Mixture of Gaussians for Distance Estimation with Missing Data<br>Emil Eirola, Amaury Lendasse, Vincent Vandewalle, Christophe Biernacki

## - To cite this version:

Emil Eirola, Amaury Lendasse, Vincent Vandewalle, Christophe Biernacki. Mixture of Gaussians for Distance Estimation with Missing Data. Neurocomputing, 2014, 131, pp.32-42. 10.1016/j.neucom.2013.07.050 . hal-00921023v2

HAL Id: hal-00921023
https://inria.hal.science/hal-00921023v2
Submitted on 30 Dec 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Mixture of Gaussians for distance estimation with missing data 

Emil Eirola ${ }^{\mathrm{a}, *}$, Amaury Lendasse ${ }^{\mathrm{a}, \mathrm{b}, \mathrm{c}, \mathrm{d}}$, Vincent Vandewalle ${ }^{\mathrm{e}, \mathrm{f}}$, Christophe<br>Biernackid ${ }^{\text {,f }}$<br>${ }^{a}$ Department of Information and Computer Science, Aalto University, FI-00076 Aalto, Finland<br>${ }^{b}$ IKERBASQUE, Basque Foundation for Science, 48011 Bilbao, Spain<br>${ }^{c}$ Computational Intelligence Group, Computer Science Faculty, University of the Basque<br>Country, Paseo Manuel Lardizabal 1, Donostia/San Sebastián, Spain<br>${ }^{d}$ Laboratoire P. Painlevé, UMR 8524 CNRS Université Lille I, Bât M3, Cité Scientifique, $F-59655$ Villeneuve d'Ascq Cedex, France<br>${ }^{e}$ EA 2694, Université Lille 2, 1 place de Verdun, 59045 Lille Cedex, France<br>${ }^{f}$ MODAL team, INRIA Lille - Nord Europe, 40 avenue Halley, 59650, Villeneuve d'Ascq, France


#### Abstract

Many data sets have missing values in practical application contexts, but the majority of commonly studied machine learning methods cannot be applied directly when there are incomplete samples. However, most such methods only depend on the relative differences between samples instead of their particular values, and thus one useful approach is to directly estimate the pairwise distances between all samples in the data set. This is accomplished by fitting a Gaussian mixture model to the data, and using it to derive estimates for the distances. A variant of the model for high-dimensional data with missing values is also studied. Experimental simulations confirm that the proposed method provides accurate estimates compared to alternative methods for estimating distances. In particular, using the mixture model for estimating distances is on average more accurate than using the same model to impute any missing values and then calculating distances. The experimental evaluation additionally shows that more accurately estimating distances leads to improved prediction performance for classification and regression tasks when used as inputs for a neural network.


Keywords: missing data, distance estimation, mixture model

[^0]
## 1. Introduction

Missing values are a common phenomenon when dealing with real world data sets. The effects of incomplete data cannot be ignored when designing machine learning procedures. Values could be missing for a variety of reasons depending on the source of the data, including measurement error, device malfunction, operator failure, etc. However, many modelling approaches start with the assumption of a data set with a certain number of samples, and a fixed set of measurements for each sample. Such methods cannot be applied directly if some measurements are missing. Simply discarding the samples or variables which have missing components often means throwing out a large part of data that could be useful for the model. It is relevant to look for better ways of dealing with missing values in such scenarios.

In this paper, the particular problem of estimating distances between samples in a data set with missing values is studied. Being able to appropriately estimate distances between samples, or between samples and prototypes, has numerous applications. It directly enables the use of several powerful statistical and machine learning methods which are based only on distances and not the direct values, including nearest neighbours ( $k$-NN), support vector machines (SVM) with Gaussian kernels, or radial basis function (RBF) neural networks [1].

There are several paradigms for dealing with missing data used in conjunction with machine learning methods [2]. Conditional mean imputation, which is optimal in terms of minimising the mean squared error of the imputed values, suffers from leading to biased derived statistics of the data. For instance, estimates of variance or distances are negatively biased. Random draw imputation is more appropriate for generating a representative example of a fully imputed data set, but has too much variability in estimates of any single values, or distances between particular samples, to be accurate. Multiple imputation (drawing several representative imputations of the data, analysing each set separately, and combining the results) [3] can result in unbiased and accurate estimates after a sufficiently high number of draws, but it is not always straightforward to determine the posterior distribution to draw from [4]. In the context of machine learning, repeating the analysis several times is however impractical as training and analysing a sophisticated model tends to be computationally expensive.

The conceptually simplest approach to dealing with incomplete data is to fill in the missing values before commencing any further analysis. Many methods have been suggested for imputation with the intent to appropriately conform to the distribution of the data. These include imputation by nearest neighbours [5], or the improved incomplete-case $k$-NN imputation [6]. An alternative approach is to study the input density indirectly through conditional distributions, by fully conditional specification [7]. However, the uncertainty of the imputed values is often not explicitly modelled in most imputation methods, and hence ignored in the further analysis, potentially leading to biased results.

One option for automatic preprocessing of data for machine learning is feature selection, which has been studied in the context of incomplete data in
various ways $[8,9]$. A possibility for integrating the imputation of missing values with learning a prediction model is presented in the MLEM2 rule induction algorithm [10]. Another suggested alternative is to use nearest neighbours to simultaneously conduct classification and imputation [11].

Finite mixture models have proven to be a versatile and powerful modelling tool in a wide variety of applications. Particularly mixture models of Gaussians have been studied extensively to describe the distributions of data sets. The general approach to estimate the model parameters from data is maximum likelihood (ML) estimation by the EM algorithm [12]. This has been extended to estimating Gaussian mixture models for data sets with missing values [13, 14]. Recently, Gaussian mixture models as applied to missing data problems have been studied extensively $[15,16,17]$.

A way to use mixtures of Gaussians for training neural networks on data with missing values has previously been proposed in [18], involving finding the average gradient of the relevant parameters by integrating over the conditional distribution of missing values. However, the authors only specify widths for the Gaussian components separately for each dimension in their implementation. This simplifies the analysis greatly, effectively ignoring correlations by restricting the covariance matrices to be diagonal. The suggested procedure specifically applies to training the network by back-propagation, and cannot directly be used for other machine learning methods. Another, more limited, approach to directly allow incomplete samples to be used in back-propagation is to flag input neurons corresponding to unknown attributes as protected, temporarily restricting them from being modified [19]. Further suggested approaches to using a mixture of Gaussians to model the input density for machine learning include forming hidden Markov models for speech recognition by integrating over the density [20]. Another analysis accounting for the uncertainty of missing values using a single multivariate Gaussian in clinical trials is to be found in [21].

A variant of Gaussian mixture models for high-dimensional data has been proposed as high-dimensional data clustering (HDDC) [22]. Its use for data with missing values is studied in this paper.

In spite of the large body of work to approach the problem of incomplete data from different angles, very few authors explicitly consider the distances between the samples. The only specific method to be found to directly estimate these distances is the partial distance strategy [23], which is popularly used and efficient but too simple to adequately account for more sophisticated distributions of the data.

In this paper, we propose to apply mixtures of Gaussians to the task of estimating pairwise distance between samples. Using a Gaussian mixture model is appropriate for this problem, as it 1) can be optimised efficiently even in the presence of missing values, 2) allows one to derive estimates of pairwise distances, 3) is flexible enough to cover any distribution of samples, and 4) is sufficiently sophisticated to provide non-linear imputation.

The sequel of this paper is organised as follows. Section 2 reviews the EM algorithm for mixtures of Gaussians, and introduces the extension to missing
data, including how to efficiently conduct the computations and how to extend the method to high-dimensional data. Section 3 presents the estimation of pairwise distances. Section 4 contains comparison experiments on simulations of data with missing values, including results of neural networks for supervised learning tasks using the estimated distances. Section 5 concludes the paper.

## 2. EM for mixture of Gaussians with missing data

### 2.1. The standard EM algorithm

Before studying the case with missing data, we present the conventional EM algorithm for fitting a mixture of Gaussians [1, Section 9.2]. Given data $\mathbf{X}$ consisting of a set of $N$ observations $\left\{\boldsymbol{x}_{i}\right\}_{i=1}^{N}$, we wish to model the data using a mixture of $K$ Gaussian distributions. The log-likelihood is given by

$$
\begin{equation*}
\log \mathcal{L}(\theta)=\log p(\mathbf{X} \mid \theta)=\sum_{i=1}^{N} \log \left(\sum_{k=1}^{K} \pi_{k} \mathcal{N}\left(\boldsymbol{x}_{i} \mid \boldsymbol{\mu}_{k}, \boldsymbol{\Sigma}_{k}\right)\right) \tag{1}
\end{equation*}
$$

where $\mathcal{N}\left(\boldsymbol{x} \mid \boldsymbol{\mu}_{k}, \boldsymbol{\Sigma}_{k}\right)$ is the probability density function of the multivariate normal distribution, and $\theta=\left\{\pi_{k}, \boldsymbol{\mu}_{k}, \boldsymbol{\Sigma}_{k}\right\}_{k=1}^{K}$ is the set of parameters to be determined. The log-likelihood can be maximised by applying the EM-algorithm. Initialisation consists in choosing values for the means $\boldsymbol{\mu}_{k}$, covariances $\boldsymbol{\Sigma}_{k}$, and mixing coefficients $\pi_{k}$ for each component $k\left(0<\pi_{k}<1, \sum_{k=1}^{K} \pi_{k}=1\right)$. The Estep is to evaluate the probabilities $t_{i k}$ that $\boldsymbol{x}_{i}$ is generated by the $k$ th Gaussian using the current parameter values:

$$
\begin{equation*}
t_{i k}=\frac{\pi_{k} \mathcal{N}\left(\boldsymbol{x}_{i} \mid \boldsymbol{\mu}_{k}, \boldsymbol{\Sigma}_{k}\right)}{\sum_{j=1}^{K} \pi_{j} \mathcal{N}\left(\boldsymbol{x}_{i} \mid \boldsymbol{\mu}_{j}, \boldsymbol{\Sigma}_{j}\right)} . \tag{2}
\end{equation*}
$$

In the M-step, the parameters are re-estimated with the updated probabilities:

$$
\begin{equation*}
\boldsymbol{\mu}_{k}=\frac{1}{N_{k}} \sum_{i=1}^{N} t_{i k} \boldsymbol{x}_{i}, \quad \boldsymbol{\Sigma}_{k}=\frac{1}{N_{k}} \sum_{i=1}^{N} t_{i k}\left(\boldsymbol{x}_{i}-\boldsymbol{\mu}_{k}\right)\left(\boldsymbol{x}_{i}-\boldsymbol{\mu}_{k}\right)^{T}, \quad \pi_{k}=\frac{N_{k}}{N}, \tag{3}
\end{equation*}
$$

where $N_{k}=\sum_{i=1}^{N} t_{i k}$. The E and M-steps are alternated repeatedly until convergence is observed in the log-likelihood or parameter values.

### 2.2. Missing data extension of the EM algorithm

An important consideration when dealing with missing data is the missingdata mechanism. We assume that a missing value represents a value which is defined and exists, but for an unspecified reason is unknown. Following the conventions of [2], the assumption here is that data are Missing-at-Random (MAR): $P\left(M \mid x_{\mathrm{obs}}, x_{\mathrm{mis}}\right)=P\left(M \mid x_{\mathrm{obs}}\right)$, i.e., the event $M$ of a measurement being missing is independent from the value it would take ( $x_{\text {mis }}$ ), conditional on the observed data ( $x_{\text {obs }}$ ). The stronger assumption of Missing-Completely-at-Random (MCAR) is not necessary, as MAR is an ignorable missing-data
mechanism in the sense that maximum likelihood estimation still provides a consistent estimator [2].

The standard EM algorithm for fitting Gaussian mixture models has been extended to handle data with missing values $[13,14]$. The input data $\mathbf{X}$ is now a set of observations $\left\{\boldsymbol{x}_{i}\right\}_{i=1}^{N}$ such that for each sample there is an index set $O_{i} \subseteq\{1, \ldots, d\}$ enumerating the observed variables. The indices in the complement set $M_{i}$ correspond to missing values in the data sample $\boldsymbol{x}_{i}$. In the case with missing values, the observed log-likelihood can be written as

$$
\begin{equation*}
\log \mathcal{L}(\theta)=\log p\left(\mathbf{X}^{O} \mid \theta\right)=\sum_{i=1}^{N} \log \left(\sum_{k=1}^{K} \pi_{k} \mathcal{N}\left(\boldsymbol{x}_{i}^{O_{i}} \mid \boldsymbol{\mu}_{k}, \boldsymbol{\Sigma}_{k}\right)\right) \tag{4}
\end{equation*}
$$

where $\mathbf{X}^{O}=\left\{\boldsymbol{x}_{i}^{O} i\right\}_{i=1}^{N}$, and as a shorthand of notation, $\mathcal{N}\left(\boldsymbol{x}_{i}^{O_{i}} \mid \boldsymbol{\mu}_{k}, \boldsymbol{\Sigma}_{k}\right)$ is also used for the marginal multivariate normal distribution probability density of the observed values of $\boldsymbol{x}_{i}$.

In the EM algorithm, in order to account for the missing data, some additional expectations need to be computed in the E-step. These include the conditional expectations of the missing components of a sample with respect to each Gaussian component $k$, and their conditional covariance matrices, i.e., $\tilde{\boldsymbol{\mu}}_{i k}^{M_{i}}=\mathrm{E}\left[\boldsymbol{x}_{i}^{M_{i}} \mid \boldsymbol{x}_{i}^{O_{i}}\right]$, and $\tilde{\boldsymbol{\Sigma}}_{i k}^{M M_{i}}=\operatorname{Var}\left[\boldsymbol{x}_{i}^{M_{i}} \mid \boldsymbol{x}_{i}^{O_{i}}\right]$, where the mean and covariance are calculated under the assumption that $\boldsymbol{x}_{i}$ originates from the $k$ th Gaussian. For convenience, we also define corresponding imputed data vectors $\tilde{\boldsymbol{x}}_{i k}$ and full covariance matrices $\tilde{\boldsymbol{\Sigma}}_{i k}$ which are padded with zeros for the known components. Then the E-step is:

$$
\begin{align*}
t_{i k} & =\frac{\pi_{k} \mathcal{N}\left(\boldsymbol{x}_{i}^{O_{i}} \mid \boldsymbol{\mu}_{k}, \boldsymbol{\Sigma}_{k}\right)}{\sum_{j=1}^{K} \pi_{j} \mathcal{N}\left(\boldsymbol{x}_{i}^{O_{i}} \mid \boldsymbol{\mu}_{j}, \boldsymbol{\Sigma}_{j}\right)},  \tag{5}\\
\tilde{\boldsymbol{\mu}}_{i k}^{M_{i}} & =\boldsymbol{\mu}_{k}^{M_{i}}+\boldsymbol{\Sigma}_{k}^{M O_{i}}\left(\boldsymbol{\Sigma}_{k}^{O O_{i}}\right)^{-1}\left(\boldsymbol{x}_{i}^{O_{i}}-\boldsymbol{\mu}_{k}^{O_{i}}\right), \quad \tilde{\boldsymbol{x}}_{i k}=\binom{\boldsymbol{x}_{i}^{O_{i}}}{\tilde{\boldsymbol{\mu}}_{i k}^{M_{i}}},  \tag{6}\\
\tilde{\boldsymbol{\Sigma}}_{i k}^{M M_{i}} & =\boldsymbol{\Sigma}_{k}^{M M_{i}}-\boldsymbol{\Sigma}_{k}^{M O_{i}}\left(\boldsymbol{\Sigma}_{k}^{O O_{i}}\right)^{-1} \boldsymbol{\Sigma}_{k}^{O M_{i}},
\end{align*} \quad \tilde{\boldsymbol{\Sigma}}_{i k}=\left(\begin{array}{cc}
\mathbf{0}^{O O_{i}} & \mathbf{0}^{O M_{i}}  \tag{7}\\
\mathbf{0}^{M O_{i}} & \tilde{\boldsymbol{\Sigma}}_{i k}^{M M_{i}}
\end{array}\right) . . ~ l
$$

The notation $\boldsymbol{\mu}_{k}^{M_{i}}$ refers to using only the elements from the vector $\boldsymbol{\mu}_{k}$ specified by the index set $M_{i}$, and similarly for $\boldsymbol{x}_{i}^{O_{i}}$, etc. For matrices, $\boldsymbol{\Sigma}_{k}^{M O_{i}}$ refers to elements in the rows specified by $M_{i}$ and columns by $O_{i}$, and so on. The expressions for the parameters in equations (6) and (7) originate from the observation that the conditional distribution of the missing components also follows a multivariate normal distribution, with these parameters [24, Thm. 2.5.1].

The M-step remains functionally equivalent, the only changes being that the component means are estimated from the imputed data vectors and the covariance matrix estimate requires an additional term to include the covariances concerning the imputed values.
$\boldsymbol{\mu}_{k}=\frac{1}{N_{k}} \sum_{i=1}^{N} t_{i k} \tilde{\boldsymbol{x}}_{i k}, \quad \boldsymbol{\Sigma}_{k}=\frac{1}{N_{k}} \sum_{i=1}^{N} t_{i k}\left[\left(\tilde{\boldsymbol{x}}_{i k}-\boldsymbol{\mu}_{k}\right)\left(\tilde{\boldsymbol{x}}_{i k}-\boldsymbol{\mu}_{k}\right)^{T}+\tilde{\boldsymbol{\Sigma}}_{i k}\right], \quad \pi_{k}=\frac{N_{k}}{N}$.

### 2.3. High-dimensional data

As the number of free parameters grows with the square of the data dimension, in high-dimensional cases it is often not possible to fit a conventional Gaussian mixture model, or even a model with a single Gaussian component. A version of Gaussian mixture models, high-dimensional data clustering (HDDC), has been proposed for this scenario [22] where in the M-step, the covariances matrices are replaced by a reduced representation. Applying this idea to the case of missing data is possible by modifying the covariance matrices of each component after calculating them in the M-step.

From the different variants of HDDC, the experiments in this paper allow all the parameters of the reduced representation to be determined freely for each Gaussian (the model $\left[a_{i j} b_{i} Q_{i} d_{i}\right]$ in the notation of [22]). Following [22], the number of significant eigenvalues is selected by the scree test [25], where the dimension is selected when the subsequent eigenvalues have a difference smaller than a specified threshold ( 0.001 of the trace of the covariance matrix).

In brief (see [22] for details and derivations), the reduced representation entails taking the eigenvalue decomposition of the covariance matrix

$$
\boldsymbol{\Sigma}_{k}=\mathbf{Q}_{k} \boldsymbol{\Lambda}_{k} \mathbf{Q}_{k}^{T}
$$

and modifying all the eigenvalues $\lambda_{k j}$ apart from some of the largest ones. The scree test is used to find the number of eigenvalues $d_{k}$ to be kept exactly. The remaining eigenvalues are replaced by their arithmetic mean. In other words, the covariance matrix $\boldsymbol{\Sigma}_{k}$ is replaced by a matrix $\boldsymbol{\Sigma}_{k}^{\prime}=\mathbf{Q}_{k} \boldsymbol{\Lambda}_{k}^{\prime} \mathbf{Q}_{k}^{T}$ where $\boldsymbol{\Lambda}_{k}^{\prime}$ is a diagonal matrix with the elements

$$
\lambda_{k j}^{\prime}= \begin{cases}\lambda_{k j} & j \leq d_{k} \\ b_{k} & j>d_{k}\end{cases}
$$

where

$$
b_{k}=\frac{1}{d-d_{k}} \sum_{l=d_{k}+1}^{d} \lambda_{k l}=\frac{1}{d-d_{k}}\left(\operatorname{trace}\left(\boldsymbol{\Sigma}_{k}\right)-\sum_{l=1}^{d_{k}} \lambda_{k l}\right) .
$$

assuming the eigenvalues $\lambda_{k j}$ are in decreasing order. This representation implies that only the first $d_{k}$ eigenvalues and eigenvectors need to be calculated and stored, efficiently reducing the number of free parameters required to specify each Gaussian component.

### 2.4. Initialisation

In our implementation of the EM algorithm with missing data, the means $\boldsymbol{\mu}_{k}$ are initialised by a random selection from the observed data, preferring samples without missing values if available. When there are insufficient complete samples, some of the Gaussian means are initialised by incomplete samples where the missing values are imputed by the sample mean. The covariances $\boldsymbol{\Sigma}_{k}$ are initialised with the sample covariance of the data (ignoring samples with missing values). Alternatively, the covariances can be initialised as diagonal matrices, using only the sample variance of each variable.

### 2.5. Efficient EM iterations with the sweep operator

The most computationally consuming part of the EM algorithm is to deal with the inverses of the covariance matrices in equations (6) and (7), as the specific matrix to be inverted depends on the pattern of missing values for each data sample. A naive implementation would require inverting a matrix for each sample and Gaussian component every iteration, which is time consuming. A more efficient procedure for the same end result is to use the sweep operator [2]. A symmetric matrix $G$ is said to be swept on row and column $m$ when it results in a matrix $H=\operatorname{SWP}[m] G$ with elements:

$$
\begin{aligned}
h_{m m} & =-1 / g_{m m} & & \\
h_{j m}=h_{m j} & =g_{j m} / g_{m m} & & j \neq m \\
h_{j l} & =g_{j l}-g_{j m} g_{m l} / g_{m m} & & j, l \neq m
\end{aligned}
$$

Sweeping over several rows and columns can be conducted sequentially, denoted as $\operatorname{SWP}\left[m_{1}, m_{2}, \ldots, m_{t}\right]=\operatorname{SWP}\left[m_{1}\right] \operatorname{SWP}\left[m_{2}\right] \ldots \operatorname{SWP}\left[m_{t}\right]$. Furthermore, it can be shown that the sweep operator is commutative, and thus when sweeping over a set of indices it is not necessary to specify the order in which the sweeps are done. Hence it is possible to simplify and consider the process of sweeping over a set of indices, meaning sequentially sweeping over each index in the set, in any order. Now to apply the sweep operator to the present situation in order to find the conditional expectation and covariance distribution of a sample $\boldsymbol{x}_{i}$ with respect to the Gaussian component $k$, form the $(d+1) \times(d+1)$ matrix:

$$
G=\left[\begin{array}{ccc}
\boldsymbol{\Sigma}_{k}^{M M_{i}} & \boldsymbol{\Sigma}_{k}^{M O_{i}} & \boldsymbol{\mu}_{k}^{M_{i}} \\
\boldsymbol{\Sigma}_{k}^{O M_{i}} & \boldsymbol{\Sigma}_{k}^{O O_{i}} & \boldsymbol{\mu}_{k}^{O_{i}}-\boldsymbol{x}_{i}^{O_{i}} \\
\left(\boldsymbol{\mu}_{k}^{M i}\right)^{T} & \left(\boldsymbol{\mu}_{k}^{O_{i}}-\boldsymbol{x}_{i}^{O_{i}}\right)^{T} & 0
\end{array}\right]
$$

Then sweeping over all observed variables $O_{i}$ results in
$\operatorname{SWP}\left[O_{i}\right] G=\left[\begin{array}{ccc}\tilde{\boldsymbol{\Sigma}}_{i k}^{M M_{i}} & \boldsymbol{\Sigma}_{k}^{M O_{i}}\left(\boldsymbol{\Sigma}_{k}^{O O_{i}}\right)^{-1} & \tilde{\boldsymbol{\mu}}_{i k}^{M_{i}} \\ \left(\boldsymbol{\Sigma}_{k}^{O O_{i}}\right)^{-1} \boldsymbol{\Sigma}_{k}^{O M_{i}} & \left(\boldsymbol{\Sigma}_{k}^{O O_{i}}\right)^{-1} & \left(\boldsymbol{\Sigma}_{k}^{O O_{i}}\right)^{-1}\left(\boldsymbol{x}_{i}^{O_{i}}-\boldsymbol{\mu}_{k}^{O_{i}}\right) \\ \left(\tilde{\boldsymbol{\mu}}_{i k}^{M_{i}}\right)^{T} & \left(\boldsymbol{x}_{i}^{O_{i}}-\boldsymbol{\mu}_{k}^{O_{i}}\right)^{T}\left(\boldsymbol{\Sigma}_{k}^{O O_{i}}\right)^{-1} & a\end{array}\right]$
where

$$
\begin{aligned}
\tilde{\boldsymbol{\Sigma}}_{i k}^{M M_{i}} & =\boldsymbol{\Sigma}_{k}^{M M_{i}}-\boldsymbol{\Sigma}_{k}^{M O_{i}}\left(\boldsymbol{\Sigma}_{k}^{O O_{i}}\right)^{-1} \boldsymbol{\Sigma}_{k}^{O M_{i}} \\
\tilde{\boldsymbol{\mu}}_{i k}^{M_{i}} & =\boldsymbol{\mu}_{k}^{M_{i}}-\boldsymbol{\Sigma}_{k}^{M O_{i}}\left(\boldsymbol{\Sigma}_{k}^{O O_{i}}\right)^{-1}\left(\boldsymbol{\mu}_{k}^{O_{i}}-\boldsymbol{x}_{i}^{O_{i}}\right) \\
a & =-\left(\boldsymbol{\mu}_{k}^{O_{i}}-\boldsymbol{x}_{i}^{O_{i}}\right)^{T}\left(\boldsymbol{\Sigma}_{k}^{O O_{i}}\right)^{-1}\left(\boldsymbol{\mu}_{k}^{O_{i}}-\boldsymbol{x}_{i}^{O_{i}}\right)
\end{aligned}
$$

Here the top left part of the matrix is the conditional covariance of the missing values $\left(\tilde{\boldsymbol{\Sigma}}_{i k}^{M M_{i}}\right)$, the vector at the top right is the conditional mean $\left(\tilde{\boldsymbol{\mu}}_{i k}^{M_{i}}\right)$, and the scalar value in the bottom right is necessary for finding the value of the marginal probability density for evaluating the log-likelihood.

However, in most data sets there are more observed data points than missing values, and hence it is more efficient to sweep over the missing values rather
than the observed ones. This can be accomplished by applying the inverse of the sweep operator - the reverse sweep operator - denoted $H=\operatorname{RSWP}[m] G$ and defined by its elements:

$$
\begin{aligned}
h_{m m} & =-1 / g_{m m} & & \\
h_{j m}=h_{m j} & =-g_{j m} / g_{m m} & & j \neq m \\
h_{j l} & =g_{j l}-g_{j m} g_{m l} / g_{m m} & & j, l \neq m
\end{aligned}
$$

Instead of calculating $H=\operatorname{SWP}\left[O_{i}\right] G$ directly, we can find the same result by computing $H=\operatorname{RSWP}\left[M_{i}\right] G^{\prime}$ instead, where

$$
G^{\prime}=\operatorname{SWP}[1, \ldots, d] G=\left[\begin{array}{cc}
\boldsymbol{\Sigma}_{k}^{-1} & \boldsymbol{\Sigma}_{k}^{-1}\left(\boldsymbol{\mu}_{k}-\boldsymbol{x}_{i}\right) \\
\left(\boldsymbol{\mu}_{k}-\boldsymbol{x}_{i}\right)^{T} \boldsymbol{\Sigma}_{k}^{-1} & -\left(\boldsymbol{\mu}_{k}-\boldsymbol{x}_{i}\right)^{T} \boldsymbol{\Sigma}_{k}^{-1}\left(\boldsymbol{\mu}_{k}-\boldsymbol{x}_{i}\right)
\end{array}\right]
$$

where for this calculation, zeros are inserted for the missing values of $\boldsymbol{x}_{i}$. Note that forming $G^{\prime}$ does not require the matrix $\boldsymbol{\Sigma}_{k}$ to be inverted repeatedly for each sample $\boldsymbol{x}_{i}$.

For samples with no missing values, it is only necessary to evaluate ( $\boldsymbol{x}_{i}-$ $\left.\boldsymbol{\mu}_{k}\right)^{T} \boldsymbol{\Sigma}_{k}^{-1}\left(\boldsymbol{x}_{i}-\boldsymbol{\mu}_{k}\right)$ for calculating the value of the probability density. The covariance matrices $\boldsymbol{\Sigma}_{k}$ can be inverted in $O\left(K d^{3}\right)$ operations, and the vector-matrix-vector product requires $O\left(d^{2}\right)$ operations for each sample and component. Sweeping over one index in $G^{\prime}$ requires $O\left(d^{2}\right)$ operations. Every missing value in the data set requires one sweep operation for each Gaussian component. The M-step is computationally less intensive, but estimating the updated covariance matrices still has a complexity of $O\left(N K d^{2}\right)$. Adding it all up, the complexity of one iteration of the EM algorithm is $O\left(K d^{3}+N K d^{2}+K d^{2} L\right)$, where $L=\sum_{i=1}^{N}\left|M_{i}\right|$ is the total number of missing values in the data set.

### 2.6. Model selection

The number of components is selected according to the Akaike information criterion [26] with the small sample (second-order) bias adjustment [27]. Using the corrected version is crucial, as the number of parameters grows relatively large when increasing the number of components. The corrected Akaike information criterion is a function of the log-likelihood:

$$
\begin{equation*}
\mathrm{AIC}_{C}=-2 \log \mathcal{L}(\theta)+2 P+\frac{2 P(P+1)}{N-P-1} \tag{9}
\end{equation*}
$$

where $P$ is the number of free parameters. $P=K d+K-1+\frac{1}{2} K d(d+1)$ in the case of full, separate, covariance matrices for each of the $K$ components. With high-dimensional data sets, the number of parameters quickly tends to become larger than the number of available samples when increasing the number of components, and the criterion would not be valid anymore. This effect can be mitigated by imposing restrictions on the structure of the covariance matrices, but this would also make the model less powerful.

## 3. Distance estimation with missing data

The intended application of the mixture of Gaussians model is to use it for distance estimation. The problem of estimating distances between samples with missing data is non-trivial, since even perfect imputation (by the conditional expectation) results in biased estimates for the distance. Using additional knowledge about the distribution of the data leads to more accurate estimates.

In the following, we focus on calculating the expectation of the squared Euclidean ( $\ell^{2}$ ) distance. Estimating the $\ell^{2}$-norm itself could be feasible, but due to the square-root, the expressions do not simplify and separate as cleanly. Another motivation for directly estimating the squared distance is that many methods for further processing of the distance matrix actually only make use of the squared distances (e.g., RBF and SVM), while others only consider the ranking of the distances (nearest neighbours).

Given two samples $\boldsymbol{x}_{i}, \boldsymbol{x}_{j} \in \mathbb{R}^{d}$ with components $x_{i}^{l}, x_{j}^{l}(1 \leq l \leq d)$, which may contain missing values, denote again by $M_{i}$ and $O_{i}$ the set of indices of the missing and, respectively, observed components for each sample $\boldsymbol{x}_{i}$. Partition the index set into four parts based on the missing components, and the expression for the squared distance $\left\|\boldsymbol{x}_{i}-\boldsymbol{x}_{j}\right\|^{2}$ can be split accordingly:

$$
\begin{aligned}
&\left\|\boldsymbol{x}_{i}-\boldsymbol{x}_{j}\right\|^{2}=\sum_{l=1}^{d}\left(x_{i}^{l}-x_{j}^{l}\right)^{2}=\sum_{l \in O_{i} \cap O_{j}}\left(x_{i}^{l}-x_{j}^{l}\right)^{2}+\sum_{l \in O_{i} \cap M_{j}}\left(x_{i}^{l}-x_{j}^{l}\right)^{2} \\
&+\sum_{l \in M_{i} \cap O_{j}}\left(x_{i}^{l}-x_{j}^{l}\right)^{2}+\sum_{l \in M_{i} \cap M_{j}}\left(x_{i}^{l}-x_{j}^{l}\right)^{2}
\end{aligned}
$$

The missing values can be modelled as random variables, $X_{i}^{l}, l \in M_{i}$. Taking the expectation of the above expression, by the linearity of expectation:

$$
\begin{aligned}
\mathrm{E}\left[\left\|\boldsymbol{x}_{i}-\boldsymbol{x}_{j}\right\|^{2}\right]= & \sum_{l \in O_{i} \cap O_{j}}\left(x_{i}^{l}-x_{j}^{l}\right)^{2}+\sum_{l \in O_{i} \cap M_{j}}\left(\left(x_{i}^{l}-\mathrm{E}\left[X_{j}^{l}\right]\right)^{2}+\operatorname{Var}\left[X_{j}^{l}\right]\right) \\
& +\sum_{l \in M_{i} \cap O_{j}}\left(\left(\mathrm{E}\left[X_{i}^{l}\right]-x_{j}^{l}\right)^{2}+\operatorname{Var}\left[X_{i}^{l}\right]\right) \\
& +\sum_{l \in M_{i} \cap M_{j}}\left(\left(\mathrm{E}\left[X_{i}^{l}\right]-\mathrm{E}\left[X_{j}^{l}\right]\right)^{2}+\operatorname{Var}\left[X_{i}^{l}\right]+\operatorname{Var}\left[X_{j}^{l}\right]\right) .
\end{aligned}
$$

In the final summation, it is necessary to consider $X_{i}^{l}$ and $X_{j}^{l}$ to be uncorrelated, given the known values of $\boldsymbol{x}_{i}$ and $\boldsymbol{x}_{j}$. This assumption is not restrictive, and follows directly from the common approach that samples are independent draws from an unknown multivariate distribution.

It thus suffices to find the expectation and variance of each random variable separately. If the original samples $\boldsymbol{x}_{i}$ are thought to originate as independent draws from a multivariate distribution, the distributions of the random variables $X_{i}^{l}$ can be found as the conditional distribution when conditioning their joint distribution on the observed values. Then finding the expected squared distance
between two samples reduces to finding the (conditional on the observed values) expectation and variance of each missing component separately. Define $\tilde{\boldsymbol{x}}_{i}$ to be an imputed version of $\boldsymbol{x}_{i}$ where each missing value has been replaced by its conditional mean, and define $\tilde{\sigma}_{i}^{l}$ as the corresponding conditional variance:

$$
\tilde{x}_{i}^{l}=\left\{\begin{array}{ll}
\mathrm{E}\left[X_{i}^{l} \mid \boldsymbol{x}_{i}^{O_{i}}\right] & \text { if } l \in M_{i},  \tag{10}\\
x_{i}^{l} & \text { otherwise }
\end{array} \quad \tilde{\sigma}_{i}^{l}= \begin{cases}\operatorname{Var}\left[X_{i}^{l} \mid \boldsymbol{x}_{i}^{O_{i}}\right] & \text { if } l \in M_{i} \\
0 & \text { otherwise }\end{cases}\right.
$$

With these notations, the expectation of the squared distance can conveniently be expressed as:

$$
\begin{equation*}
\mathrm{E}\left[\left\|\boldsymbol{x}_{i}-\boldsymbol{x}_{j}\right\|^{2}\right]=\left\|\tilde{\boldsymbol{x}}_{i}-\tilde{\boldsymbol{x}}_{j}\right\|^{2}+s_{i}+s_{j}, \quad \text { where } \quad s_{i}=\sum_{l \in M_{i}} \tilde{\sigma}_{i}^{l} \tag{11}
\end{equation*}
$$

This form of the expression particularly emphasises how the uncertainty of the missing values is accounted for. The first term - the distance between imputed samples - already provides an estimate of the distance between $\boldsymbol{x}_{i}$ and $\boldsymbol{x}_{j}$, but including the variances of each imputed component is the deciding factor.

The conditional means and covariances can be calculated using the Gaussian mixture model. These are calculated separately for each component in the M-step, and it only remains to determine the overall conditional mean and covariance matrix. These are found weighted by the memberships as follows:

$$
\begin{equation*}
\tilde{\boldsymbol{x}}_{i}=\sum_{k=1}^{K} t_{i k} \tilde{\boldsymbol{x}}_{i k}, \quad \quad \tilde{\boldsymbol{\Sigma}}_{i}=\sum_{k=1}^{K} t_{i k}\left(\tilde{\boldsymbol{\Sigma}}_{i k}+\tilde{\boldsymbol{x}}_{i k} \tilde{\boldsymbol{x}}_{i k}^{T}\right)-\tilde{\boldsymbol{x}}_{i} \tilde{\boldsymbol{x}}_{i}^{T} \tag{12}
\end{equation*}
$$

The expression for the covariance is found by direct calculation of the second moments. In order to estimate pairwise distances, the conditional variances $\tilde{\sigma}_{i}^{l}=$ $\tilde{\boldsymbol{\Sigma}}_{i}^{l l}$ can be extracted from the diagonal of the conditional covariance matrix, or $s_{i}$ calculated directly as the trace of $\tilde{\boldsymbol{\Sigma}}_{i}$.

In summary, the complete procedure for estimating distances consists of the following steps:

1. Fit a Gaussian mixture model of $K$ components by the EM algorithm, for values of $K$ from 1 up to a chosen maximum.
2. Calculate the log-likelihood and $\mathrm{AIC}_{C}$ for each model.
3. Choose the model which minimises the $\mathrm{AIC}_{C}$, and apply it calculate conditional means and variances for each missing value in the data set.
4. Estimate distances between samples using the conditional means and variances.

## 4. Experiments

To study the effectiveness of the proposed approach, some simulated experiments are conducted to compare the algorithm to alternative methods on several data sets. The accuracy of the estimated distances is directly evaluated by different criteria, and also through the performance of neural network classifiers

Table 1: Data sets used for the experiments, with the number of samples $(N)$, number of variables (d), supervised learning task, and source.

| Name | $N$ | $d$ | Task | Source |
| :---: | :---: | :---: | :---: | :---: |
| Computer Hardware | 209 | 66 | Regression | [28] |
| Glass Identification | 214 | 9 | 6-class | [28] |
| Housing | 506 | 9 | Regression | [28] |
| Image Segmentation ${ }^{1}$ | 2310 | 14 | 7-class | [28] |
| Iris | 150 | 4 | 3-class | [28] |
| Pima Indians Diabetes | 768 | 8 | 2-class | [28] |
| Servo | 167 | 4 | Regression | [28] |
| Stocks | 950 | 9 | Regression | [29] |
| Statlog (Vehicle Silhouettes) | 846 | 18 | 4-class | [28] |
| Auto-price | 159 | 15 | Regression | [29] |
| Breast Cancer Wisconsin (Diagnostic) | 569 | 30 | 2-class | [28] |
| Breast Cancer Wisconsin (Prognostic) | 194 | 32 | Regression | [28] |
| Breast Tissue | 106 | 9 | 6 -class | [28] |
| Ecoli | 336 | 7 | 4 -class | [28] |
| Financial Ratios ${ }^{2}$ | 500 | 40 | 2-class | 5 |
| Ionosphere | 351 | 33 | 2-class | [28] |
| Parkinsons ${ }^{3}$ | 195 | 20 | 2-class | [28] |
| Spambase ${ }^{4}$ | 4601 | 57 | 2-class | [28] |
| SPECTF Heart | 267 | 44 | 2-class | [28] |
| Wine | 178 | 13 | 3-class | [28] |
| BCI | 400 | 117 | 2-class | 6 |
| COIL | 1500 | 241 | 6-class | 6 |
| Tecator | 240 | 100 | Regression | 7 |

${ }^{1}$ Discarding the attributes $3,10,14,15$, and 16 , as they are averages of other variables
${ }^{2}$ Discarding variable SF7 which is merely the sum of SF10 and LI8
${ }^{3}$ Discarding Jitter:DDP and Shimmer:DD, being duplicates of MDVP:RAP and Shimmer:APQ3
${ }^{4}$ Preprocessed by taking the logarithm of each value
${ }^{5}$ Available at http://research.ics.aalto.fi/eiml/datasets.shtml
${ }_{7}^{6}$ Available at http://olivier.chapelle.cc/ssl-book/benchmarks.html; see also [30]
${ }^{7}$ Available at http://www.dm.unibo.it/~simoncin/tecator
and regression models built using the estimated distances. Starting with a complete data set, values are removed at random with a fixed probability. As the true distances between samples are known, the methods can then be compared on how well they estimate the distances after values have been removed.

### 4.1. Data

Several different data sets are used for the experiments, and they are listed in Table 1. Some data sets contain variables which are direct linear combinations of other variables, and as this induces unnecessary computational difficulties, some variables are discarded as detailed in the table. The data sets contain binary classification, multiclass classification, and regression tasks. As the problem of pairwise distance estimation is unsupervised, the outputs for the samples are ignored while estimating distances.

To make distances meaningful, the variables in each data set are standardised to zero mean and unit variance before values are removed. This standardisation is conducted only in order to have comparable error rates between repeated
experiments, and the methods used do not depend on the variables being standardised.

### 4.2. Methods

The Gaussian mixture model approach is compared to two other methods for estimating distances:

PDS The Partial Distance Strategy [23]. Calculate the sum of squared differences of the mutually known components and scale to the missing components:

$$
\begin{equation*}
\hat{d}_{i j}^{2}=\frac{d}{\left|O_{i} \cap O_{j}\right|} \sum_{l \in O_{i} \cap O_{j}}\left(x_{i}^{l}-x_{j}^{l}\right)^{2} . \tag{13}
\end{equation*}
$$

For samples which have no known components in common, the method is not defined. For such pairs, the average of the pairwise distances which were possible to estimate is returned instead.

ICkNNI Incomplete-case k-NN imputation [6]. An improvement of completecase k-NN imputation, here any sample with a valid pattern of missing values is viable nearest neighbour. In accordance to the suggestions in [6], up to $k=5$ neighbours are considered. The imputation fails whenever there are no samples with such valid patterns. For these cases, the missing value is imputed by the sample mean for that variable.

In addition, the mixture model is compared to estimating the distribution by a single Gaussian. As another alternative, the distances are estimated after using the mixture model for imputation by the conditional mean - equivalent to discarding the variance terms of Eq. (11). Using a single Gaussian for imputation by the conditional mean is an interesting special case, as it is equivalent to a least-squares linear regression.

### 4.3. Performance criteria

The methods are evaluated by three different performance criteria. First, the methods are compared by the root mean squared error (RMSE) of all the estimated pairwise distances in the data set,

$$
C_{1}=\left(\frac{1}{\lambda} \sum_{i>j}\left(\hat{d}_{i j}-d_{i j}\right)^{2}\right)^{1 / 2}
$$

where $d_{i j}$ is the true Euclidean distance between samples $i$ and $j$ calculated without any missing data, and $\hat{d}_{i j}$ is the estimate of the distance provided by each method after removing data. The scaling factor $\lambda$ is determined so that the average is calculated only over those distances which are estimates, discarding all the cases where the distance can be calculated exactly because neither sample has any missing components: $\lambda=M N-M(M+1) / 2$, where $M$ is the number of samples having missing values.

A common application for pairwise distances is a nearest neighbour search, and thus we also consider the average (true) distance to the predicted nearest neighbour,

$$
C_{2}=\frac{1}{N} \sum_{i=1}^{N} d_{i, \mathrm{NN}(i)}, \quad \text { where } \quad \mathrm{NN}(i)=\arg \min _{j \neq i} \hat{d}_{i j}
$$

Here, $\mathrm{NN}(i)$ is the nearest neighbour of the $i$ th sample as estimated by the method, and $d_{i, \mathrm{NN}(i)}$ is the true Euclidean distance between the samples as calculated without any missing data. The criterion measures how well the method can identify samples which actually are close in the real data.

In addition, the accuracy is evaluated by the mean relative error of all pairwise distances:

$$
C_{3}=\frac{1}{\lambda} \sum_{i>j} \frac{\left|\hat{d}_{i j}-d_{i j}\right|}{d_{i j}}
$$

This criterion gives more weight to small distances, and is also an $\ell^{1}$-type error. (As some data sets contain duplicate samples, sample pairs $i, j$ where $d_{i j}=0$ are ignored when calculating this criterion.)

### 4.4. Procedure

Values are removed from the data set independently at a fixed probability $p$. For each value of $p, 100$ repetitions are conducted for the Monte Carlo simulation, and simulations are run for value of $p$ of $5 \%$ (low ratio of missing values), $20 \%$, and $50 \%$ (high ratio of missing values). The EM algorithm is run for 200 iterations, and repeated for a total of 5 times for each number of components. Runs are aborted if a covariance matrix becomes too poorly conditioned (condition number over $10^{12}$ ). The best solution in terms of loglikelihood is selected, and the number of components is selected by the $\mathrm{AIC}_{C}$ criterion.

Having 100 repetitions of the same set-up enables the use of statistical significance testing to assess the difference between the mean errors of different methods. The testing is conducted as a two-tailed paired $t$-test, with a significance level of $\alpha=0.05$. Comparing the performance of the best method to that of every other method results in a multiple hypothesis scenario, and thus the Bonferroni correction [31] is used to control the error rate.

### 4.5. Distance estimation results

The average RMSE values for the methods are presented in Table 2. The data sets are grouped into three categories as follows, depending on which mixture model is applicable. The first category includes data for which the EM algorithm converges appropriately with at least two components, and here the number of components is selected by the $\mathrm{AIC}_{C}$ criterion. For data in the second group, the EM algorithm either did not converge with two components, or the $\mathrm{AIC}_{C}$ indicated that a single component is clearly sufficient. Finally there are
the high-dimensional data sets (BCI, COIL, Tecator), for which the EM algorithm would not converge even with one component, and we apply the HDDC instead.

The most immediate observation is that using a single Gaussian or the mixture model tends to give the most accurate results for most data sets. For certain data (Computer Hardware, Ecoli, Image Segmentation), it appears that ICkNNI leads to better estimates.

For many data sets in the first category (Stocks, Iris, Statlog), the mixture model provides a clear advantage over using a single Gaussian. In other cases the differences are not significant, meaning that it may be sufficient to model the data with one Gaussian for this purpose, and indeed several runs resulted in mixture models of consisting of a single Gaussian, as evidenced by the mean $K$ values of less than two.

For most data sets, it can also be seen that including the variance terms of Equation (11) tends to lead to an improvement in the accuracy compared to only conducting imputation.

Table 3 shows the corresponding performances in terms of the true distance to the predicted nearest neighbour. This criterion is more emphasised on the accuracy of small distances, and reveals similar behaviour between the methods on most data sets. There are some notable differences, such as for the Computer Hardware data, where ICkNNI is the most accurate in terms of RMSE, but the mixture model approach is better for estimating nearest neighbours. For most data sets in the first group, the difference between using the mixture model and a single Gaussian is more pronounced than in Table 2.

Table 4 shows the mean relative error of each method. The relative performances are otherwise similar, except that in nearly all cases, using the mixture model for imputation is more accurate in terms of average relative error than directly estimating distances.

In terms of computational resources, the mixture model is more expensive than PDS or ICkNNI (which do not require multiple iterations over the data), but correspondingly delivers more accurate results. As such, the mixture model can be recommended in any situation where the computational task is feasible. Comparing imputation by the mixture model to using the same model for directly estimating distances, there is no difference in computational cost, but the directly estimated distances tend to be more accurate.

### 4.6. Regression and classification results

In order to study the effect of estimated distances on the performance of supervised learning tasks we assume a semi-supervised learning scenario. In particular, we have a full set of input samples, but labels only for a subset of them, and the task is to label the unlabelled samples. In the current situation this implies that the distances are estimated using the full set of input samples (ignoring labels), and a machine learning model is built using the distances between the labelled samples only (i.e., the "training set").

The model used is an Extreme Learning Machine model with RBF kernels $[32,33]$. The kernel centres $\boldsymbol{\mu}_{j}$ are randomly selected samples from the training
set, and widths $\sigma_{j}$ are assigned randomly. The number of neurons is fixed at $N / 10$, where $N$ is the number of samples. The model is not optimised at all so that the comparison would be as fair as possible, i.e., the models are exactly the same, the only change is the estimated distances (and the final weights of the output layer after training the ELM). The training phase only consists in finding the least-squares solution to the linear system

$$
\mathbf{H} \beta=\mathbf{T}
$$

where $\mathbf{T}$ is the target output of the labelled data and the hidden layer output matrix $\mathbf{H}$ has the elements

$$
H_{i j}=\exp \left(-\frac{\left\|\boldsymbol{x}_{i}-\boldsymbol{\mu}_{j}\right\|^{2}}{\sigma_{j}^{2}}\right)
$$

as well as a constant column of ones to account for the bias term. In the current scenario, $\mathbf{T}, \boldsymbol{\mu}_{j}$, and $\sigma_{j}$ are equal for each method. The only difference is how the distance $\left\|\boldsymbol{x}_{i}-\boldsymbol{\mu}_{j}\right\|$ is estimated considering the missing values.

In the simulations, the data samples were randomly partitioned into $75 \%$ for training and $25 \%$ for testing.

The results in terms of average prediction MSE (normalised by the variance of the target variable) for regression tasks are shown in Table 5, and prediction accuracies for classification tasks in Table 6. In a majority of cases, using a mixture model or single Gaussian to estimate distances results in more accurate predictions than PDS or ICkNNI. Comparing the relative performances to the relative accuracies in Table 2, it seems evident that, in general, more accurately estimated distances lead to a better performance in prediction tasks.

## 5. Conclusions

In this paper, we present how mixtures of Gaussians can be used to effectively estimate pairwise distances in incomplete data. The problem of estimating distances in a data set with missing values can be reduced to finding the conditional means and variances separately for each missing value. Having a Gaussian mixture model of the distribution of the data enables these quantities to be estimated. In order to fit the mixture model, certain extensions to the standard EM algorithm are presented in Section 2, including a variant specifically for high-dimensional data.

The combination of these ideas provides for a method to estimate distances, and the simulations in Section 4 show that the method is competitive, if not better, than alternative methods in terms of accuracy. The increased accuracy in estimated distances is reflected in an improved prediction performance of neural networks built using the distances.

For future work, it remains to investigate the most effective ways to extend the method to high-dimensional cases where the number of dimensions would exceed the number of samples. Furthermore, it will be interesting to study the influence of the distance estimation when used with other machine learning methods, such as SVM.

Table 2: Average RMSE of estimated pairwise distances. The best result for each row is underlined, and any results which are not statistically significantly different (two-tailed paired $t$-test, $\alpha=0.05$ ) from the best result are bolded. The values in parenthesis represent the accuracy when the distances are calculated using the particular model for imputation only. The final column shows the mean number of Gaussian components $K$ as selected by the AIC $C_{C}$ criterion, and the mean number of distinct eigenvalues $d_{k}$ for HDDC.

|  |  | PDS | ICkNNI | Single | Gaussian | Mixture model |  | Mean K |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Computer Hardware $N=209, d=6$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.676 \\ & 1.102 \\ & 1.938 \end{aligned}$ | $\frac{0.408}{\frac{0.710}{1.340}}$ | $\begin{aligned} & 0.459 \\ & 0.736 \\ & \mathbf{1 . 2 7 3} \\ & \hline \end{aligned}$ | $\begin{aligned} & (0.443) \\ & (0.737) \\ & (1.327) \end{aligned}$ | $\begin{aligned} & 0.451 \\ & \mathbf{0 . 7 3 2} \\ & 1.331 \end{aligned}$ | $\begin{gathered} (0.441) \\ (0.730) \\ (1.341) \end{gathered}$ | $\begin{aligned} & 3.82 \\ & 3.70 \\ & 3.52 \end{aligned}$ |
| Glass Identification $N=214, d=9$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.526 \\ & 0.971 \\ & 2.098 \end{aligned}$ | $\begin{aligned} & 0.337 \\ & 0.706 \\ & 1.540 \end{aligned}$ | $\begin{aligned} & 0.226 \\ & 0.524 \\ & \mathbf{1 . 1 9 7} \end{aligned}$ | $\begin{aligned} & (\mathbf{0 . 2 2 1 )} \\ & (0.522) \\ & (1.281) \end{aligned}$ | $\begin{aligned} & 0.231 \\ & 0.519 \\ & \hline 1.197 \end{aligned}$ | $\begin{aligned} & (0.228) \\ & (0.532) \\ & (1.265) \end{aligned}$ | $\begin{aligned} & 2.31 \\ & 2.51 \\ & 2.45 \end{aligned}$ |
| Housing $N=506, d=13$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.514 \\ & 1.001 \\ & 2.269 \end{aligned}$ | $\begin{aligned} & \frac{0.329}{0.672} \\ & 1.593 \end{aligned}$ | $\begin{aligned} & 0.338 \\ & 0.597 \\ & \mathbf{1} .066 \\ & \hline \end{aligned}$ | $\begin{aligned} & (0.348) \\ & (0.650) \\ & (1.330) \end{aligned}$ | $\begin{aligned} & 0.331 \\ & 0.587 \\ & \hline 1.104 \end{aligned}$ | $\begin{aligned} & (0.338) \\ & (0.619) \\ & (1.245) \end{aligned}$ | $\begin{aligned} & 3.34 \\ & 3.26 \\ & 3.21 \end{aligned}$ |
| Image Segmentation $N=2310, d=14$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.632 \\ & 1.254 \\ & 2.801 \end{aligned}$ | $\begin{aligned} & 0.376 \\ & \hline 0.774 \\ & 1.788 \end{aligned}$ | $\begin{aligned} & 0.389 \\ & \mathbf{0 . 7 2 0} \\ & \hline \mathbf{1 . 3 6 4} \end{aligned}$ | $\begin{aligned} & (0.397) \\ & (0.759) \\ & (1.560) \end{aligned}$ | $\begin{aligned} & 0.399 \\ & 0.778 \\ & 1.439 \end{aligned}$ | $\begin{gathered} (0.395) \\ (0.778) \\ (1.490) \end{gathered}$ | $\begin{aligned} & 4.24 \\ & 4.12 \\ & 3.99 \end{aligned}$ |
| $\begin{aligned} & \text { Iris } \\ & N=150, d=4 \end{aligned}$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.440 \\ & 0.676 \\ & 1.106 \end{aligned}$ | $\begin{aligned} & 0.246 \\ & 0.391 \\ & 0.898 \end{aligned}$ | $\begin{aligned} & 0.253 \\ & 0.379 \\ & 0.766 \end{aligned}$ | $\begin{aligned} & (0.257) \\ & (0.401) \\ & (0.897) \end{aligned}$ | $\begin{aligned} & 0.219 \\ & 0.335 \\ & \hline \mathbf{0 . 7 3 8} \end{aligned}$ | $\begin{aligned} & (0.225) \\ & (0.358) \\ & (0.845) \end{aligned}$ | $\begin{aligned} & 3.00 \\ & 2.49 \\ & 2.79 \end{aligned}$ |
| Pima Indians $N=768, d=8$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.540 \\ & 0.957 \\ & 1.910 \end{aligned}$ | $\begin{aligned} & 0.413 \\ & 0.692 \\ & 1.329 \end{aligned}$ | $\begin{aligned} & 0.388 \\ & \begin{array}{l} 0.600 \\ 0.973 \end{array} \end{aligned}$ | $\begin{aligned} & (0.418) \\ & (0.714) \\ & (1.441) \end{aligned}$ | $\begin{aligned} & 0.395 \\ & 0.610 \\ & 0.998 \end{aligned}$ | $\begin{aligned} & (0.416) \\ & (0.691) \\ & (1.335) \end{aligned}$ | $\begin{aligned} & 4.74 \\ & 4.47 \\ & 4.27 \end{aligned}$ |
| Servo $N=167, d=4$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.575 \\ & 0.864 \\ & 1.217 \end{aligned}$ | $\begin{aligned} & 0.463 \\ & 0.629 \\ & 1.089 \end{aligned}$ | $\begin{aligned} & \frac{0.383}{0.509} \\ & \underline{0.761} \end{aligned}$ | $\begin{aligned} & (0.425) \\ & (0.610) \\ & (1.122) \end{aligned}$ | $\begin{aligned} & 0.399 \\ & 0.536 \\ & 0.835 \end{aligned}$ | $\begin{aligned} & (0.437) \\ & (0.620) \\ & (1.110) \end{aligned}$ | $\begin{aligned} & 3.41 \\ & 3.86 \\ & 4.14 \end{aligned}$ |
| Stocks $N=950, d=9$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.350 \\ & 0.649 \\ & 1.516 \end{aligned}$ | $\begin{aligned} & \mathbf{0 . 0 6 2} \\ & 0.179 \\ & 0.864 \end{aligned}$ | $\begin{aligned} & 0.183 \\ & 0.343 \\ & 0.780 \end{aligned}$ | $\begin{aligned} & (0.182) \\ & (0.352) \\ & (0.897) \end{aligned}$ | $\begin{array}{r} 0.076 \\ \mathbf{0 . 1 4 7} \\ \hline \mathbf{0 . 4 7 9} \\ \hline \end{array}$ | $\begin{aligned} & (0.077) \\ & (0.150) \\ & (0.533) \end{aligned}$ | $\begin{aligned} & 9.44 \\ & 8.72 \\ & 6.69 \end{aligned}$ |
| Statlog $N=846, d=18$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.343 \\ & 0.728 \\ & 1.779 \end{aligned}$ | $\begin{aligned} & 0.199 \\ & 0.562 \\ & 1.831 \end{aligned}$ | $\begin{aligned} & 0.156 \\ & 0.308 \\ & 0.708 \end{aligned}$ | $\begin{aligned} & (0.161) \\ & (0.331) \\ & (0.813) \end{aligned}$ | $\begin{aligned} & \frac{0.142}{0.287} \\ & \underline{0.645} \end{aligned}$ | $\begin{aligned} & (0.147) \\ & (0.305) \\ & (0.731) \end{aligned}$ | $\begin{aligned} & 2.62 \\ & 2.60 \\ & 2.11 \end{aligned}$ |
|  |  | PDS | ICkNNI | Single Gaussian |  |  |  |  |
| Auto-price $N=159, d=15$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.368 \\ & 0.750 \\ & 1.820 \end{aligned}$ | $\begin{aligned} & 0.263 \\ & 0.733 \\ & 1.687 \end{aligned}$ | $\begin{aligned} & \frac{0.236}{0.455} \\ & \hline 0.871 \end{aligned}$ | $\begin{aligned} & (0.239) \\ & (0.485) \\ & (1.015) \end{aligned}$ |  |  |  |
| Breast Cancer (Diag.) $N=569, d=30$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.358 \\ & 0.806 \\ & 1.913 \end{aligned}$ | $\begin{aligned} & 0.260 \\ & 1.086 \\ & 2.551 \end{aligned}$ | $\begin{aligned} & 0.141 \\ & \begin{array}{l} 0.344 \\ 0.802 \end{array} \end{aligned}$ | $\begin{aligned} & (0.142) \\ & (0.352) \\ & (0.871) \end{aligned}$ |  |  |  |
| $\begin{aligned} & \text { Breast Cancer (Prog.) } \\ & N=194, d=32 \end{aligned}$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.349 \\ & 0.790 \\ & 1.872 \end{aligned}$ | $\begin{aligned} & 0.298 \\ & 1.071 \\ & 2.612 \end{aligned}$ | $\begin{aligned} & \frac{0.166}{0.383} \\ & 0.843 \end{aligned}$ | $\begin{aligned} & (0.169) \\ & (0.395) \\ & (0.911) \end{aligned}$ |  |  |  |
| Breast Tissue $N=106, d=9$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.424 \\ & 0.784 \\ & 1.783 \end{aligned}$ | $\begin{aligned} & 0.243 \\ & 0.586 \\ & 1.594 \end{aligned}$ | $\begin{aligned} & 0.221 \\ & \begin{array}{l} 0.415 \\ \hline 0.959 \end{array} \end{aligned}$ | $\begin{aligned} & (0.222) \\ & (0.420) \\ & (0.996) \end{aligned}$ |  |  |  |
| Ecoli $N=336, d=7$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.717 \\ & 1.239 \\ & 2.300 \end{aligned}$ | $\frac{0.432}{\frac{0.737}{1.525}}$ | $\begin{aligned} & 0.441 \\ & 0.745 \\ & 1.387 \\ & \hline \end{aligned}$ | $\begin{aligned} & (0.440) \\ & (0.767) \\ & (1.591) \end{aligned}$ |  |  |  |
| Financial Ratios $N=500, d=40$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.682 \\ & 1.516 \\ & 3.207 \end{aligned}$ | $\begin{aligned} & 0.483 \\ & 1.447 \\ & 3.152 \end{aligned}$ | $\begin{aligned} & 0.391 \\ & \mathbf{0 . 9 1 0} \\ & \mathbf{1 . 7 4 1} \\ & \hline \end{aligned}$ | $\begin{aligned} & \left(\frac{(0.387)}{(0.905)}\right. \\ & (1.794) \end{aligned}$ |  |  |  |
| Ionosphere $N=351, d=33$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.276 \\ & 0.624 \\ & 1.550 \end{aligned}$ | $\begin{aligned} & 0.257 \\ & 1.099 \\ & 2.603 \end{aligned}$ | $\begin{aligned} & 0.243 \\ & 0.584 \\ & 1.112 \end{aligned}$ | $\begin{aligned} & (0.223) \\ & (\underline{0.514}) \\ & (1.073) \end{aligned}$ |  |  |  |
| Parkinsons $N=195, d=20$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.346 \\ & 0.742 \\ & 1.780 \end{aligned}$ | $\begin{aligned} & 0.264 \\ & 0.901 \\ & 2.066 \end{aligned}$ | $\begin{aligned} & 0.188 \\ & \begin{array}{l} 0.385 \\ \hline 0.772 \end{array} \end{aligned}$ | $\begin{aligned} & (0.193) \\ & (0.409) \\ & (0.892) \end{aligned}$ |  |  |  |
| Spambase $N=4601, d=57$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.988 \\ & 2.205 \\ & 4.664 \end{aligned}$ | $\begin{aligned} & 0.695 \\ & 1.859 \\ & 3.940 \end{aligned}$ | $\begin{aligned} & \frac{0.657}{1.364} \\ & \underline{2.236} \end{aligned}$ | $\begin{aligned} & (0.688) \\ & (1.609) \\ & (3.280) \end{aligned}$ |  |  |  |
| SPECTF Heart $N=267, d=44$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.334 \\ & 0.763 \\ & 1.815 \end{aligned}$ | $\begin{aligned} & 0.331 \\ & 1.190 \\ & 2.970 \end{aligned}$ | $\begin{aligned} & \frac{0.203}{0.491} \\ & 1.066 \end{aligned}$ | $\begin{aligned} & (0.206) \\ & (0.495) \\ & (1.141) \end{aligned}$ |  |  |  |
| Wine $N=178, d=13$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.361 \\ & 0.734 \\ & 1.783 \end{aligned}$ | $\begin{aligned} & 0.267 \\ & 0.623 \\ & 1.469 \end{aligned}$ | $\begin{aligned} & \frac{0.248}{0.469} \\ & \hline 0.881 \end{aligned}$ | $\begin{aligned} & (0.262) \\ & (0.540) \\ & (1.173) \end{aligned}$ |  |  |  |
|  |  | PDS | ICkNNI |  |  | HDDC mixture |  | Mean $d_{k}$ |
| $\begin{aligned} & \mathrm{BCI} \\ & N=400, d=117 \end{aligned}$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.284 \\ & 0.649 \\ & 1.518 \end{aligned}$ | $\begin{aligned} & 0.548 \\ & 1.815 \\ & 4.768 \end{aligned}$ |  |  | $\begin{aligned} & \frac{0.108}{0.249} \\ & \hline \underline{0.610} \end{aligned}$ | $\begin{aligned} & (0.113) \\ & (0.273) \\ & (0.692) \end{aligned}$ | $\begin{aligned} & 23.31 \\ & 23.96 \\ & 18.19 \end{aligned}$ |
| COIL $N=1500, d=241$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \end{gathered}$ | $\begin{aligned} & 0.404 \\ & 0.926 \\ & 2.123 \end{aligned}$ | $\begin{aligned} & 0.689 \\ & 2.538 \\ & 6.709 \end{aligned}$ |  |  | $\begin{aligned} & \frac{0.202}{0.488} \\ & \hline 1.067 \end{aligned}$ | $\begin{aligned} & (0.224) \\ & (0.630) \\ & (1.625) \end{aligned}$ | $\begin{aligned} & 26.23 \\ & 26.10 \\ & 21.23 \end{aligned}$ |
| Tecator $N=240, d=100$ | $\begin{gathered} 5 \% \\ 20 \% \\ 50 \% \\ \hline \end{gathered}$ | $\begin{aligned} & 0.052 \\ & 0.120 \\ & 0.282 \end{aligned}$ | $\begin{aligned} & 1.163 \\ & 2.366 \\ & 5.218 \\ & \hline \end{aligned}$ |  |  | $\begin{aligned} & 0.004 \\ & 0.008 \\ & 0.017 \end{aligned}$ | $\begin{array}{r} (0.004) \\ (0.008) \\ (0.016) \end{array}$ | $\begin{aligned} & 4.00 \\ & 4.00 \\ & 4.00 \\ & \hline \end{aligned}$ |

16

Table 3: Average of the mean distance to the estimated nearest neighbour. The best result for each row is underlined, and any results which are not statistically significantly different (two-tailed paired $t$-test, $\alpha=0.05$ ) from the best result are bolded. The values in parenthesis represent the accuracy when the distances are calculated using the particular model for imputation only. The final column shows the mean number of Gaussian components $K$ as selected by the $\mathrm{AIC}_{C}$ criterion, and the mean number of distinct eigenvalues $d_{k}$ for HDDC.


Table 4: Average relative error of estimated pairwise distances. The best result for each row is underlined, and any results which are not statistically significantly different (two-tailed paired $t$-test, $\alpha=0.05$ ) from the best result are bolded. The values in parenthesis represent the accuracy when the distances are calculated using the particular model for imputation only. The final column shows the mean number of Gaussian components $K$ as selected by the $\mathrm{AIC}_{C}$ criterion, and the mean number of distinct eigenvalues $d_{k}$ for HDDC.

|  |  | PDS | ICkNNI | Single | Gaussian | Mixture model | Mean |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  |  |  |  |  |

Table 5: Average normalised MSE of ELM predictions for regression tasks. The best result for each row is underlined, and any results which are not statistically significantly different (two-tailed paired $t$-test, $\alpha=0.05$ ) from the best result are bolded. The values in parenthesis represent the accuracy when the distances are calculated using the particular model for imputation only. The final column shows the mean number of Gaussian components $K$ as selected by the $\mathrm{AIC}_{C}$ criterion, and the mean number of distinct eigenvalues $d_{k}$ for HDDC.

|  |  | PDS | ICkNNI | Single | Gaussian | Mixtur | e model | Mean $K$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Computer Hardware$N=209, d=6$ |  | 0.382 | 0.351 | 0.354 | (0.351) | 0.353 | (0.353) | 3.82 |
|  | $20 \%$ | 0.427 | $\frac{0.351}{0.383}$ | 0.385 | (0.381) | 0.375 | (0.382) | 3.70 |
|  | 50\% | 0.558 | 0.540 | 0.491 | (0.531) | 0.486 | (0.532) | 3.52 |
| Housing$N=506, d=13$ | 5\% | 0.242 | 0.199 | 0.199 | (0.199) | 0.198 | (0.198) | 3.34 |
|  | 20\% | 0.342 | 0.279 | 0.255 | (0.256) | 0.255 | (0.258) | 3.26 |
|  | 50\% | 0.567 | 0.593 | $\underline{0.419}$ | (0.446) | 0.433 | (0.461) | 3.21 |
| Servo$N=167, d=4$ | 5\% | 0.444 | 0.359 | 0.358 | (0.350) | 0.360 | (0.353) | 3.41 |
|  | 20\% | 0.644 | 0.532 | 0.533 | (0.522) | 0.540 | (0.539) | 3.86 |
|  | 50\% | 0.798 | 0.816 | $\underline{0.772}$ | (0.799) | 0.793 | (0.837) | 4.14 |
| Stocks$N=950, d=9$ | 5\% | 0.039 | 0.019 | 0.028 | (0.027) | 0.020 | (0.020) | 9.44 |
|  | 20\% | 0.085 | 0.028 | 0.062 | (0.059) | 0.026 | (0.026) | 8.72 |
|  | 50\% | 0.302 | 0.226 | 0.230 | (0.218) | 0.101 | (0.107) | 6.69 |
|  |  | PDS | ICkNNI | Single | Gaussian |  |  |  |
| Auto-price$N=159, d=15$ | 5\% | 0.253 | 0.223 | 0.221 | (0.221) |  |  |  |
|  | 20\% | 0.315 | 0.245 | 0.224 | (0.224) |  |  |  |
|  | 50\% | 0.438 | 0.483 | 0.265 | (0.262) |  |  |  |
| $\begin{aligned} & \text { Breast Cancer (Prog.) } \\ & N=194, d=32 \end{aligned}$ | 5\% | 0.949 | 0.940 | 0.939 | (0.940) |  |  |  |
|  | 20\% | 0.974 | 0.943 | 0.939 | (0.944) |  |  |  |
|  | 50\% | 1.028 | 0.975 | $\underline{0.944}$ | (0.961) |  |  |  |
|  |  | PDS | ICkNNI |  |  | HDDC mixture |  | Mean $d_{k}$ |
| Tecator$N=240, d=100$ | 5\% | 0.285 | 0.364 |  |  | 0.290 | (0.290) | 4.00 |
|  | 20\% | 0.348 | 0.471 |  |  | 0.331 | (0.333) | 4.00 |
|  | 50\% | 0.519 | 0.788 |  |  | 0.427 | (0.520) | 4.00 |

Table 6: Average ELM correct classification rate for classification tasks. The best result for each row is underlined, and any results which are not statistically significantly different (twotailed paired $t$-test, $\alpha=0.05$ ) from the best result are bolded. The values in parenthesis represent the accuracy when the distances are calculated using the particular model for imputation only. The final column shows the mean number of Gaussian components $K$ as selected by the $\mathrm{AIC}_{C}$ criterion, and the mean number of distinct eigenvalues $d_{k}$ for HDDC.

[1] C. M. Bishop, Pattern Recognition and Machine Learning, Springer, 2006.
[2] R. J. A. Little, D. B. Rubin, Statistical Analysis with Missing Data, WileyInterscience, second edition, 2002.
[3] D. B. Rubin, Multiple Imputation for Nonresponse in Surveys, Wiley Series in Probability and Statistics, Wiley, 1987.
[4] C. K. Enders, Applied Missing Data Analysis, Methodology In The Social Sciences, Guilford Press, 2010.
[5] E. R. Hruschka, E. R. Hruschka Jr., N. F. F. Ebecken, Evaluating a nearestneighbor method to substitute continuous missing values, in: AI 2003: Advances in Artificial Intelligence, volume 2903 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2003, pp. 723-734.
[6] J. Van Hulse, T. M. Khoshgoftaar, Incomplete-case nearest neighbor imputation in software measurement data, Information Sciences (2011). In press, available online 9 January 2011, DOI: 10.1016/j.ins.2010.12.017.
[7] S. Van Buuren, J. P. Brand, C. G. Groothuis-Oudshoorn, D. B. Rubin, Fully conditional specification in multivariate imputation, Journal of Statistical Computation and Simulation 76 (2006) 1049-1064.
[8] A. Aussem, S. R. de Morais, A conservative feature subset selection algorithm with missing data, Neurocomputing 73 (2010) 585-590.
[9] G. Doquire, M. Verleysen, Feature selection with missing data using mutual information estimators, Neurocomputing 90 (2012) 3-11.
[10] J. W. Grzymala-Busse, W. J. Grzymala-Busse, Handling missing attribute values, in: O. Maimon, L. Rokach (Eds.), Data Mining and Knowledge Discovery Handbook, Springer US, second edition, 2010, pp. 33-51.
[11] P. J. García-Laencina, J.-L. Sancho-Gómez, A. R. Figueiras-Vidal, M. Verleysen, K nearest neighbours with mutual information for simultaneous classification and missing data imputation, Neurocomputing 72 (2009) 1483-1493.
[12] A. P. Dempster, N. M. Laird, D. B. Rubin, Maximum likelihood from incomplete data via the EM algorithm, Journal of the Royal Statistical Society. Series B (Methodological) 39 (1977) pp. 1-38.
[13] Z. Ghahramani, M. Jordan, Learning From Incomplete Data, Technical Report, Lab Memo No. 1509, CBCL Paper No. 108, MIT AI Lab, 1995.
[14] L. Hunt, M. Jorgensen, Mixture model clustering for mixed data with missing information, Computational Statistics \& Data Analysis 41 (2003) 429-440.
[15] T. I. Lin, J. C. Lee, H. J. Ho, On fast supervised learning for normal mixture models with missing information, Pattern Recognition 39 (2006) 1177-1187.
[16] R. J. Steele, N. Wang, A. E. Raftery, Inference from multiple imputation for missing data using mixtures of normals, Statistical Methodology 7 (2010) 351-365.
[17] O. Delalleau, A. C. Courville, Y. Bengio, Efficient EM training of Gaussian mixtures with missing data, CoRR abs/1209.0521 (2012).
[18] V. Tresp, S. Ahmad, R. Neuneier, Training neural networks with deficient data, in: Advances in Neural Information Processing Systems 6, Morgan Kaufmann, 1994, pp. 128-135.
[19] Z. Viharos, L. Monostori, T. Vincze, Training and application of artificial neural networks with incomplete data, in: T. Hendtlass, M. Ali (Eds.), Developments in Applied Artificial Intelligence, volume 2358 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2002, pp. 649-659.
[20] A. Morris, M. Cooke, P. Green, Some solutions to the missing feature problem in data classification, with application to noise robust ASR, in: Acoustics, Speech and Signal Processing, 1998. Proceedings of the 1998 IEEE International Conference on, volume 2, pp. 737-740.
[21] C. Beunckens, G. Molenberghs, M. G. Kenward, Direct likelihood analysis versus simple forms of imputation for missing data in randomized clinical trials, Clinical Trials 2 (2005) 379-386.
[22] C. Bouveyron, S. Girard, C. Schmid, High-dimensional data clustering, Computational Statistics \& Data Analysis 52 (2007) 502-519.
[23] J. K. Dixon, Pattern recognition with partly missing data, Systems, Man and Cybernetics, IEEE Transactions on 9 (1979) 617-621.
[24] T. W. Anderson, An Introduction to Multivariate Statistical Analysis, Wiley-Interscience, New York, third edition, 2003.
[25] R. B. Cattell, The scree test for the number of factors, Multivariate Behavioral Research 1 (1966) 245-276.
[26] H. Akaike, A new look at the statistical model identification, Automatic Control, IEEE Transactions on 19 (1974) 716-723.
[27] C. M. Hurvich, C.-L. Tsai, Regression and time series model selection in small samples, Biometrika 76 (1989) 297-307.
[28] A. Asuncion, D. J. Newman, UCI machine learning repository, 2012. University of California, Irvine, School of Information and Computer Sciences, http://archive.ics.uci.edu/ml/.
[29] L. Torgo, Regression datasets, 2012. University of Porto, http://www.liaad.up.pt/~1torgo/Regression/DataSets.html.
[30] O. Chapelle, B. Schölkopf, A. Zien (Eds.), Semi-Supervised Learning, MIT Press, Cambridge, MA, 2006.
[31] J. Shaffer, Multiple hypothesis testing, Annual review of psychology 46 (1995) 561-584.
[32] G.-B. Huang, Q.-Y. Zhu, C.-K. Siew, Extreme learning machine: Theory and applications, Neurocomputing 70 (2006) 489-501.
[33] G.-B. Huang, C.-K. Siew, Extreme learning machine with randomly assigned RBF kernels, International Journal of Information Technology 11 (2005) 16-24.


[^0]:    * Corresponding author

    Email address: emil.eirola@aalto.fi (Emil Eirola)

