{w}_{i}}\right]\right)^{-1}, \\
\overline{\mathbf{x}}_{j} & =\frac{1}{v_{y}} \Sigma_{\mathbf{x}_{j}} \sum_{i \in O_{j}} \overline{\mathbf{w}}_{i}\left(y_{i j}-\bar{m}_{i}\right), \quad j=1, \ldots, n .
\end{aligned}
$$

The update of the bias term and matrix $\mathbf{W}$ :

$$
\begin{align*}
\bar{m}_{i} & =\frac{v_{m}}{\left|O_{i}\right|\left(v_{m}+v_{y} /\left|O_{i}\right|\right)} \sum_{j \in O_{i}}\left[y_{i j}-\overline{\mathbf{w}}_{i}^{\mathrm{T}} \overline{\mathbf{x}}_{j}\right],  \tag{49}\\
\widetilde{m}_{i} & =\frac{v_{y} v_{m}}{\left|O_{i}\right|\left(v_{m}+v_{y} /\left|O_{i}\right|\right)},  \tag{50}\\
\Sigma_{\mathbf{w}_{i}} & =v_{y}\left(v_{y} \operatorname{diag}\left(v_{w, k}^{-1}\right)+\sum_{j \in O_{i}}\left[\overline{\mathbf{x}}_{\overline{\mathbf{x}}} \overline{\mathbf{x}}_{j}^{\mathrm{T}}+\Sigma_{\mathbf{x}_{j}}\right]\right)^{-1}, \\
\overline{\mathbf{w}}_{i} & =\frac{1}{v_{y}} \Sigma_{\mathbf{w}_{i}} \sum_{j \in O_{i}} \overline{\mathbf{x}}_{j}\left(y_{i j}-\bar{m}_{i}\right), \quad i=1, \ldots, d
\end{align*}
$$

and the variance parameters:

$$
\begin{align*}
v_{y} & =\frac{1}{N} \sum_{i j \in O}\left[\left(y_{i j}-\overline{\mathbf{w}}_{i}^{\mathrm{T}} \overline{\mathbf{x}}_{j}-\bar{m}_{i}\right)^{2}+\widetilde{m}_{i}+\overline{\mathbf{w}}_{i}^{\mathrm{T}} \Sigma_{\mathbf{x}_{j}} \overline{\mathbf{w}}_{i}+\overline{\mathbf{x}}_{j}^{\mathrm{T}} \Sigma_{\mathbf{w}_{i}} \overline{\mathbf{x}}_{j}+\operatorname{tr}\left(\Sigma_{\mathbf{x}_{j}} \Sigma_{\mathbf{w}_{i}}\right)\right], \\
v_{w, k} & =\frac{1}{d} \sum_{i=1}^{d}\left(\bar{w}_{i k}^{2}+\widetilde{w}_{i k}\right), \quad v_{m}=\frac{1}{d} \sum_{i=1}^{d}\left(\bar{m}_{i}^{2}+\widetilde{m}_{i}\right), \tag{51}
\end{align*}
$$

where $\widetilde{w}_{i k}$ is the $k$-th element on the diagonal of $\Sigma_{\mathbf{w}_{i}}$.

## Appendix D. VBPCA with Fully Factorial Approximation (VBPCAd)

VBPCAd uses the fully factorial posterior approximation

$$
q(\mathbf{W}, \mathbf{X}, \mathbf{m})=\prod_{i=1}^{d} \mathcal{N}\left(m_{i} ; \bar{m}_{i}, \widetilde{m}_{i}\right) \prod_{i=1}^{d} \prod_{k=1}^{c} \mathcal{N}\left(w_{i k} ; \bar{w}_{i k}, \widetilde{w}_{i k}\right) \prod_{k=1}^{c} \prod_{j=1}^{n} \mathcal{N}\left(x_{k j} ; \bar{x}_{k j}, \widetilde{x}_{k j}\right) .
$$

The VB cost function is then the sum of the following terms:

$$
C_{\mathrm{vb}}=\sum_{i j \in O} C_{y i j}+\sum_{i=1}^{d} C_{m i}+\sum_{i=1}^{d} \sum_{k=1}^{c} C_{w i k}+\sum_{k=1}^{c} \sum_{j=1}^{n} C_{x k j},
$$

where the individual terms are

$$
\begin{aligned}
C_{y i j} & =\frac{1}{2 v_{y}}\left[\left(y_{i j}-\overline{\mathbf{w}}_{i}^{\mathrm{T}} \overline{\mathbf{x}}_{j}-\bar{m}_{i}\right)^{2}+\widetilde{m}_{i}+\sum_{k=1}^{c}\left(\widetilde{w}_{i k} \bar{x}_{k j}^{2}+\bar{w}_{i k}^{2} \widetilde{x}_{k j}+\widetilde{w}_{i k} \widetilde{x}_{k j}\right)\right]+\frac{1}{2} \log 2 \pi v_{y}, \\
C_{m i} & =\left\langle\log \frac{q\left(m_{i}\right)}{p\left(m_{i}\right)}\right\rangle=\frac{\bar{m}_{i}^{2}+\widetilde{m}_{i}}{2 v_{m}}-\frac{1}{2} \log \frac{\widetilde{m}_{i}}{v_{m}}-\frac{1}{2}, \\
C_{w i k} & =\left\langle\log \frac{q\left(w_{i k}\right)}{p\left(w_{i k}\right)}\right\rangle=\frac{\bar{w}_{i k}^{2}+\widetilde{w}_{i k}}{2 v_{w, k}}-\frac{1}{2} \log \frac{\widetilde{w}_{i k}}{v_{w, k}}-\frac{1}{2}, \\
C_{x k j} & =\left\langle\log \frac{q\left(x_{k j}\right)}{p\left(x_{k j}\right)}\right\rangle=\frac{1}{2}\left(\bar{x}_{k j}^{2}+\widetilde{x}_{k j}\right)-\frac{1}{2} \log \widetilde{x}_{k j}-\frac{1}{2} .
\end{aligned}
$$

Variational parameters $\bar{m}_{i}, \widetilde{m}_{i}, v_{w, k}$ and $v_{m}$ are updated to minimize the cost function using the update rules in (49)-(51). Parameters $\widetilde{w}, \widetilde{x}$ can be updated to minimize the cost function:

$$
\begin{align*}
& \widetilde{w}_{i k}=v_{y}\left[\frac{v_{y}}{v_{w, k}}+\sum_{j \in O_{i}}\left(\bar{x}_{k j}^{2}+\widetilde{x}_{k j}\right)\right]^{-1},  \tag{52}\\
& \widetilde{x}_{k j}=v_{y}\left[v_{y}+\sum_{i \in O_{j}}\left(\bar{w}_{i k}^{2}+\widetilde{w}_{i k}\right)\right]^{-1} \tag{53}
\end{align*}
$$

The update rule for $v_{y}$ is

$$
v_{y}=\frac{1}{N} \sum_{i j \in O}\left[\left(y_{i j}-\overline{\mathbf{w}}_{i}^{\mathrm{T}} \overline{\mathbf{x}}_{j}-\bar{m}_{i}\right)^{2}+\widetilde{m}_{i}+\sum_{k=1}^{c}\left(\widetilde{w}_{i k} \bar{x}_{k j}^{2}+\bar{w}_{i k}^{2} \widetilde{x}_{k j}+\widetilde{w}_{i k} \widetilde{x}_{k j}\right)\right] .
$$

Minimizing the cost function w.r.t. each parameter $\bar{w}_{i j}$ keeping the others fixed would lead to a slow algorithm (e.g., Honkela et al., 2003). Instead, derivatives required for gradient-based optimization are:

$$
\begin{aligned}
& \frac{\partial C_{\mathrm{vb}}}{\partial \bar{w}_{i k}}=\frac{\bar{w}_{i k}}{v_{w, k}}+\frac{1}{v_{y}} \sum_{j \in O_{i}}\left[-\left(y_{i j}-\overline{\mathbf{w}}_{i}^{\mathrm{T}} \overline{\mathbf{x}}_{j}-\bar{m}_{i}\right) \bar{x}_{k j}+\bar{w}_{i k} \widetilde{x}_{k j}\right], \\
& \frac{\partial C_{\mathrm{vb}}}{\partial \bar{x}_{k j}}=\bar{x}_{k j}+\frac{1}{v_{y}} \sum_{i \in O_{j}}\left[-\left(y_{i j}-\overline{\mathbf{w}}_{i}^{\mathrm{T}} \overline{\mathrm{x}}_{j}-\bar{m}_{i}\right) \bar{w}_{i k}+\widetilde{w}_{i k} \bar{x}_{k j}\right] .
\end{aligned}
$$

The second-order derivatives which can be used to speed up learning, as explained in Section 6.1, coincide with the inverse of the updated variances given in (52)-(53): $\partial^{2} C_{\mathrm{vb}} / \partial \bar{w}_{i k}^{2}=\widetilde{w}_{i k}^{-1}$ and $\partial^{2} C_{\mathrm{vb}} / \partial \bar{x}_{k j}^{2}=\widetilde{x}_{k j}^{-1}$.

## Appendix E. MAPPCA

The cost function $C_{\text {MAP }}$ is given in (39). The alternating optimization procedure can be implemented as follows:

$$
\begin{aligned}
\mathbf{x}_{j} & =\left(v_{y} \mathbf{I}+\sum_{i \in O_{j}} \mathbf{w}_{i} \mathbf{w}_{i}^{\mathrm{T}}\right)^{-1} \sum_{i \in O_{j}} \mathbf{w}_{i}\left(y_{i j}-m_{i}\right), \quad j=1, \ldots, n, \\
m_{i} & =\frac{v_{m}}{\left|O_{i}\right|\left(v_{m}+v_{y} /\left|O_{i}\right|\right)} \sum_{j \in O_{i}}\left[y_{i j}-\mathbf{w}_{i}^{\mathrm{T}} \mathbf{x}_{j}\right], \\
\mathbf{w}_{i} & =\left(v_{y} \operatorname{diag}\left(v_{w, k}^{-1}\right)+\sum_{j \in O_{i}} \mathbf{x}_{j} \mathbf{x}_{j}^{\mathrm{T}}\right)^{-1} \sum_{j \in O_{i}} \mathbf{x}_{j}\left(y_{i j}-m_{i}\right), \quad i=1, \ldots, d, \\
v_{y} & =\frac{1}{N} \sum_{i j \in O}\left(y_{i j}-\mathbf{w}_{i}^{\mathrm{T}} \mathbf{x}_{j}-m_{i}\right)^{2}, \\
v_{w, k} & =\frac{1}{d} \sum_{i=1}^{d} w_{i k}^{2}, \quad \quad v_{m}=\frac{1}{d} \sum_{i=1}^{d} m_{i}^{2} .
\end{aligned}
$$

Gradient-based learning discussed in Section 6.1 requires the following derivatives:

$$
\begin{aligned}
\frac{\partial C_{\mathrm{MAP}}}{\partial w_{i k}} & =\frac{w_{i k}}{v_{w, k}}+\frac{1}{v_{y}} \sum_{j \in O_{i}}\left[-\left(y_{i j}-\mathbf{w}_{i}^{\mathrm{T}} \mathbf{x}_{j}-m_{i}\right) x_{k j}\right] \\
\frac{\partial C_{\mathrm{MAP}}}{\partial x_{k j}} & =x_{k j}+\frac{1}{v_{y}} \sum_{i \in O_{j}}\left[-\left(y_{i j}-\mathbf{w}_{i}^{\mathrm{T}} \mathbf{x}_{j}-m_{i}\right) w_{i k}\right], \\
\frac{\partial^{2} C_{\mathrm{MAP}}}{\partial w_{i k}^{2}} & =\frac{1}{v_{w, k}}+\frac{1}{v_{y}} \sum_{j \in O_{i}} x_{k j}^{2}, \\
\frac{\partial^{2} C_{\mathrm{MAP}}}{\partial x_{k j}^{2}} & =1+\frac{1}{v_{y}} \sum_{i \in O_{j}} w_{i k}^{2} .
\end{aligned}
$$

## References

T. W. Anderson. Maximum likelihood estimates for a multivariate normal distribution when some observations are missing. Journal of the American Statistical Association, 52:200-203, 1957.
R. Bell and Y. Koren. Scalable collaborative filtering with jointly derived neighborhood interpolation weights. In Proceedings of the IEEE International Conference on Data Mining (ICDM 2007), 2007.
R. Bell, Y. Koren, and C. Volinsky. The BellKor solution to the Netflix Prize. Available at http: //www.netflixprize.com/, 2007.
C. M. Bishop. Variational principal components. In Proceedings of the 9th International Conference on Artificial Neural Networks (ICANN99), pages 509-514, 1999.
C. M. Bishop. Pattern Recognition and Machince Learning. Springer, Cambridge, 2006.
W. J. Boscardin and X. Zhang. Modeling the covariance and correlation matrix of repeated measures. In A. Gelman and X.-L. Meng, editors, Applied Bayesian Modeling and Causal Inference from an Incomplete-Data Perspective. John Wiley \& Sons, New York, 2004.
A. Christoffersson. The One Component Model with Incomplete Data. PhD thesis, Uppsala University, 1970.
A. Cichocki and S. Amari. Adaptive Blind Signal and Image Processing - Learning Algorithms and Applications. Wiley, 2002.
C. L. de Ligny, G. H. E. Nieuwdorp, W. K. Brederode, W. E. Hammers, and J. C. van Houwelingen. An application of factor analysis with missing data. Technometrics, 23(1):91-95, 1981.
R. E. Dear. A principal components missing data method for multiple regression models. Technical Report SP-86, Santa Monica: Systems Development Corporation, 1959.
A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. Journal of the Royal Statistical Society, Series B (Methodological), 39(1):1-38, 1977.
K. Diamantaras and S. Kung. Principal Component Neural Networks - Theory and Application. Wiley, 1996.
S. Funk. Netflix update: Try this at home. Available at http://sifter.org/~simon/journal/ 20061211.html, December 2006.
Z. Ghahramani and M. I. Jordan. Learning from incomplete data. Technical report CBCL-108, Massachusetts Institute of Technology, 1994.
B. Grung and R. Manne. Missing values in principal component analysis. Chemometrics and Intelligent Laboratory Systems, 42(1):125-139, 1998.
S. Haykin. Modern Filters. Macmillan, 1989.
G. E. Hinton and D. van Camp. Keeping neural networks simple by minimizing the description length of the weights. In Proceedings of the 6th Annual ACM Conference on Computational Learning Theory, pages 5-13, Santa Cruz, CA, USA, 1993.
P. Hoff. Model averaging and dimension selection for the singular value decomposition. Journal of the American Statistical Association, 102(478):674-685, 2008.
T. Hofmann. Latent semantic models for collaborative filtering. ACM Transactions on Information Systems, 22(1):89-115, 2004.
A. Honkela, H. Valpola, and J. Karhunen. Accelerating cyclic update algorithms for parameter estimation by pattern searches. Neural Processing Letters, 17(2):191-203, 2003.
A. Ilin and H. Valpola. On the effect of the form of the posterior approximation in variational learning of ICA models. Neural Processing Letters, 22(2):183-204, 2005.
I. T. Jolliffe. Principal Component Analysis. Springer-Verlag, 2nd edition, 2002.
A. Kaplan, Y. Kushnir, M. Cane, and M. Blumenthal. Reduced space optimal analysis for historical datasets: 136 years of Atlantic sea surface temperatures. Journal of Geophysical Research, 102: 27835-27860, 1997.
Y. J. Lim and Y. W. Teh. Variational Bayesian approach to movie rating prediction. In Proceedings of of KDD Cup and Workshop 2007, held during the Thirteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Jose, California, 2007.
R. J. A. Little and D. B. Rubin. Statistical Analysis with Missing Data. J. Wiley \& Sons, 1987.
J. Luttinen and A. Ilin. Transformations in variational Bayesian factor analysis to speed up learning. Neurocomputing, 73(7-9):1093-1102, 2010.
T. Minka. Automatic choice of dimensionality for PCA. In V. Tresp T. Leen, T. Dietterich, editor, Advances in Neural Information Processing Systems 13, pages 598-604. MIT Press, Cambridge, MA, USA, 2001.
R. M. Neal and G. E. Hinton. A view of the EM algorithm that justifies incremental, sparse, and other variants. In Michael I. Jordan, editor, Learning in Graphical Models, pages 355-368. The MIT Press, Cambridge, MA, USA, 1999.

Netflix. Netflix prize webpage, 2007. http://www.netflixprize.com/.
S. Oba, M. Sato, I. Takemasa, M. Monden, K. Matsubara, and S. Ishii. A Bayesian missing value estimation method for gene expression profile data. Bioinformatics, 19(16):2088-2096, 2003.
E. Oja. Subspace Methods of Pattern Recognition. Research Studies Press and J. Wiley, 1983.
A. Paterek. Improving regularized singular value decomposition for collaborative filtering. In Proceedings of KDD Cup and Workshop, 2007.
K. Pearson. On lines and planes of closest fit to systems of points in space. Philosophical Magazine, 2(6):559-572, 1901.
T. Raiko and H. Valpola. Missing values in nonlinear factor analysis. In Proceedings of the 8th International Conference on Neural Information Processing (ICONIP'01), pages 822-827, Shanghai, 2001.
T. Raiko, A. Ilin, and J. Karhunen. Principal component analysis for large scale problems with lots of missing values. In Proceedings of the 18th European Conference on Machine Learning (ECML 2007), pages 691-698, Warsaw, Poland, 2007a.
T. Raiko, H. Valpola, M. Harva, and J. Karhunen. Building blocks for variational Bayesian learning of latent variable models. Journal of Machine Learning Research, 8:155-201, 2007b.
T. Raiko, A. Ilin, and J. Karhunen. Principal component analysis for sparse high-dimensional data. In Proceedings of the 14th International Conference on Neural Information Processing (ICONIP 2007), pages 566-575, Kitakyushu, Japan, 2008.
S. Roweis. EM algorithms for PCA and SPCA. In Advances in Neural Information Processing Systems 10, pages 626-632, Cambridge, MA, 1998. MIT Press.
R. Salakhutdinov and A. Mnih. Bayesian probabilistic matrix factorization using Markov chain Monte Carlo. In Proceedings of the 25th International Conference on Machine Learning (ICML2008), Helsinki, Finland, 2008a.
R. Salakhutdinov and A. Mnih. Probabilistic matrix factorization. In Advances in Neural Information Processing Systems 20, 2008b.
R. Salakhutdinov, A. Mnih, and G. Hinton. Restricted Boltzmann machines for collaborative filtering. In Proceedings of the 24th International Conference on Machine Learning (ICML-2007), 2007.
N. Srebro and T. Jaakkola. Weighted low-rank approximations. In Proceedings of the 20th International Conference on Machine Learning (ICML-2003), Washington DC, 2003.

The Ensemble. Website of the runner-up in the Netflix competition, 2009. http://www. the-ensemble.com/.
M. Tipping and C. M. Bishop. Probabilistic principal component analysis. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 61(3):611-622, 1999.
C. S. Wallace. Classification by minimum-message-length inference. In S. G. Aki, F. Fiala, and W. W. Koczkodaj, editors, Advances in Computing and Information - ICCI '90, volume 468 of Lecture Notes in Computer Science, pages 72-81. Springer, Berlin, 1990.
T. Wiberg. Computation of principal components when data are missing. In Johannes Gordesch and Peter Naeve, editors, COMPSTAT 1976: Proceedings in Computational Statistics, 2nd symposium held in Berlin (West), pages 229-236. Wien: Physica-Verlag, 1976.
G. Young. Maximum likelihood estimation and factor analysis. Psychometrika, 6(1):49-53, 1941.

