

Nonparametric imputation by data depth

Pavlo Mozharovskyi*

CREST, Ensai, Université Bretagne Loire

and

Julie Josse

École Polytechnique, CMAP

and

François Husson

IRMAR, Applied Mathematics Unit, Agrocampus Ouest, Rennes

August 6, 2018

Abstract

We present single imputation method for missing values which borrows the idea of data depth—a measure of centrality defined for an arbitrary point of a space with respect to a probability distribution or data cloud. This consists in iterative maximization of the depth of each observation with missing values, and can be employed with any properly defined statistical depth function. For each single iteration, imputation reverts to optimization of quadratic, linear, or quasiconcave functions that are solved analytically by linear programming or the Nelder-Mead method. As it accounts for the underlying data topology, the procedure is distribution free, allows imputation close to the data geometry, can make prediction in situations where local imputation (k -nearest neighbors, random forest) cannot, and has attractive robustness and asymptotic properties under elliptical symmetry. It is shown that a special case—when using the Mahalanobis depth—has direct connection to well-known methods for the multivariate normal model, such as iterated regression and regularized PCA. The methodology is extended to multiple imputation for data stemming from an elliptically symmetric distribution. Simulation and real data studies show good results compared with existing popular alternatives. The method has been implemented as an R-package. Supplementary materials for the article are available online.

Keywords: Elliptical symmetry, Outliers, Tukey depth, Zonoid depth, Local depth, Nonparametric imputation, Convex optimization.

—*The major part of this project has been conducted during the postdoc of Pavlo Mozharovskyi at Agrocampus Ouest (Rennes) granted by *Centre Henri Lebesgue* due to program PIA-ANR-11-LABX-0020-01.

1 Introduction

Missing data is a ubiquitous problem in statistics. Non-responses to surveys, machines that break and stop reporting, and data that have not been recorded, impede analysis and threaten the validity of inference. A common strategy (Little and Rubin, 2002) for dealing with missing values is single imputation, replacing missing entries with plausible values to obtain a completed data set, which can then be analyzed.

There are two main families of parametric imputation methods: “joint” and “conditional” modeling, see e.g., Josse and Reiter (2018) for a literature overview. Joint modeling specifies a joint distribution for the data, the most popular being the normal multivariate distribution. The parameters of the distribution, here the mean and the covariance matrix, are then estimated from the incomplete data using an algorithm such as expectation maximization (EM) (Dempster et al., 1977). The missing entries are then imputed with the conditional mean, i.e., the conditional expectation of the missing values, given observed values and the estimated parameters. An alternative is to impute missing values using a principal component analysis (PCA) model which assumes data are generated as a low rank structure corrupted by Gaussian noise. This method is closely connected to the literature on matrix completion Josse and Husson (2012), Hastie et al. (2015), and has shown good imputation capacity due to the plausibility of the low rank assumption (Udell and Townsend, 2017). The conditional modeling approach (van Buuren, 2012) consists in specifying one model for each variable to be imputed, and considers the others variables as explanatory. This procedure is iterated until predictions stabilize. Nonparametric imputation methods have also been developed such as imputation by k -nearest neighbors (k NN) (see Troyanskaya et al., 2001, and references therein) or random forest (Stekhoven and Bühlmann, 2012).

Most imputation methods are defined under the missing (completely) at random (M(C)AR) assumption, which means that the probability of having missing values does not depend on missing data (nor on observed data). Gaussian and PCA imputations are sensitive to outliers and deviations from distributional assumptions, whereas nonparametric methods such as k NN and random forest cannot extrapolate.

Here we propose a family of nonparametric imputation methods based on the notion of a statistical depth function (Tukey, 1975). Data depth is a data-driven multivariate measure of centrality that describes data with respect to location, scale, and shape based on a multivariate ordering. It has been applied in multivariate data analysis (Liu et al., 1999), classification (Jörnsten, 2004, Lange et al., 2014), multivariate risk measurement (Casco and Molchanov, 2007), and robust linear programming (Bazovkin and Mosler, 2015), but has never been applied in the context of missing data. Depth based imputation provides excellent predictive properties and has the advantages of both global and local imputation methods. It imputes close to the data geometry, while still accounting for global features. In addition, it allows robust imputation in both outliers and heavy-tailed distributions.

Figures 1 and 2 motivate our proposed depth-based imputation by contrasting it to classical methods. First, 150 points are drawn from a bivariate normal distribution with mean $\mu_1 = (1, 1)^T$ and covariance $\Sigma_1 = \begin{pmatrix} 1 & 1 \\ 1 & 4 \end{pmatrix}$ and 30% of the entries are removed completely at random in both variables; points with one missing entry are indicated by dotted lines while solid lines provide (oracle) imputation using distribution parameters. The imputation assuming a joint Gaussian distribution using EM estimates is shown by rhombi (Figure 1, left). Zonoid depth-based imputation, represented by filled circles, shows that the sample is not necessarily normal, and that this uncertainty increases as we move to the fringes of the data cloud, where imputed points deviate from the conditional mean towards the unconditional one. Second, the missing




Figure 1: Bivariate normal distribution with 30% MCAR (left) and with MAR in the second coordinate for values > 3.5 (right); imputation using maximum zonoid depth (filled circles), conditional mean imputation using EM estimates (rhombi), and random forest imputation (triangles).

values are generated as follows: the first coordinate is removed when the second coordinate > 3.5 (Figure 1, right). Here, the depth-based imputation allows extrapolation when predicting missing values, while the random forest imputation (triangles) gives, as expected, rather poor results.

In Figure 2 (left), we draw 500 points, 425 from the same normal distribution as above, with 15% of MCAR values and 75 outliers from the Cauchy distribution with the same center and shape matrix and without missing values. In Figure 2 (right), we depict 1000 points drawn from Cauchy distribution with 15% MCAR. As expected, imputation with conditional mean based on EM estimates (rhombi) is rather random. Depth-based imputation with Tukey depth (filled circles) has robust imputed values that are close to the (distribution's) regression lines reflecting data geometry.

The paper is organized as follows. Section 2 describes the algorithm for imputing by data depth and derives its theoretical properties under ellipticity. Section 3 describes the special case of imputation with Mahalanobis depth, emphasizing its relationship to existing imputation methods by regression and PCA, and imputation with zonoid and Tukey depths. For each of them, we suggest an efficient optimization strategy. Next, to go beyond ellipticity, we propose imputation with local depth (Paindaveine and Bever, 2013) appropriate to data with non-convex support. Section 4 provides a comparative simulation and real data study. Section 5 extends the proposed approach to multiple imputation in order to perform statistical inference with missing values. Section 6 concludes the article, gathering together some useful remarks. Proofs are available in the supplementary materials.




Figure 2: Left: Mixture of normal (425 points, 15% MCAR) and Cauchy (75 points) samples. Right: 1000 Cauchy distributed points with 15% MCAR. Imputation with Tukey depth (filled circles) and conditional mean imputation using EM estimates (rhombi).

2 Imputation by depth maximization

2.1 Imputation by iterative regression

Let X be a random vector in \mathbb{R}^d and denote $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top$ a sample. For a point $\mathbf{x}_i \in \mathbf{X}$, we denote $miss(i)$ and $obs(i)$ the sets of its coordinates containing missing and observed values, $|miss(i)|$ and $|obs(i)|$ their corresponding cardinalities.

Let the rows \mathbf{x}_i be i.i.d. draws from $\mathbf{N}(\boldsymbol{\mu}_X, \boldsymbol{\Sigma}_X)$. One of the simplest conditional methods for imputing missing values consists in the following iterative regression imputation: (1) initialize missing values arbitrary, using unconditional mean imputation; (2) impute missing values in one variable by the values predicted by the regression model of this variable with the remaining variables taken as explanatory ones, (3) iterate through variables containing missing values until convergence. Here, at each step, each point \mathbf{x}_i with missing values at a coordinate j is imputed with the univariate conditional mean $\mathbb{E}[X_j | X_{\{1, \dots, d\} \setminus \{j\}} = \mathbf{x}_{i, \{1, \dots, d\} \setminus \{j\}}, \boldsymbol{\mu}_X = \boldsymbol{\mu}_X, \boldsymbol{\Sigma}_X = \boldsymbol{\Sigma}_X]$ with the moment estimates $\boldsymbol{\mu}_X = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i$ and $\boldsymbol{\Sigma}_X = \frac{1}{n-1} \sum_{i=1}^n (\mathbf{x}_i - \boldsymbol{\mu}_X)(\mathbf{x}_i - \boldsymbol{\mu}_X)^\top$.

After convergence, each point \mathbf{x}_i with missing values in $miss(i)$ is imputed with the multivariate conditional mean

$$\begin{aligned} & \mathbb{E}[\mathbf{X} | \mathbf{X}_{obs(i)} = \mathbf{x}_{i, obs(i)}, \boldsymbol{\mu}_X = \boldsymbol{\mu}_X, \boldsymbol{\Sigma}_X = \boldsymbol{\Sigma}_X] \\ &= \boldsymbol{\mu}_X_{miss(i)} + \boldsymbol{\Sigma}_X_{miss(i), obs(i)} \boldsymbol{\Sigma}_X_{obs(i), obs(i)}^{-1} (\mathbf{x}_{i, obs(i)} - \boldsymbol{\mu}_X_{obs(i)}) \end{aligned} \quad (1)$$

The last expression is the closed-form solution to

$$\min_{\mathbf{z}_{miss(i)} \in \mathbb{R}^{|miss(i)|}, \mathbf{z}_{obs(i)} = \mathbf{x}_{obs(i)}} d_M(\mathbf{z}, \boldsymbol{\mu}_X | \boldsymbol{\Sigma}_X)$$

with $d_M^2(\mathbf{z}, \boldsymbol{\mu}_X | \boldsymbol{\Sigma}_X) = (\mathbf{z} - \boldsymbol{\mu}_X)^\top \boldsymbol{\Sigma}^{-1} (\mathbf{z} - \boldsymbol{\mu}_X)$ being the squared Mahalanobis distance from \mathbf{z} to $\boldsymbol{\mu}_X$. Minimizing the Mahalanobis distance can be seen as maximizing a centrality measure—the Mahalanobis depth:

$$\max_{\mathbf{z}_{\text{miss}(i)} \in \mathbb{R}^{|\text{miss}(i)|}, \mathbf{z}_{\text{obs}(i)} = \mathbf{x}_{\text{obs}(i)}} D_n^M(\mathbf{z} | \mathbf{X})$$

where the Mahalanobis depth of $\mathbf{x} \in \mathbb{R}^d$ w.r.t. \mathbf{X} is defined as follows.

Definition 1. (Mahalanobis, 1936) $D^M(\mathbf{x} | \mathbf{X}) = 1 + (\mathbf{x} - \boldsymbol{\mu}_X)^\top \boldsymbol{\Sigma}_X^{-1} (\mathbf{x} - \boldsymbol{\mu}_X)^{-\frac{1}{2}}$, where $\boldsymbol{\mu}_X$ and $\boldsymbol{\Sigma}_X$ are the location and shape parameters of \mathbf{X} .

In its empirical version $D_n^M(\mathbf{x} | \mathbf{X})$, the parameters are replaced by their estimates.

The Mahalanobis depth is the simplest instance of a statistical depth function. We now generalize the iterative imputation algorithm to other depths.

2.2 Imputation by depth maximization

2.2.1 Definition of data depth

Definition 2. (Zuo and Serfling, 2000a) A bounded non-negative mapping $D(\cdot | \mathbf{X})$ from \mathbb{R}^d to \mathbb{R} is called a statistical depth function if it is (P1) affine invariant, i.e., $D(\mathbf{x} | \mathbf{X}) = D(\mathbf{A}\mathbf{x} + \mathbf{b} | \mathbf{A}\mathbf{X} + \mathbf{b})$ for any invertible $d \times d$ matrix \mathbf{A} and any $\mathbf{b} \in \mathbb{R}^d$; (P2) maximal at the symmetry center, i.e., $D(\mathbf{c} | \mathbf{X}) = \sup_{\mathbf{x} \in \mathbb{R}^d} D(\mathbf{x} | \mathbf{X})$ for any \mathbf{X} halfspace symmetric around \mathbf{c} (A random vector \mathbf{X} having distribution P_X is said to be halfspace symmetric around (a center) $\mathbf{c} \in \mathbb{R}^d$ if $P_X(H) \geq \frac{1}{2}$ for every halfspace H containing \mathbf{c} .); (P3) monotone w.r.t. the deepest point, i.e., for any \mathbf{X} having \mathbf{c} as a deepest point, $D(\mathbf{x} | \mathbf{X}) \leq D(a\mathbf{c} + (1-a)\mathbf{x} | \mathbf{X})$ for any $\mathbf{x} \in \mathbb{R}^d$ and $a \in [0, 1]$; (P4) vanishing at infinity, i.e., $\lim_{\|\mathbf{x}\| \rightarrow 0} D(\mathbf{x} | \mathbf{X}) = 0$.

Additionally, we require (P5) *quasiconcavity* of $D(\cdot | \mathbf{X})$, upper-level sets (or depth-trimmed regions) $D_a(\mathbf{X}) = \{\mathbf{x} \in \mathbb{R}^d : D(\mathbf{x} | \mathbf{X}) \geq a\}$ to be convex, a useful property for optimization. We denote $D_n(\mathbf{X})$ the corresponding empirical depth. See also Zuo and Serfling (2000b) for a reference on depth contours.

2.2.2 Imputation by depth maximization

We suggest a unified framework to impute missing values by depth maximization, which extends iterative regression imputation. More precisely, consider the following iterative scheme: (1) initialize missing values arbitrarily using unconditional mean imputation; (2) impute a point \mathbf{x} containing missing coordinates with the point \mathbf{y} maximizing data depth conditioned on observed values \mathbf{x}_{obs} :

$$\mathbf{y} = \underset{\mathbf{z}_{\text{miss}} \in \mathbb{R}^{|\text{miss}|}, \mathbf{z}_{\text{obs}} = \mathbf{x}_{\text{obs}}}{\text{argmax}} D_n(\mathbf{z} | \mathbf{X}); \quad (2)$$

(3) iterate until convergence.

The solution of (2) can be non-unique (see Figure 1 in the supplementary materials for an illustration) and the depth value may become zero immediately beyond the convex hull of the

support of the distribution. To avoid these problems, we suggest *imputation by depth* (ID) of an \mathbf{x} which has missing values with $\mathbf{y} = ID(\mathbf{x}, D_n(\cdot|\mathbf{X}))$:

$$ID(\mathbf{x}, D_n(\cdot|\mathbf{X})) = \text{ave}_{\mathbf{u} \in \mathbf{R}^d, \mathbf{u}_{obs}=\mathbf{x}_{obs}} \left\{ \|\mathbf{u} - \mathbf{v}\| \mid \mathbf{v} \in D_{n,a^*}(\mathbf{X}) \right\}, \quad (3)$$

$$\text{with } a^* = \inf_{a \in (0;1)} a \mid D_{n,a}(\mathbf{X}) \cap \{\mathbf{z} \mid \mathbf{z} \in \mathbf{R}^d, \mathbf{z}_{obs} = \mathbf{x}_{obs}\} = \emptyset,$$

where ave is the averaging operator. The imputation by iterative maximization of depth is summarized in Algorithm 1. The complexity of Algorithm 1 is $O(N_s n_{miss} \Omega(D))$. It depends on the data geometry and on the missing values (through the number of outer-loop iterations N_s necessary to achieve s -convergence), the number of points containing missing values n_{miss} , and the depth-specific complexities for solving (3) $\Omega(D)$ are detailed in subsections of Section 3.

Algorithm 1 Single imputation

```

1: function IMPUTE.DEPTH.SINGLE( $\mathbf{X}$ )
2:    $\mathbf{Y} \leftarrow \mathbf{X}$ 
3:    $\boldsymbol{\mu} \leftarrow \hat{\boldsymbol{\mu}}^{(obs)}(\mathbf{X})$  D Calculate mean, ignoring missing values
4:   for  $i = 1 : n$  do
5:     if  $miss(i) \neq \emptyset$  then
6:        $\mathbf{y}_{i,miss(i)} \leftarrow \boldsymbol{\mu}_{miss(i)}$  D Impute with unconditional mean
7:    $I \leftarrow 0$ 
8:   repeat D Iterate until convergence or maximal iteration
9:      $I \leftarrow I + 1$ 
10:     $\mathbf{Z} \leftarrow \mathbf{Y}$ 
11:    for  $i = 1 : n$  do
12:      if  $miss(i) \neq \emptyset$  then  $\Sigma$ 
13:         $\mathbf{y}_i \leftarrow ID(\mathbf{y}_i, D_n(\cdot|\mathbf{Z}))$  D Impute with maximum depth
14:    until  $\max_{i \in \{1, \dots, n\}, j \in \{1, \dots, d\}} |\mathbf{y}_{i,j} - \mathbf{z}_{i,j}| < s$  or  $I = I_{max}$ 
15:    return  $\mathbf{Y}$ 

```

2.2.3 Theoretical properties for elliptical distributions

An elliptical distribution is defined as follows (see Fang et al. (1990), and Liu and Singh (1993) in the data depth context).

Definition 3. A random vector X in \mathbf{R}^d is elliptical if and only if there exists a vector $\boldsymbol{\mu}_X \in \mathbf{R}^d$ and $d \times d$ symmetric and positive semi-definite invertible matrix $\boldsymbol{\Sigma}_X = \boldsymbol{\Lambda}\boldsymbol{\Lambda}^\top$ such that for a random vector U uniformly distributed on the unit sphere \mathbf{S}^{d-1} and a non-negative random variable R , it holds that $X \stackrel{D}{=} \boldsymbol{\mu}_X + R\boldsymbol{\Lambda}U$. We then say that $X \sim E_d(\boldsymbol{\mu}_X, \boldsymbol{\Sigma}_X, F_R)$, where F_R is the cumulative distribution function of the generating variate R .

Theorem 1 shows that for an elliptical distribution, imputation of one point with a quasi-concave uniformly consistent depth converges to the center of the conditional distribution when conditioning on the observed values. Theorem 1 is illustrated in Figure 2 in the supplementary materials.

From Proposition 1 it follows that the minimum of the covariance determinant is unique and the determinant itself decreases at each iteration. Thus, to impute points with missing coordinates one-by-one and iterate until convergence constitutes the block coordinate descent method, which can be proved to numerically converge due to Proposition 2.7.1 from Bertsekas (1999) (as long as Σ_X is invertible).

Further, Theorem 3 states that imputation using the maximum Mahalanobis depth, iterative (multiple-output) regression, and regularized PCA (Josse and Husson, 2012) with $S = d - 1$ dimensions, all converge to the same imputed sample.

Theorem 3. *Suppose that we impute $X = (X_{miss}, X_{obs})$ in \mathbb{R}^d with Y so that $y_i = \operatorname{argmax}_{z: \text{obs}(i)=y_{\text{obs}(i)}} DM(z|Y)$ for each i with $|\text{miss}(i)| > 0$ and $y_i = x_i$ otherwise. Then for each such y_i it also holds that:*

- x_i is imputed with the **conditional mean**:

$$y_{i, \text{miss}(i)} = \mu_{Y \text{ miss}(i)} + \Sigma_{Y \text{ miss}(i), \text{obs}(i)} \Sigma_{Y \text{ obs}(i), \text{obs}(i)}^{-1} (x_{\text{obs}(i)} - \mu_{Y \text{ obs}(i)})$$

which is equivalent to **single- and multiple-output regression**,

- Y is a **stationary point** of $|\Sigma_X(X_{miss})|$: $\frac{\partial |\Sigma_X|}{\partial X_{miss}}(Y_{miss}) = \mathbf{0}$, and
- each missing coordinate j of x_i is imputed with **regularized PCA** as in Josse & Husson (2012) with any $0 < \sigma^2 \leq \lambda_d$ and with $X - \mu_X = U \Lambda^{\frac{1}{2}} V^T$ the singular value decomposition (SVD): $y_{i,j} = \sum_{s=1}^d U_{i,s} \frac{\lambda_s - \sigma^2}{\lambda_s} V_{j,s} + \mu_{Y j}$.

The first point of the theorem sheds light on the connection between imputation by Mahalanobis depth and the iterative regression imputation of Section 2.1. When the Mahalanobis depth is used in Algorithm 1, each x_i with missingness in $\text{miss}(i)$ is imputed by the multivariate conditional mean as in equation (1), and thus lies in the $|d - |\text{miss}(i)||$ -dimensional multiple-output regression subspace of $X_{\cdot, \text{miss}(i)}$ on $X_{\cdot, \text{obs}(i)}$. This subspace is obtained as the intersection of the single-output regression hyperplanes $X_{\cdot j}$ on $X_{\cdot, \{1, \dots, d\} \setminus \{j\}}$ for all $j \in \text{miss}(i)$ corresponding to missing coordinates. The third point strengthens the method as imputation with regularized PCA has proved to be highly efficient in practice due to its sticking to low-rank structure of importance and ignoring noise.

The complexity of imputing a single point with the Mahalanobis depth is $O(nd^2 + d^3)$. Despite its good properties, it is not robust to outliers. However, robust estimates for μ_X and Σ_X can be used, e.g., the minimum covariance determinant ones (MCD, see Rousseeuw and Van Driessen, 1999).

3.2 Zonoid depth

Koshevoy and Mosler (1997) define a zonoid trimmed region, with $a \in (0, 1]$, as

$$D_a^z(X) = \left\{ \int_{\mathbb{R}^d} x g(x) dP_X(x) : g : \mathbb{R}^d \rightarrow [0, \frac{1}{a}] \text{ measurable and } \int_{\mathbb{R}^d} g(x) dP_X(x) = 1 \right\}$$

and for $a = 0$ as $D_0^z(X) = \text{cl} \left(\bigcup_{a \in (0, 1]} D_a^z(X) \right)$, where cl denotes the closure. Its empirical version can be defined as

$$D_{n,a}^z(X) = \left\{ \sum_{i=1}^n \lambda_i x_i : \sum_{i=1}^n \lambda_i = 1, \lambda_i \geq 0, a \lambda_i \leq \frac{1}{n} \forall i \in \{1, \dots, n\} \right\}.$$

Definition 4. (Koshevoy and Mosler, 1997) *The zonoid depth of \mathbf{x} w.r.t. X is defined as*

$$D^z(\mathbf{x}|X) = \begin{cases} \sup_{\mathbf{a} \in \text{conv}(\text{supp}(X))} \mathbf{a}^\top \mathbf{x} & \text{if } \mathbf{x} \in \text{conv}(\text{supp}(X)) \\ 0 & \text{otherwise.} \end{cases}$$

For a comprehensive reference on the zonoid depth, the reader is referred to Mosler (2002).

Imputation of a point \mathbf{x}_i in Algorithm 1 is then performed by a slight modification of the linear programming for computation of zonoid depth with variables γ and $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_n)^\top$:

$$\min \gamma \quad \text{s.t.} \quad \mathbf{X}_{\cdot, \text{obs}(i)}^\top \boldsymbol{\lambda} = \mathbf{x}_{i, \text{obs}(i)}, \boldsymbol{\lambda}^\top \mathbf{1}_n = 1, \gamma \mathbf{1}_n - \boldsymbol{\lambda} \geq \mathbf{0}_n, \boldsymbol{\lambda} \geq \mathbf{0}_n.$$

Here $\mathbf{X}_{\cdot, \text{obs}(i)}$ stands for the completed $n \times |\text{obs}(i)|$ data matrix containing columns corresponding only to non-missing coordinates of \mathbf{x}_i , and $\mathbf{1}_n$ (respectively $\mathbf{0}_n$) is a vector of ones (respectively zeros) of length n . In the implementation, we use the simplex method, which is known for being fast despite its exponential complexity. This implies that, for each point \mathbf{x}_i , imputation is performed by the weighted mean:

$$\mathbf{y}_{i, \text{miss}(i)} = \mathbf{X}_{\cdot, \text{miss}(i)}^\top \boldsymbol{\lambda},$$

the average of the maximum number of equally weighted points. Additional insight on the position of imputed points with respect to the sample can be gained by inspecting the optimal weights λ_i . Zonoid imputation is related to local methods such as k NN imputation, as only some of the weights are positive.

3.3 Tukey depth

Definition 5. (Tukey, 1975) *The Tukey depth of \mathbf{x} w.r.t. X is defined as $D^T(\mathbf{x}|X) = \inf\{P_X(H) : H \text{ a closed halfspace, } \mathbf{x} \in H\}$.*

In the empirical version, the probability is substituted by the portion of X giving $D^T(\mathbf{x}|X) = \min_{\mathbf{u}} \frac{1}{n} \cdot |\{i : \mathbf{x}_i^\top \mathbf{u} \geq \mathbf{x}^\top \mathbf{u}, i = 1, \dots, n\}|$. For more information on Tukey depth see Donoho and Gasko (1992).

With nonparametric imputation by Tukey depth, one can expect that after convergence of Algorithm 1, for each point initially containing missing values, it holds that $\mathbf{y}_i = \arg\max_{\mathbf{z}_{\text{obs}} = \mathbf{x}_{\text{obs}}} \min_{\mathbf{u} \in S^{d-1}} \sum_k \mathbf{y}_k^\top \mathbf{u} \geq \mathbf{z}^\top \mathbf{u}, k \in \{1, \dots, n\}$. Thus, imputation is performed according to the maximin principle based on criteria involving indicator functions, which implies robustness of the solution. Note that as the Tukey depth is not continuous, the searched-for maximum (2) may be non-unique (see Figure 1 (top right) in the supplementary materials), and we impute with the barycenter of the maximizing arguments (3). Due to the combinatorial nature of the Tukey depth, to speed up implementation, we run $2d$ times the Nelder-Mead downhill-simplex algorithm, and take the average over the solutions. The imputation is illustrated in Figure 3 in the supplementary materials.

The Tukey depth can be computed exactly (Dyckerhoff and Mozharovskyi, 2016) with complexity $O(n^{d-1} \log n)$, although to avoid computational burden we also implement its approximation with random directions (Dyckerhoff, 2004) having complexity $O(kn)$, with k denoting the number of random directions. All of the experiments are performed with exactly computed Tukey depth, unless stated otherwise.

3.4 Beyond ellipticity: local depth

Imputation with the so-called “global depth” from Definition 2 may be appropriate in applications even if the data moderately deviate from ellipticity (see Section 4.2.6). However, it can fail when the distribution has non-convex support or several modes. A solution is to use the local depth in Algorithm 1.

Definition 6. (Paindaveine and Bever, 2013) For a depth $D(\cdot|X)$, the β -local depth is defined as $LD^\beta(\cdot, X) : \mathbb{R}^d \rightarrow \mathbb{R}^+ : \mathbf{x} \mapsto LD^\beta(\mathbf{x}, X) = D(\mathbf{x}|X^{\beta, \mathbf{x}})$ with $X^{\beta, \mathbf{x}}$ the conditional distribution of X conditioned on $\bigcap_{a \geq 0, P_Y(D_a(Y)) \geq \beta} D_a(Y)$, where Y has the distribution $P_Y = \frac{1}{2} P_X + \frac{1}{2} P_{2x}$.

The locality level β should be chosen in a data-driven way, for instance by cross-validation. An important advantage of this approach is that any depth satisfying Definition 2 can be plugged in to the local depth. We suggest using the Nelder-Mead algorithm to enable imputation with maximum local depth regardless of the chosen depth notion.

3.5 Dealing with outsiders

A number of depths that exploit the geometry of the data are equal to zero beyond $\text{conv}(X)$, including the zonoid and Tukey depths. Although (3) deals with this situation, for a finite sample it means that points with missing values having the maximal value in at least one of the observed coordinates will never move from the initial imputation because they will become vertices of the $\text{conv}(X)$. For the same reason, other points to be imputed and lying exactly on the $\text{conv}(X)$ will not move much during imputation iterations. As such points are not numerous and would need to move quite substantially to influence imputation quality, we impute them—during the initial iterations—using the spatial depth function (Vardi and Zhang, 2000), which is everywhere non-negative. This resembles the so-called “outsider treatment” introduced by Lange et al. (2014). Another possibility is to extend the depth beyond $\text{conv}(X)$, see e.g., Einmahl et al. (2015) for the Tukey depth.

4 Experimental study

4.1 Choice of competitors

We assess the prediction abilities of Tukey, zonoid, and Mahalanobis depth imputation, and the robust Mahalanobis depth imputation using MCD mean and covariance estimates, with the robustness parameter chosen in an optimal way due to knowledge of the simulation setting. We measure their performance against the competitors: conditional mean imputation based on EM estimates of the mean and covariance matrix; regularized PCA imputation with rank 1 and 2; two nonparametric imputation methods: random forest (using the default implementation in the R-package `missForest`), and k NN imputation choosing k from $\{1, \dots, 15\}$, minimizing the imputation error over 10 validation sets as in Stekhoven and Bühlmann (2012). Mean and oracle (if possible) imputations are used to benchmark the results.

4.2 Simulated data

4.2.1 Elliptical setting with Student- t distribution

We generate 100 points according to an elliptical distribution (Definition 3) with $\boldsymbol{\mu}_2 = (1, 1, 1)^\top$ and the shape $\boldsymbol{\Sigma}_2 = (\hat{1}, 1, 1)^\top, (1, 4, 4)^\top, (1, 4, 8)^\top$, where F is the univariate Student- t distribution ranging in number of degrees of freedom (d.f.) from the Gaussian to the Cauchy: $t = \infty, 10, 5, 3, 2, 1$. For each of the 1000 simulations, we remove 5%, 15% and 25% of values completely at random (MCAR), and compute the median and the median absolute deviation from the median (MAD) of the root mean square error (RMSE) of each imputation method. Table 1 presents the results for 25% missing values. The conclusions with other percentages (see the supplementary materials) are the same, but as expected, performances decrease with increasing percentage of missing data.

Distr.	D^{Tukey}	D^{Zon}	D^{Mah}	$D^{MCD,75}$	EM	regPCA1	regPCA2	kNN	RF	mean	oracle
$t = \infty$	1.675 (0.205)	<i>1.609</i> (0.1893)	1.613 (0.1851)	1.991 (0.291)	1.575 (0.1766)	1.65 (0.1846)	1.613 (0.1856)	1.732 (0.2066)	1.763 (0.2101)	2.053 (0.2345)	1.536 (0.1772)
$t = 10$	1.871 (0.2445)	1.81 (0.2395)	<i>1.801</i> (0.2439)	2.214 (0.3467)	1.755 (0.2379)	1.836 (0.2512)	<i>1.801</i> (0.2433)	1.923 (0.2647)	1.96 (0.2759)	2.292 (0.2936)	1.703 (0.2206)
$t = 5$	2.143 (0.3313)	2.089 (0.3331)	<i>2.079</i> (0.3306)	2.462 (0.4323)	2.026 (0.3144)	2.108 (0.3431)	<i>2.08</i> (0.3307)	2.235 (0.3812)	2.259 (0.3656)	2.612 (0.3896)	1.949 (0.3044)
$t = 3$	2.636 (0.5775)	2.603 (0.5774)	2.62 (0.5745)	2.946 (0.6575)	2.516 (0.5537)	2.593 (0.561)	2.619 (0.5741)	2.757 (0.5874)	2.79 (0.5856)	3.165 (0.6042)	2.384 (0.5214)
$t = 2$	3.563 (1.09)	3.73 (1.236)	3.738 (1.183)	3.989 (1.287)	3.567 (1.146)	3.692 (1.186)	3.738 (1.19)	3.798 (1.133)	3.849 (1.19)	4.341 (1.252)	3.175 (0.9555)
$t = 1$	16.58 (13.71)	19.48 (16.03)	19.64 (16.2)	18.5 (12.4)	18.5 (15.46)	18.22 (15.02)	19.61 (16.1)	17.59 (14.59)	17.48 (14.33)	20.32 (16.36)	13.55 (10.71)

Table 1: Median and MAD of the RMSEs of the imputation for a sample of 100 points drawn from elliptically symmetric Student- t distributions with $\boldsymbol{\mu}_2$ and $\boldsymbol{\Sigma}_2$ with 25% of MCAR values, over 1000 repetitions. Bold values indicate the best results, italics the second best.

As expected, the behavior of the different imputation methods changes with the number of d.f., as does the overall leadership trend. For the Cauchy distribution, robust methods perform best: Mahalanobis depth-based imputation using MCD estimates, followed closely by the one using Tukey depth. For 2 d.f., when the first moment exists but not the second, EM- and Tukey-depth-based imputations perform similarly, with a slight advantage to the Tukey depth in terms of MAD. For larger numbers of d.f., when two first moments exist, EM takes the lead. It is followed by the group of regularized PCA methods, and Mahalanobis- and zonoid-depth-based imputation. Note that the Mahalanobis depth and regularized PCA with rank two perform similarly (the small difference can be explained by numerical precision considerations), see Theorem 3. Both nonparametric methods perform poorly, being “unaware” of the ellipticity of the underlying distribution, but give reasonable results for the Cauchy distribution because of insensitivity to correlation. By default, we present the results obtained with spatial depth for the outsiders. For the Tukey depth, implementation is also available using the extension by Einmahl et al. (2015).

4.2.2 Contaminated elliptical setting

We then modify the above setting by adding 15% of outliers (which do not contain missing values) that stem from the Cauchy distribution with the same parameters $\boldsymbol{\mu}_2$ and $\boldsymbol{\Sigma}_2$. As expected, Table 2 shows that the best RMSEs are obtained by the robust imputation methods: Tukey depth and Mahalanobis depth with MCD estimates. Being restricted to a neighborhood, nonparametric methods often impute based on non-outlying points, and thus perform less well as the preceding

Distr.	D^{luk}	D^{zon}	D^{Mah}	$D_{MCD}^{Mah, \mathbf{75}}$	EM	regPCA1	regPCA2	kNN	RF	mean	oracle
$t \infty$	1.751 (0.2317)	1.86 (0.3181)	1.945 (0.4299)	<i>1.81</i> (0.239)	1.896 (0.3987)	1.958 (0.4495)	1.945 (0.4328)	1.859 (0.2602)	1.86 (0.2332)	2.23 (0.3304)	1.563 (0.1849)
$t 10$	1.942 (0.2976)	2.087 (0.4295)	2.165 (0.5473)	<i>2.022</i> (0.3128)	2.112 (0.5226)	2.196 (0.5729)	2.165 (0.5479)	2.051 (0.3143)	2.047 (0.3043)	2.48 (0.4163)	1.733 (0.2266)
$t 5$	2.178 (0.3556)	2.333 (0.4924)	2.421 (0.6026)	<i>2.231</i> (0.381)	2.376 (0.5715)	2.398 (0.6035)	2.421 (0.5985)	2.315 (0.3809)	2.325 (0.3946)	2.766 (0.528)	1.939 (0.2979)
$t 3$	2.635 (0.6029)	2.864 (0.7819)	2.935 (0.8393)	<i>2.664</i> (0.5877)	2.828 (0.7773)	2.916 (0.8221)	2.93 (0.8384)	2.797 (0.6045)	2.838 (0.6228)	3.34 (0.7721)	2.356 (0.4946)
$t 2$	3.763 (1.17)	4.082 (1.535)	4.136 (1.501)	<i>3.783</i> (1.224)	4.036 (1.518)	4.09 (1.585)	4.14 (1.503)	3.955 (1.265)	4.026 (1.354)	4.623 (1.561)	3.323 (1.04)
$t 1$	17.17 (13.27)	20.43 (15.99)	20.27 (15.91)	16.46 (12.94)	19.01 (15.21)	19.81 (16.15)	20.53 (16.28)	18.96 (14.73)	19.04 (14.62)	21.04 (15.56)	14.44 (11.33)

Table 2: Median and MAD of the RMSEs of the imputation for 100 points drawn from elliptically symmetric Student- t distributions, with μ_2 and Σ_2 contaminated with 15% of outliers, and 25% of MCAR values on non-contaminated data, repeated 1000 times.

group. The rest of the included imputation methods cannot deal with the contaminated data and perform rather poorly.

4.2.3 The MAR setting

We next generate highly correlated Gaussian data by setting $\mu_3 = (1, 1, 1)$ and the covariance matrix to $\Sigma_3 = (1, 1.75, 2)^T, (1.75, 4, 4)^T, (2, 4, 8)^T$. We insert missing values according to the MAR mechanism: the first and third variables are missing depending on the value of the second variable. Figure 3 (left) shows the boxplots of the RMSEs. As we expected, semiparametric




Figure 3: Left: RMSEs for different imputation methods for 100 points drawn from a correlated 3-dimensional Gaussian distribution with μ_3 and Σ_3 with MAR values (see implementation for details), over 1000 repetitions. Right: 1000 points drawn from a 6-dimensional Gaussian distribution with μ_4 and Σ_4 contaminated with 15% of outliers, and 15% of MCAR values on non-contaminated data, over 500 repetitions.

methods (EM, regularized PCA and Mahalanobis depth) perform close to the oracle imputation. The good performance of the rank 1 regularized PCA can be explained by the high correlation between variables. The zonoid depth imputes well despite having no parametric knowledge. Nonparametric methods are unable to capture the correlation, while robust methods “throw away” points possibly containing valuable information.

4.2.4 The low-rank model

We consider as a stress-test an extremely contaminated low-rank model by adding to a 2-dimensional low-rank structure a Cauchy-distributed noise. Generally, while capturing any structure is rather meaningless in this setting (confirmed by the high MADs in Table 3), the performance of the methods is “proportional to the way they ignore” dependency information. For this reason, mean imputation as well as nonparametric methods perform best. The Tukey and zonoid depths perform second best by accounting only for fundamental features of the data. This can be also said about the regularized PCA when keeping the first principal component only. The remaining methods try to reflect the data structure, but are distracted either by the low rank or the heavy-tailed noise.

	D^{Tukey}	D^{Zonoid}	D^{Mah}	$D_{MCD,75}^{MCD}$	EM	regPCA1	regPCA2	kNN	RF	mean
Median RMSE	0.4511	0.4536	0.4795	0.5621	0.4709	0.4533	0.4664	0.4409	0.4444	<i>0.4430</i>
Mad of RMSE	0.3313	0.3411	0.3628	0.4355	0.3595	0.3461	0.3554	0.3302	0.3389	<i>0.3307</i>

Table 3: Medians and MADs of the RMSE for a rank-two model in \mathbb{R}^4 of 50 points with Cauchy noise and 20% of missing values according to MCAR, over 1000 repetitions.

4.2.5 Contamination in higher dimensions

To check the resistance to outliers in higher dimensions, we consider a simulation setting similar to that of Section 4.2.2, in dimension 6, with a normal multivariate distribution with $\boldsymbol{\mu}_4 = (0, \dots, 0)^\top$ and a Toeplitz covariance matrix $\boldsymbol{\Sigma}_4$ (having $\sigma_{i,j} = 2^{-|i-j|}$ as entries). The data are contaminated with 15% of outliers and have 15% of MCAR values on non-contaminated data. The Tukey depth is approximated using 1000 random directions. Figure 3 (right) shows that the Tukey depth imputation has high predictive quality, comparable to that of the random forest imputation even with only 1000 random directions.

4.2.6 Skewed distributions and distributions with non-convex support

First, let us consider only a slight deviation from ellipticity. We simulate 150 points from a skewed normal distribution (Azzalini and Capitanio, 1999), insert 15% MCAR values, and impute them with global (Tukey, zonoid and Mahalanobis) depths and their local versions (see Section 3.4). This is shown in Figure 4. In this setting, both global and local imputation perform similarly.

Further, let us consider an extreme departure from ellipticity with the moon-shaped example from Paindaveine and Bever (2013). We generate 150 bivariate observations from $(X_1, X_2)^\top$ with $X_1 \sim U(-1, 1)$ and $X_2|X_1 = x_1 \sim U(1.5(1 - x_1^2), 2(1 - x_1^2))$, and introduce 15% of MCAR values on X_2 , see Figure 5 (left). Figure 5 (right) shows boxplots of the RMSE for single imputation using local Tukey, zonoid and Mahalanobis depths. If the depth and value of β are properly chosen (this can be achieved by cross-validation), the local-depth imputation considerably outperforms the classical methods as well as the global depth.

4.3 Real data

We validate the proposed methodology on three real data sets taken from the UCI Machine Learning Repository (Dua and Karra Taniskidou, 2017) and on the Cows data set. We thus consider




Figure 4: Left: An example of Tukey depth imputation (pluses). Right: boxplots of RMSEs of the prediction for 150 points drawn from a skewed distribution with 15% MCAR, over 100 repetitions; ld.* stands for the local depth with $\beta = 0.8$.




Figure 5: Left: Comparison of global (crosses) and local (pluses) Tukey depth imputation. Right: boxplots of RMSEs of predictions for 150 points drawn from the moon-shaped distribution with 15% MCAR values in the second coordinate, over 100 repetitions; ld.* stands for the local depth with $\beta = 0.2$.

Banknotes ($n = 100$, $d = 3$), Glass ($n = 76$, $d = 3$), Blood Transfusion ($n = 502$, $d = 3$, [Yeh et al., 2009](#)), and Cows ($n = 3454$, $d = 6$). For details on the experimental design, see the implementation. Figure 6 shows boxplots of the RMSEs for the ten imputation methods considered. The zonoid depth is stable across data sets and provides the best results.

Observations in Banknotes are clustered in two groups, which explains the poor performance of the mean and one-dimensional regularized PCA imputation. The zonoid depth searches for a compromise between local and global features and performs the best. The Tukey depth captures the data geometry, but under-exploits information on points' location. Methods imputing by conditional mean (Mahalanobis depth, EM-based, and regularized PCA imputation) perform similarly and reasonably well while imputing in two-dimensional affine subspaces. The Glass data is challenging as it highly deviates from ellipticity, and part of the data lie sparsely in part

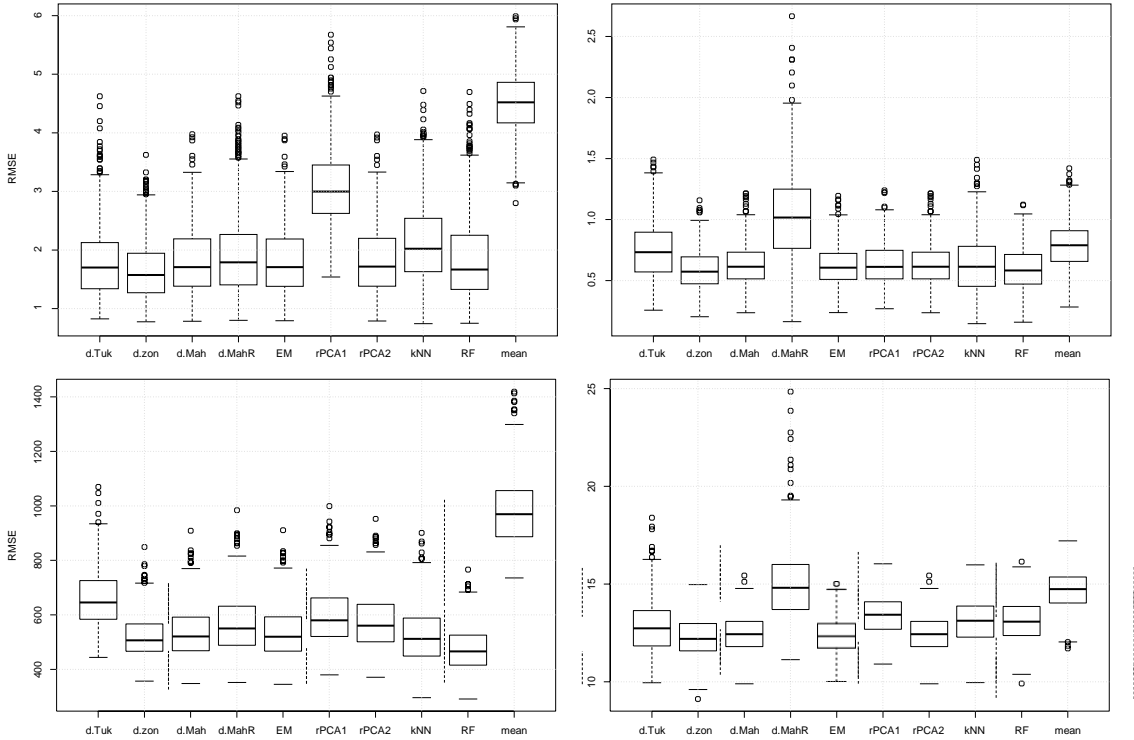


Figure 6: RMSEs for the Banknotes (top, left), Glass (top, right), Cows (bottom, left), and Blood Transfusion (bottom, right) data sets with 15% (5% for Cows) of MCAR values over 500 repetitions.

of the space, but do not seem to be outlying. Thus, the mean, and robust Mahalanobis and Tukey depth imputation perform poorly. Accounting for local geometry, random forest and zonoid depth perform slightly better. For the Cows data, which is larger-dimensional, the best results are obtained with random forest, but followed closely by zonoid depth imputation which reflects the data structure. The Tukey depth with 1000 directions struggles, while the satisfactory results of EM suggest that the data are close to elliptical. The Blood Transfusion data visually resemble a tetrahedron dispersed from one of its vertices. Thus, mean imputation can be substantially improved. Nonparametric methods and rank one regularized PCA perform poorly because they disregard dependency between dimensions. Better imputation is delivered by those capturing correlation: the depth- and EM-based methods.

Table 4 shows the time taken by different imputation methods. Zonoid imputation is very fast, and the approximation scheme by [Dyckerhoff \(2004\)](#) allows for a scalable application of the Tukey depth.

5 Multiple imputation for the elliptical family

When the objective is to predict missing entries as well as possible, single imputation is well suited. When analyzing complete data, it is important to go further, so as to better reflect the uncertainty in predicting missing values. This can be done with multiple imputation (MI) ([Little](#)

Data set	D^{Tuk}	D^{zon}	D^{Mah}	D_{MCD}^{Mah}
Higher dimension (Section 4.2.5) ($n = 1000, d = 6$)	3230*	1210	0.102	1.160
Banknotes ($n = 100, d = 3$)	81.2	0.376	0.010	0.126
Glass ($n = 76, d = 3$)	12.6	0.143	0.008	0.085
Cows ($n = 3454, d = 6$)	4490*	14300	0.212	2.37
Blood transfusion ($n = 502, d = 3$)	26400	51.3	0.46	0.775

Table 4: Median (in seconds, over 35 runs) execution time for depth-based imputation. * indicates approximate Tukey depth with 1000 random directions.

and Rubin, 2002) where several plausible values are generated for each missing entry, leading to several imputed data sets. MI then applies a statistical method to each imputed data set, and aggregates the results for inference. Under the Gaussian assumption, the generation of several imputed data sets is achieved by drawing missing values from the Gaussian conditional distribution of the missing entries, e.g., imputing \mathbf{x}_{miss} by draws from $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ conditional on \mathbf{x}_{obs} , with the mean and covariance matrix estimated by EM. This method is called stochastic EM. The objective is to impute close to the underlying distribution. However, this is not enough to perform *proper* (Little and Rubin, 2002) multiple imputation, since uncertainty in the imputation model’s parameters must also be reflected. This is usually obtained either using a bootstrap or Bayesian approach, see e.g., Schafer (1997), Efron (1994), van Buuren (2012) for more details.

The generic framework of depth-based single imputation developed above allows for multiple imputation to be extended to the more general elliptical framework. We first show how to reflect the uncertainty due to the distribution (Section 5.1), then apply bootstrap to reflect model uncertainty, and state the complete algorithm (Section 5.2).

5.1 Stochastic single depth-based imputation

The extension of stochastic EM to the elliptically symmetric distribution consists in drawing from a conditional distribution that is also elliptical. For this we design a Monte Carlo Markov chain (MCMC), see Figure 7 for an illustration of a single iteration. First, starting with a point with missing values and observed values \mathbf{x}_{obs} , we impute it with $\boldsymbol{\mu}^*$ by maximizing its depth (3), see Figure 7 (right). Then, for each \mathbf{y} with $\mathbf{y}_{obs} = \mathbf{x}_{obs}$ it holds that $D(\mathbf{y} | \mathbf{X}) \leq D(\boldsymbol{\mu}^* | \mathbf{X})$. The cumulative distribution function (with the normalization constant omitted as it is used, exceptionally, for drawing random variables) of the depth of the random vector Y corresponding to \mathbf{y} can be written as

$$F_{D(Y|X)}(\mathbf{y}) = \int_0^{\mathbf{y}} \frac{f_{D(X|X)}(\mathbf{z})}{d_M(\mathbf{z})} \frac{d_M(\mathbf{z})}{d_M(\mathbf{z}) - d_M^2(\boldsymbol{\mu}^* | \mathbf{X})} \frac{\Sigma \Sigma_{|miss(x)|-1}}{d_M(\mathbf{z}) - d_M^2(\boldsymbol{\mu}^* | \mathbf{X})} dz, \quad (4)$$

where $f_{D(X|X)}$ denotes the density of the depth for a random vector X w.r.t. itself, and $d_M(\mathbf{z})$ is the Mahalanobis distance to the center as a function of depth (see the supplementary materials for the derivation of this). For the specific case of the Mahalanobis depth, $d_M(\mathbf{x}) = 1/x - 1$. Then, we draw a quantile Q uniformly on $[0, F_{D(Y|X)}(D(\boldsymbol{\mu}^* | \mathbf{X}))]$ that gives the value of

the depth of \mathbf{y} as $a = F_{D(Y|X)}^{-1}(Q)$, see Figure 7 (left). a defines a depth contour, which is depicted as an ellipsoid in Figure 7 (right). Finally, we draw \mathbf{y} uniformly in the intersection of this contour with the hyperplane of missing coordinates: $\mathbf{y} \in \partial D_a(\mathbf{X}) \cap \{\mathbf{z} \in \mathbb{R}^d \mid z_{obs(x)} = x_{obs}\}$. This is done by drawing \mathbf{u} uniformly on $S^{|\text{miss}(x)|-1}$ and transforming it using a conditional scatter matrix, obtaining $\mathbf{u}^* \in \mathbb{R}^d$, where $\mathbf{u}_{\text{miss}(x)}^* = \mathbf{\Lambda} \mathbf{u}$ (with $\mathbf{\Lambda}(\mathbf{\Lambda})^\top = \mathbf{\Sigma}_{\text{miss}(x), \text{miss}(x)} - \mathbf{\Sigma}_{\text{miss}(x), \text{obs}(x)} \mathbf{\Sigma}_{\text{obs}(x), \text{obs}(x)}^{-1} \mathbf{\Sigma}_{\text{obs}(x), \text{miss}(x)}$) and $\mathbf{u}_{\text{obs}(x)}^* = \mathbf{0}$. Such a \mathbf{u}^* is uniformly distributed on the conditional depth contour. Then \mathbf{x} is imputed as $\mathbf{y} = \boldsymbol{\mu}^* + \beta \mathbf{u}^*$, where β is a scalar obtained as the positive solution of $\boldsymbol{\mu}^* + \beta \mathbf{u}^* \in \partial D_a(\mathbf{X})$ (e.g., the quadratic equation $(\boldsymbol{\mu}^* + \beta \mathbf{u}^* - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu}^* + \beta \mathbf{u}^* - \boldsymbol{\mu}) = d^2$ in the case of the Mahalanobis depth), see Figure 7 (right).

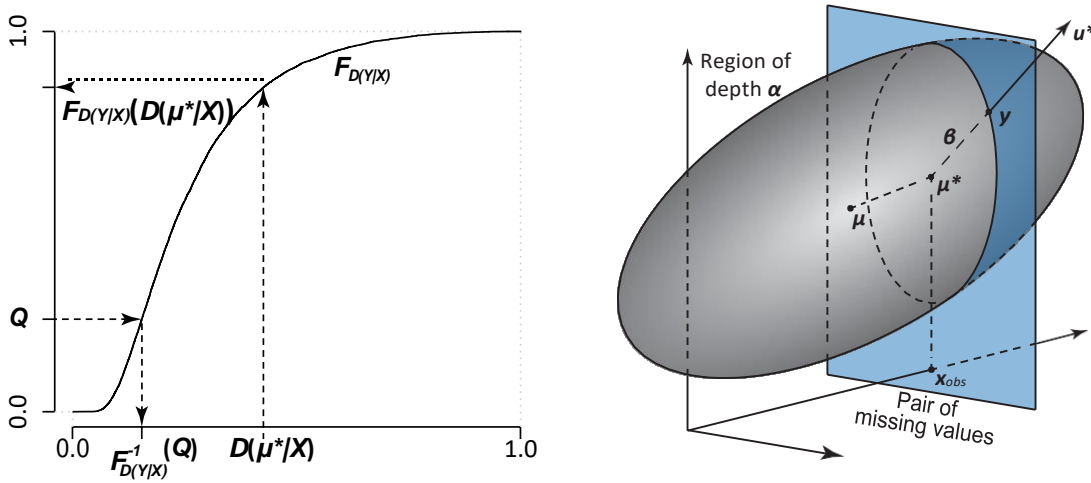


Figure 7: Illustration of an application of (4) to impute by drawing from the conditional distribution of an elliptical distribution. Drawing the depth $D = F_{D(Y|X)}^{-1}(Q)$ via the depth cumulative distribution function $F_{D(Y|X)}$ (left) and locating the corresponding imputed point \mathbf{y} (right).

5.2 Depth-based multiple imputation

We use a bootstrap approach to reflect uncertainty due to the estimation of the underlying semi-parametric model. The depth-based procedure for multiple imputation (called DMI), detailed in Algorithm 2 consists of the following steps: first, a sequence of indices $\mathbf{b} = (b_1, \dots, b_n)$ is drawn from $b_i \sim U(1, \dots, n)$ for $i = 1, \dots, n$, and this sequence is used to obtain new incomplete data set $\mathbf{X}_{\mathbf{b}\cdot} = (\mathbf{x}_{b_1}, \dots, \mathbf{x}_{b_n})$. Then, on each incomplete data set, the stochastic single depth imputation method described in Section 5.1 is applied.

Algorithm 2 Depth-based multiple imputation

```
1: function IMPUTE.DEPTH.MULTIPLE( $\mathbf{X}$ , num.burnin, num.sets)
2:   for  $m = 1 : \text{num.sets}$  do
3:      $\mathbf{Y}^{(m)} \leftarrow \text{IMPUTE.DEPTH.SINGLE}(\mathbf{X})$            D Start MCMC with a single
      imputation
4:      $\mathbf{b} = (b_1, \dots, b_n) = U(1, \dots, n) \overset{\Sigma}{\dots}, U(1, \dots, n)$    D Draw bootstrap
      sequence
5:     for  $k = 1 : (\text{num.burnin} + 1)$  do
6:        $\Sigma \leftarrow \Sigma(\mathbf{Y}^{(m)})$ 
           $\leftarrow \Sigma_{\mathbf{b}, \cdot}$ 
7:       Estimate  $f_{D(\mathbf{X}|\mathbf{X})}$  using  $\mathbf{Y}^{(m)}$ .
8:       for  $i = 1 : n$  do
9:         if  $\text{miss}(i) \neq \emptyset$  then
10:           $\boldsymbol{\mu}^* \leftarrow \text{IMPUTE.DEPTH.SINGLE}(\mathbf{x}_i, \mathbf{Y}_{\mathbf{b}, \cdot}^{(m)})$  D Single-impute point
11:           $\mathbf{u} \leftarrow U(\mathcal{S}^{|\text{miss}(i)|-1})$ 
12:           $\mathbf{u}_{\text{miss}(i)}^* \leftarrow \mathbf{u}\boldsymbol{\Lambda}$            D Calculate random direction
13:           $\mathbf{u}_{\text{obs}(i)}^* \leftarrow \mathbf{0}$ 
14:          Calculate  $F_{D(\mathbf{Y}|\mathbf{X})}$ 
15:           $Q \leftarrow U[0, F_{D(\mathbf{Y}|\mathbf{X})}^{-1}(D(\boldsymbol{\mu}^* | \mathbf{Y}_{\mathbf{b}, \cdot}^{(m)}))] \overset{\Sigma}{\Sigma}$            D Draw depth
16:           $\alpha \leftarrow F_{D(\mathbf{Y}|\mathbf{X})}^{-1}(Q)$ 
17:           $\beta \leftarrow$  positive solution of  $\boldsymbol{\mu}^* + \beta\mathbf{u}^* \in \partial D_{\alpha}(\mathbf{Y}_{\mathbf{b}, \cdot}^{(m)})$ .
18:           $\mathbf{y}_{i, \text{miss}(i)}^{(m)} \leftarrow \boldsymbol{\mu}_{\text{miss}(i)}^* + \beta\mathbf{u}_{\text{miss}(i)}^*$            D Impute one point
19:   return  $\mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(\text{num.sets})}$   $\Sigma$ 
```

5.3 Experiments

5.3.1 Stochastic single depth-based imputation preserves quantiles

We generate 500 points from an elliptical Student- t distribution with 4 degrees of freedom, with $\boldsymbol{\mu}_4 = (-1, -1, -1, -1)^\top$ and $\boldsymbol{\Sigma}_4 = \text{diag}(0.5, 0.5, 1, 1)$, $(0.5, 1, 1, 1)^\top$, $(1, 1, 4, 4)^\top$, $(1, 1, 4, 10)^\top$, adding 30% of MCAR values, and compare the imputation with stochastic EM, stochastic PCA (Josse and Husson, 2012), and stochastic depth imputation (Section 5.1). On the completed data, we calculate quantiles for each variable and compare them with those obtained for the initial complete data. Table 5 shows the medians of the results over 2000 simulations for the first variable; the results are the same for the other variables. The stochastic EM and PCA methods, which generate noise from the normal model, do not lead to accurate quantile estimates. The proposed method gives excellent results with only slight deviations in the tail of the distribution due to difficulties in reflecting the density's shape. Although such an outcome is expected, it considerably broadens the scope of practice in comparison to the deep-rooted Gaussian imputation.

Quantile:	0.5	0.75	0.85	0.9	0.95	0.975	0.99	0.995
complete	-1.0013	-0.4632	-0.1231	0.1446	0.6398	1.2017	2.0661	2.8253
stoch. EM	-1.0008	-0.4225	-0.0649	0.2114	0.6902	1.2022	1.9782	2.6593
stoch. PCA	-0.9999	-0.4248	-0.0650	0.2121	0.7048	1.2398	2.0417	2.7481
depth	-0.9996	-0.4643	-0.1232	0.1491	0.6509	1.2142	2.0827	2.8965

Table 5: Median (over 2000 repetitions) quantiles of the imputed variable X_1 obtained from an elliptical sample of 500 points drawn from the Student- t distribution with 3 d.f. with 30% of MCAR values.

5.3.2 Inference with missing values

We explore the performance of DMI for inference with missing values by estimating coefficients of a regression model. Data are generated according to the following model: $Y = \beta^T(1, X^T)^T + \epsilon$, with $\beta = (0.5, 1, 3)^T$ and $X \sim N(1, 1)^T, (1, 1)^T, (1, 4)^T$, then 30% of MCAR values are introduced. We employ DMI and perform multiple imputation using the R-packages `Amelia` and `mice` under their default settings, generating 20 imputed data sets. For each, we run the regression model to estimate the parameters and their variance, and combine the results according to Rubin’s rules (Little and Rubin, 2002). Here competitors are in a favourable setting as they are based on Gaussian distribution assumptions. We indicate the medians, the coverage of the 95% confidence interval, and the width of this interval, for the estimates of β in Table 6 with a sample size of 500, over 2000 simulations. In addition to the missing data, one difficulty comes from the high correlation (≈ 0.988) between two of the variables.

	β_0			β_1			β_2		
	med	cov	width	med	cov	width	med	cov	width
20 multiply-imputed data sets									
Amelia	0.487	0.931	0.489	1.01	0.941	0.399	2.998	0.929	0.206
mice	0.519	0.984	1.6	1.081	0.98	1.807	2.881	0.982	1.502
regPCA	0.495	0.971	0.853	1.04	0.964	0.751	2.957	0.936	0.334
DMI	0.504	0.971	0.613	0.989	0.979	0.519	3.003	0.97	0.26

Table 6: Medians (med), 95% coverage (cov), and width of the confidence intervals (width) for the regression parameters based on 20 imputed data sets over 2000 repetitions, for a sample of 500 observations from a regression model with 30% MCAR values.

`Amelia` has minor under-coverage problems, and `mice` provides biased coefficients and has large over-coverage issues as it is based on regression imputations that are unstable in the presence of high correlation. DMI, on the other hand, suffers from a slight amount of over-coverage, but in general provides valid inference.

6 Conclusions

The depth imputation framework we propose here fills the gap between global imputation of regression- and PCA-based methods, and the local imputation of methods such as random forest and k NN. It reflects uncertainty in the distribution assumption by imputing data close to the data geometry, is robust in the sense of the distribution and outliers, and still functions with MAR data. When used with the Mahalanobis depth, using data depth as a concept, the link between iterative regression, regularized PCA, and imputation with values that minimize the determinant of the covariance matrix, was established. Our empirical study shows the effectiveness of the suggested methodology for various elliptic distributions and real data. In addition, the method has been naturally extended to multiple imputation for the elliptical family, broadening the scope of existing tools for multiple imputation.

The methodology is general, i.e., any reasonable notion of data depth can be used, which then determines the imputation properties. In the empirical study, the zonoid depth behaves well in general, and for real data in particular. However if robustness is an issue, the Tukey depth may be preferable. The projection depth (Zuo and Serfling, 2000a) is an appropriate choice if only a few points contain missing values in a data set that is substantially outlier-contaminated. This specific case is not included in the article, but imputation based on projection depth is implemented in the associated R package. To reflect multimodality of the data, the suggested framework has been used with localized depths, see e.g. Paindaveine and Bever (2013).

A serious issue with data depths is their computation. Using approximate versions of data depths (which can also be found in the implementation) is a first step to handling larger data sets. Our methodology has been implemented as the R-package `imputeDepth`. Source code of the package and of the experiment-reproducing files can be downloaded from <https://github.com/julierennes/imputeDepth>.

Acknowledgments

The authors thank co-editor in chief Regina Liu, the associate editor, and three anonymous reviewers for their insightful remarks. The authors also gratefully thank Davy Paindaveine and Bernard Delyon for fruitful discussions as well as Germain Van Bever for providing the original code for the local depth, which served as a starting point for its implementation in the R-package `imputeDepth`.

Supplementary materials

Supplementary materials These contain additional figures, results on experiments with other percentages of missing values, as well as proofs.
(NIDD_Supplement.pdf)

R-package and codes for reproducing experiments The archive contains the source code of the accompanying R-package `imputeDepth` as well as scripts for reproducing the experiments and figures of the article. The most up-to-date content of this archive can be found online using the link <https://github.com/julierennes/imputeDepth>.
(NIDD_codes.zip).

References

- Azzalini, A. and Capitanio, A. (1999), ‘Statistical applications of the multivariate skew normal distribution’, *Journal of the Royal Statistical Society: Series B* **61**(3), 579–602.
- Bazovkin, P. and Mosler, K. (2015), ‘A general solution for robust linear programs with distortion risk constraints’, *Annals of Operations Research* **229**(1), 103–120.
- Bertsekas, P. D. (1999), *Nonlinear programming. Second edition*, MIT Press.
- Cascos, I. and Molchanov, I. (2007), ‘Multivariate risks and depth-trimmed regions’, *Finance and Stochastics* **11**(3), 373–397.
- Dempster, A. P., Laird, N. M. and Rubin, D. B. (1977), ‘Maximum likelihood from incomplete data via the em algorithm’, *Journal of the Royal Statistical Society. Series B (Methodological)* **39**(1), 1–38.
- Donoho, D. L. and Gasko, M. (1992), ‘Breakdown properties of location estimates based on halfspace depth and projected outlyingness’, *The Annals of Statistics* **20**(4), 1803–1827.
- Dua, D. and Karra Taniskidou, E. (2017), ‘UCI machine learning repository’.
URL: <http://archive.ics.uci.edu/ml>
- Dyckerhoff, R. (2004), ‘Data depths satisfying the projection property’, *Advances in Statistical Analysis* **88**(2), 163–190.
- Dyckerhoff, R. and Mozharovskyi, P. (2016), ‘Exact computation of the halfspace depth’, *Computational Statistics and Data Analysis* **98**, 19–30.
- Efron, B. (1994), ‘Missing data, imputation, and the bootstrap’, *Journal of the American Statistical Association* **89**(426), 463–475.
- Einmahl, J. H. J., Li, J. and Liu, R. Y. (2015), ‘Bridging centrality and extremity: Refining empirical data depth using extreme value statistics’, *Ann. Statist.* **43**(6), 2738–2765.
- Fang, K., Kotz, S. and Ng, K. (1990), *Symmetric multivariate and related distributions*, Monographs on statistics and applied probability, Chapman and Hall.
- Hastie, T., Mazumder, R., Lee, D. J. and Zadeh, R. (2015), ‘Matrix completion and low-rank svd via fast alternating least squares’, *Journal of Machine Learning Research* **16**, 3367–3402.
- Jörnsten, R. (2004), ‘Clustering and classification based on the $\{L\}$ data depth’, *Journal of Multivariate Analysis* **90**(1), 67–89. Special Issue on Multivariate Methods in Genomic Data Analysis.
- Josse, J. and Husson, F. (2012), ‘Handling missing values in exploratory multivariate data analysis methods’, *Journal de la Société Française de Statistique* **153**(2), 79–99.
- Josse, J. and Reiter, J. P. (2018), ‘Introduction to the special section on missing data’, *Statist. Sci.* **33**(2), 139–141.

- Koshevoy, G. and Mosler, K. (1997), ‘Zonoid trimming for multivariate distributions’, *The Annals of Statistics* **25**(5), 1998–2017.
- Lange, T., Mosler, K. and Mozharovskiy, P. (2014), ‘Fast nonparametric classification based on data depth’, *Statistical Papers* **55**(1), 49–69.
- Little, R. and Rubin, D. (2002), *Statistical analysis with missing data*, Wiley series in probability and mathematical statistics. Probability and mathematical statistics, Wiley.
- Liu, R. Y., Parelius, J. M. and Singh, K. (1999), ‘Multivariate analysis by data depth: descriptive statistics, graphics and inference, (with discussion and a rejoinder by liu and singh)’, *The Annals of Statistics* **27**(3), 783–858.
- Liu, R. Y. and Singh, K. (1993), ‘A quality index based on data depth and multivariate rank tests’, *Journal of the American Statistical Association* **88**(401), 252–260.
- Mahalanobis, P. C. (1936), ‘On the generalised distance in statistics’, **2**(1), 49–55.
- Mosler, K. (2002), *Multivariate Dispersion, Central Regions, and Depth: The Lift Zonoid Approach*, Lecture Notes in Statistics, Springer New York.
- Nagy, S. (2017), ‘Monotonicity properties of spatial depth’, *Statistics and Probability Letters* **129**, 373 – 378.
- Paindaveine, D. and Bever, G. V. (2013), ‘From depth to local depth: A focus on centrality’, *Journal of the American Statistical Association* **108**(503), 1105–1119.
- Rousseeuw, P. J. and Van Driessen, K. (1999), ‘A fast algorithm for the minimum covariance determinant estimator’, *Technometrics* **41**(3), 212–223.
- Schafer, J. (1997), *Analysis of Incomplete Multivariate Data*, Chapman & Hall/CRC Monographs on Statistics & Applied Probability, CRC Press.
- Stekhoven, D. J. and Bühlmann, P. (2012), ‘MissForest – non-parametric missing value imputation for mixed-type data.’, *Bioinformatics* **28**(1), 112–118.
- Troyanskaya, O., Cantor, M., Sherlock, G., Brown, P., Hastie, T., Tibshirani, R., Botstein, D. and Altman, R. B. (2001), ‘Missing value estimation methods for dna microarrays’, *Bioinformatics* **17**(6), 520–525.
- Tukey, J. W. (1975), Mathematics and the Picturing of Data, in R. D. James, ed., ‘International Congress of Mathematicians 1974’, Vol. 2, pp. 523–532.
- Udell, M. and Townsend, A. (2017), ‘Nice latent variable models have log-rank’, *arXiv:1705.07474* .
- van Buuren, S. (2012), *Flexible Imputation of Missing Data (Chapman & Hall/CRC Interdisciplinary Statistics)*, Chapman and Hall/CRC.
- Vardi, Y. and Zhang, C.-H. (2000), ‘The multivariate L1-median and associated data depth’, *Proceedings of the National Academy of Sciences* **97**(4), 1423–1426.

- Yeh, I.-C., Yang, K.-J. and Ting, T.-M. (2009), 'Knowledge discovery on RFM model using bernoulli sequence', *Expert Systems with Applications* **36**(3, Part 2), 5866–5871.
- Zuo, Y. and Serfling, R. (2000a), 'General notions of statistical depth function', *The Annals of Statistics* **28**(2), 461–482.
- Zuo, Y. and Serfling, R. (2000b), 'Structural properties and convergence results for contours of sample statistical depth functions', *The Annals of Statistics* **28**(2), 483–499.

Supplementary Materials to the article “Nonparametric imputation by data depth”

by Pavlo Mozharovskyi, Julie Josse and François Husson

1 Additional figures




Figure 8: A Gaussian sample consisting of 250 points and a hyperplane of two missing coordinates (top, left), and the function $f(z_{miss})$ to be optimized on each single iteration of Algorithm 1, for the smaller rectangle, for Tukey (top, right), zonoid (bottom, left), and Mahalanobis (bottom, right) depth. For the Tukey depth the maximum is not unique, and forms a polygon.

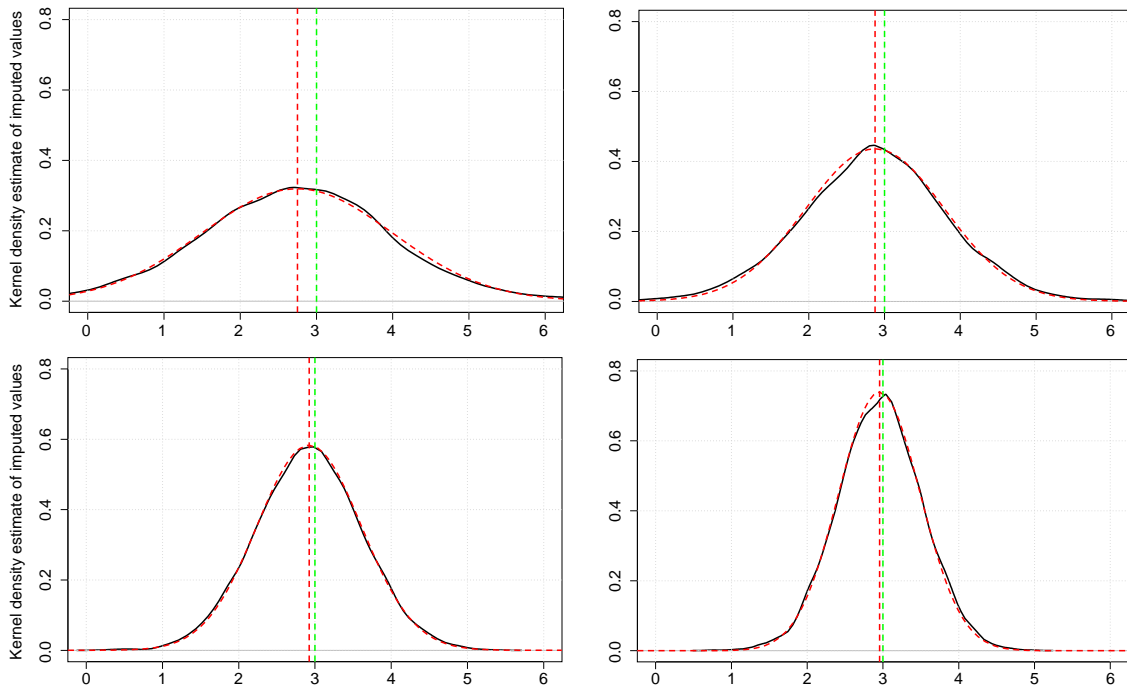


Figure 9: Samples of size 100 (top, left), 200 (top, right), 500 (bottom, left), and 1000 (bottom, right) are drawn from the bivariate Cauchy distribution with the location and scatter parameters $\boldsymbol{\mu}_1$ and $\boldsymbol{\Sigma}_1$ from the introduction. Single point with one missing coordinate is imputed with the Tukey depth. Its kernel density estimate (solid) and the best approximating Gaussian curve (dashed) over 10,000 repetitions are plotted. The population's conditional center given the observed value equals 3.




Figure 10: Illustration of imputation with the Tukey depth. When imputing the point with a missing second coordinate (left), the maximum of the constrained Tukey depth is non-unique (the red line segment), and an average over the optimal arguments (the red point) is used in equation (3) (right).

2 Simulation results with other percentages of missing values

When varying the percentage of missing values, the general trend remains unchanged. The small differences seen can be summarized as follows: with a decreasing percentage of missing values, the difference between EM and Mahalanobis depth imputation (and thus also the rank two PCA one) shrinks, and indeed the latter performs comparably to EM for 5% missingness. For the same percentage and the Cauchy distribution, nonparametric methods (k NN and random forest) perform comparably to the Tukey depth due to a sufficient quantity of available observations and an absence of correlation structure (outliers are generated from Cauchy distribution as well).

Distr.	D^{tuk}	D^{zon}	D^{Man}	D^{Mah}	EM	regPCA1	regPCA2	k NN	RF	mean	oracle
t_{∞}	1.577 (0.2345)	1.547 (0.2128)	1.532 (0.216)	1.537 (0.2199)	1.518 (0.2129)	1.596 (0.2327)	1.532 (0.2159)	1.684 (0.2422)	1.681 (0.2445)	2.058 (0.2774)	1.487 (0.2004)
t_{10}	1.748 (0.287)	1.718 (0.2838)	1.693 (0.2737)	1.709 (0.2827)	1.69 (0.2826)	1.769 (0.3039)	1.692 (0.2741)	1.853 (0.3168)	1.871 (0.3085)	2.275 (0.3757)	1.642 (0.2771)
t_5	1.993 (0.378)	1.971 (0.3602)	1.956 (0.3799)	1.956 (0.361)	1.933 (0.361)	2.017 (0.3759)	1.956 (0.3796)	2.125 (0.4126)	2.134 (0.3976)	2.565 (0.4732)	1.874 (0.3492)
t_3	2.417 (0.5996)	2.434 (0.6032)	2.39 (0.5792)	2.333 (0.5571)	2.362 (0.5808)	2.431 (0.5734)	2.39 (0.5793)	2.55 (0.6154)	2.592 (0.612)	3.045 (0.6943)	2.235 (0.5319)
t_2	3.31 (1.191)	3.373 (1.273)	3.431 (1.314)	3.192 (1.148)	3.366 (1.249)	3.437 (1.343)	3.422 (1.289)	3.538 (1.33)	3.555 (1.321)	4.155 (1.466)	2.986 (1.063)
t_1	13.19 (10.83)	15.13 (12.06)	15.17 (11.74)	13.39 (10.32)	14.86 (11.57)	14.82 (11.64)	15.22 (11.94)	14.09 (11.28)	13.91 (11.06)	16.77 (13.22)	11.17 (8.901)

Table 7: Median and MAD of the RMSEs of the imputation for a sample of 100 points drawn from elliptically symmetric Student- t distributions with μ_2 and Σ_2 having 15% of MCAR values, over 1000 repetitions.

Distr.	$D^{T uk}$	D^{zon}	D^{Mah}	$D_{MCD,75}^{Mah}$	EM	regPCA1	regPCA2	kNN	RF	mean	oracle
$t \infty$	1.656 (0.2523)	1.754 (0.3142)	1.853 (0.4062)	<i>1.671</i> (0.2906)	1.817 (0.3937)	1.855 (0.4158)	1.853 (0.4071)	1.793 (0.2974)	1.762 (0.282)	2.182 (0.3821)	1.514 (0.2121)
$t 10$	1.859 (0.3062)	1.973 (0.4031)	2.048 (0.511)	<i>1.865</i> (0.3917)	2.027 (0.4933)	2.05 (0.5069)	2.044 (0.5119)	1.995 (0.3861)	1.968 (0.3402)	2.45 (0.4778)	1.677 (0.279)
$t 5$	2.09 (0.4275)	2.23 (0.5122)	2.31 (0.6217)	<i>2.109</i> (0.4841)	2.267 (0.6006)	2.348 (0.6504)	2.31 (0.6219)	2.255 (0.4543)	2.233 (0.4476)	2.749 (0.6089)	1.91 (0.3742)
$t 3$	2.507 (0.6389)	2.697 (0.7977)	2.772 (0.8516)	<i>2.541</i> (0.7133)	2.737 (0.8243)	2.791 (0.9306)	2.791 (0.8495)	2.707 (0.6946)	2.699 (0.7254)	3.32 (0.964)	2.239 (0.5497)
$t 2$	3.462 (1.35)	3.68 (1.577)	3.733 (1.6)	<i>3.517</i> (1.39)	3.669 (1.589)	3.807 (1.648)	3.807 (1.601)	3.709 (1.413)	3.762 (1.444)	4.476 (1.794)	3.061 (1.136)
$t 1$	11.81 (9.738)	14.12 (12.09)	14.22 (12.05)	<i>12.34</i> (9.631)	13.78 (11.48)	13.73 (11.01)	14.31 (12.12)	12.58 (10.36)	13.64 (11.44)	15.73 (12.65)	10.37 (8.249)

Table 8: Median and MAD of the RMSEs of the imputation for 100 points drawn from elliptically symmetric Student- t distributions with μ_2 and Σ_2 contaminated with 15% outliers, and 15% of MCAR values on non-contaminated data, over 1000 repetitions.

Distr.	$D^{T uk}$	D^{zon}	D^{Mah}	$D_{MCD,75}^{Mah}$	EM	regPCA1	regPCA2	kNN	RF	mean	oracle
$t \infty$	1.464 (0.3694)	1.454 (0.3713)	1.447 (0.3595)	1.453 (0.3663)	<i>1.449</i> (0.3593)	1.529 (0.401)	1.447 (0.3594)	1.571 (0.3892)	1.581 (0.3946)	2.009 (0.486)	1.404 (0.3399)
$t 10$	1.649 (0.4316)	1.597 (0.4285)	1.57 (0.4163)	1.572 (0.4203)	1.57 (0.4206)	1.665 (0.4502)	1.57 (0.4163)	1.755 (0.4565)	1.754 (0.46)	2.2 (0.5737)	1.529 (0.4278)
$t 5$	1.816 (0.5134)	1.799 (0.5129)	1.757 (0.49)	1.758 (0.4991)	1.757 (0.4899)	1.876 (0.5499)	1.757 (0.4901)	1.955 (0.555)	1.972 (0.5345)	2.402 (0.7318)	1.712 (0.4869)
$t 3$	2.213 (0.7882)	2.184 (0.8159)	2.147 (0.8016)	2.101 (0.7618)	2.139 (0.8)	2.242 (0.7782)	2.147 (0.801)	2.37 (0.8563)	2.343 (0.8357)	2.844 (1.011)	2.054 (0.7649)
$t 2$	2.837 (1.249)	2.919 (1.342)	2.813 (1.309)	2.68 (1.196)	2.8 (1.287)	2.911 (1.311)	2.813 (1.31)	3.03 (1.325)	2.99 (1.331)	3.578 (1.554)	2.529 (1.133)
$t 1$	7.806 (6.351)	8.718 (7.135)	8.911 (7.127)	8.286 (6.602)	8.9 (7.124)	9.118 (7.334)	8.935 (7.137)	8.135 (6.605)	8.138 (6.563)	10.99 (8.952)	6.367 (5.12)

Table 9: Median and MAD of the RMSEs of the imputation for a sample of 100 points drawn from elliptically symmetric Student- t distributions, with μ_2 and Σ_2 having 5% of MCAR values, over 1000 repetitions.

Distr.	$D^{T uk}$	D^{zon}	D^{Mah}	$D_{MCD,75}^{Mah}$	EM	regPCA1	regPCA2	kNN	RF	mean	oracle
$t \infty$	1.552 (0.3693)	1.613 (0.4107)	1.709 (0.4867)	<i>1.553</i> (0.4379)	1.701 (0.4788)	1.769 (0.5248)	1.709 (0.4877)	1.695 (0.407)	1.603 (0.3924)	2.167 (0.5981)	1.406 (0.3171)
$t 10$	1.706 (0.4415)	1.778 (0.5106)	1.874 (0.6104)	<i>1.73</i> (0.4912)	1.861 (0.6032)	1.906 (0.6053)	1.875 (0.6111)	1.884 (0.5182)	1.823 (0.4797)	2.398 (0.7059)	1.564 (0.3938)
$t 5$	1.868 (0.5565)	1.951 (0.5843)	2.038 (0.6859)	<i>1.877</i> (0.5679)	2.027 (0.6806)	2.172 (0.7747)	2.039 (0.6819)	2.077 (0.6256)	1.995 (0.6102)	2.57 (0.8625)	1.698 (0.491)
$t 3$	2.243 (0.8064)	2.348 (0.8694)	2.421 (0.9166)	2.226 (0.8258)	2.42 (0.9345)	2.525 (1.019)	2.421 (0.9237)	2.429 (0.8484)	2.392 (0.8521)	3.05 (1.171)	2.016 (0.7047)
$t 2$	2.902 (1.375)	3.032 (1.498)	3.183 (1.566)	2.933 (1.421)	3.163 (1.558)	3.196 (1.565)	3.188 (1.582)	3.142 (1.472)	3.071 (1.43)	4.073 (2.007)	2.55 (1.129)
$t 1$	7.464 (5.916)	8.487 (6.869)	8.531 (7.081)	8.334 (6.867)	8.5 (6.988)	8.675 (7.261)	8.541 (7.117)	7.958 (6.509)	8.1 (6.922)	10.82 (8.802)	6.245 (4.874)

Table 10: Median and MAD of the RMSEs of the imputation for 100 points drawn from elliptically symmetric Student- t distributions with μ_2 and Σ_2 contaminated with 15% of outliers, with 5% MCAR values on non-contaminated data, over 1000 repetitions.

3 Proofs

Proof of Theorem 1:

Due to the fact that $D_{n,a}(X) \xrightarrow[n \rightarrow \infty]{a.s.} D_a(X)$, in what follows we focus on the population version only. For $X \sim E_a(\mu_X, \Sigma_X, F_R)$ allow the transform $X \mapsto Z = \mathbf{R}\Sigma^{1/2}(X - \mu)$, with \mathbf{R} being a rotation operator such that w.l.o.g. $x \mapsto z$, such that missing values still constitute a $|\text{miss}(x)|$ -dimensional affine space parallel to missing coordinates' axes. Since contours $D_a(Z)$ are concentric spheres centered at the origin, $D_a^*(Z)$ in (3) is of the form $\{\mathbf{v} \mid \mathbf{v} = \mathbf{z}^l + \beta \mathbf{r}, \beta \geq 0\}$ with $\mathbf{z}_{obs}^l(\mathbf{z}) = \mathbf{z}_{obs}$ and $\mathbf{z}_{miss}^l(\mathbf{z}) = \mathbf{0}_{|\text{miss}(x)|}$, and $\mathbf{r} \in \mathbf{S}^{|\text{miss}(x)|-1}$, a unit sphere in the

linear span of $\text{miss}(z)$. Because of the fact that $P\{\mathbf{x} \in \mathbb{R}^d \mid D(\mathbf{x}|X) = \bar{\alpha}\} = 0, \beta = 0$ almost surely and thus z is imputed with $z^j = \mathbf{R}\boldsymbol{\Sigma}^{-1/2}(\mathbf{y} - \boldsymbol{\mu})$. \square

Proof of Theorem 2:

(The challenge here is that the resulting distribution is not elliptical.)

For $X \sim \mathbb{E}_d(\boldsymbol{\mu}_X, \boldsymbol{\Sigma}_X, F_R)$ allow the transform $X \rightarrow Z = \mathbf{R}\boldsymbol{\Sigma}^{-1/2}(X - \boldsymbol{\mu})$, with \mathbf{R} being a rotation operator such that w.l.o.g. $\mathbf{x} = z$, such that $\text{miss}(z) = 1$. (Z has spherical density contours and missing values are in the first coordinate only.)

Let $Z = (Z^j)^\top$ with $Z^j \sim \mathbb{E}_{-1}(\mathbf{0}, \mathbf{I}, F_R)$, where \mathbf{I} is the diagonal matrix. Consider a random vector $U \sim (1-p)Z + pZ^j$ which is a mixture of d - and $(d-1)$ -dimensional spherical distributions. Z corresponds to the imputed missing values—let us now show that this is true. Due to the fact that $D_{n,a}(U) \xrightarrow[n \rightarrow \infty]{a.s.} D_a(U)$, in what follows we focus on the population version only. Missing values constitute one-dimensional affine subspaces parallel to the first coordinate. Thus, due to the affine invariance property (P1 in Definition 2), $D_a(U) \cap \{\mathbf{u} \in \mathbb{R}^d \mid u_1 \geq 0\} = D_a(U) \cap \{\mathbf{u} \in \mathbb{R}^d \mid u_1 \leq 0\} \times (-1, 0, \dots, 0)^\top$. To see this, it suffices to note that the symmetric reflection of U w.r.t. the linear space normal to $(1, 0, \dots, 0)^\top$ equals U . Now, for $\lambda \in \mathbb{R}$ let $\mathbf{u} = (\lambda, \mathbf{u}_2, \dots, \mathbf{u}_d)^\top$ be this one-dimensional affine subspace of missingness for a point. In (3), $\text{ave } D_a(U) \cap \mathbf{u} = \text{ave } D_a(\boldsymbol{\Sigma}U) \cap \{\mathbf{v} \in \mathbb{R}^d \mid v_1 > 0\} \cap \mathbf{u} \cup D_a(U) \cap \{\mathbf{v} \in \mathbb{R}^d \mid v_1 < 0\} \cap \mathbf{u} \cup D_a(U) \cap \{\mathbf{v} \in \mathbb{R}^d \mid v_1 = 0\} \cap \mathbf{u} = D_a(U) \cap \{\mathbf{v} \in \mathbb{R}^d \mid v_1 = 0\} \cap \mathbf{u} = (0, \mathbf{u}_2, \dots, \mathbf{u}_d) = \mathbf{R}\boldsymbol{\Sigma}^{-1/2}(\mathbf{y} - \boldsymbol{\mu})$ (with the obvious correspondence between \mathbf{u} and \mathbf{y}). \square

Proof of Corollary 1:

(P1)–(P5) are obviously satisfied for the Tukey, zonoid and Mahalanobis depths. In Theorem 1, $D_{n,a}(X) \xrightarrow[n \rightarrow \infty]{a.s.} D_a(X)$ is clearly satisfied for the Mahalanobis depth, following Corollary 3.11 by Mosler (2002) for the zonoid depth, and by Theorem 4.2 in Zuo and Serfling (2000b) for the Tukey depth. In Theorem 2, the same logic holds for the Mahalanobis and zonoid depths, but not for the Tukey depth as Z is not elliptical. Using techniques similar those in the proof of Theorem 3.4 in Zuo and Serfling (2000b), one can show that $P\{\mathbf{x} \in \mathbb{R}^d \mid D(\mathbf{x}|Z) = \bar{\alpha}\} = 0$, from which, together with the vanishing at infinity property (P4) and $\sup_{\mathbf{x} \in \mathbb{R}^d} |D_n(\mathbf{x}|Z) - D(\mathbf{x}|Z)| \xrightarrow[n \rightarrow \infty]{a.s.} 0$ (see Donoho and Gasko, 1992), it follows that $D_{n,a}(Z) \xrightarrow[n \rightarrow \infty]{a.s.} D_a(Z)$. \square

Proof of Proposition 1: w.l.o.g. we restrict ourselves to the case $i = 1$. Let Z be X transformed in such a way that it is an $n \times d$ matrix with $\boldsymbol{\mu}_Z = \mathbf{0}$ and $z_{1, \text{miss}(1)} = \boldsymbol{\Sigma}_Z z_{\text{miss}(1), \text{obs}(1)} \boldsymbol{\Sigma}_Z^{-1} z_{\text{obs}(1), \text{obs}(1)}$. Denote the argument $\mathbf{a} = (0, \dots, 0, \mathbf{y}^\top)^\top \in \mathbb{R}^d$. Replacing z_1 with $z_1 + \mathbf{a}$ and subtracting the column-wise average $\frac{\mathbf{a}}{n}$ from each row gives the covariance matrix estimate:

$$\begin{aligned} n\boldsymbol{\Sigma}_Z(\mathbf{y}) &= \mathbf{Z}^\top \mathbf{Z} - z_1 z_1^\top + (z_1 + \mathbf{a})(z_1 + \mathbf{a})^\top - \frac{1}{n} \mathbf{a} \mathbf{a}^\top \\ &= \mathbf{Z}^\top \mathbf{Z} + 2z_1 \mathbf{a}^\top + \frac{n-1}{n} \mathbf{a} \mathbf{a}^\top. \end{aligned}$$

Since $\mathbf{z}_1^\top (\mathbf{Z}^\top \mathbf{Z})^{-1} \mathbf{a} = 0$ due to Mahalanobis orthogonality, by simple algebra for the determinant, one obtains:

$$\begin{aligned}
 n^d |\boldsymbol{\Sigma}_{\mathbf{Z}}(\mathbf{y})| &= |\mathbf{Z}^\top \mathbf{Z} + \frac{1}{2} \mathbf{z}_1 \mathbf{a}^\top \frac{1}{2} + \frac{n-1}{n} \mathbf{a} \mathbf{a}^\top \frac{n-1}{n}| \\
 &= |\mathbf{Z}^\top \mathbf{Z} + \frac{1}{2} \mathbf{z}_1 \mathbf{a}^\top \frac{1}{2} + \frac{n-1}{n} \mathbf{a}^\top \mathbf{Z}^\top \mathbf{Z} + \frac{1}{2} \mathbf{z}_1 \mathbf{a}^\top \frac{1}{2} \frac{n-1}{n} \mathbf{a}^\top \frac{1}{2} \frac{n-1}{n}| \\
 &= |\mathbf{Z}^\top \mathbf{Z}| \cdot |1 + \frac{1}{2} \mathbf{a}^\top (\mathbf{Z}^\top \mathbf{Z})^{-1} \mathbf{z}_1 \frac{1}{2} + \frac{n-1}{n} \mathbf{a}^\top \mathbf{Z}^\top \mathbf{Z}^{-1} \mathbf{a} \frac{n-1}{n} - \frac{\frac{n-1}{n} \mathbf{a}^\top (\mathbf{Z}^\top \mathbf{Z})^{-1} \mathbf{z}_1 \frac{1}{2} \cdot \frac{1}{2} \mathbf{a}^\top (\mathbf{Z}^\top \mathbf{Z})^{-1} \mathbf{a} \frac{n-1}{n}}{1 + \frac{1}{2} \mathbf{a}^\top (\mathbf{Z}^\top \mathbf{Z})^{-1} \mathbf{z}_1 \frac{1}{2}}| \\
 &= |\mathbf{Z}^\top \mathbf{Z}| \cdot |1 + \frac{n+1}{n} \frac{1}{2} \mathbf{a}^\top (\mathbf{Z}^\top \mathbf{Z})^{-1} \mathbf{a} \frac{1}{2}|.
 \end{aligned}$$

Thus $|\boldsymbol{\Sigma}_{\mathbf{Z}}(\mathbf{y})|$ is a quadratic function of \mathbf{y} , which is clearly minimized in $\mathbf{y} = (0, \dots, 0)^\top$. Q

Proof of Theorem 3: The first point can be checked by elementary algebra. The second point follows from the coordinate-wise application of Proposition 1. For the third point, it suffices to prove the single-output regression case. The regularized PCA algorithm will converge if

$$y_{id} = \sum_{s=1}^d u_{is} \sqrt{\lambda_s} v_{ds} = \sum_{s=1}^d u_{is} (\sqrt{\lambda_s} - \frac{\sigma}{\sqrt{\lambda_s}}) v_{ds}$$

for any $\sigma^2 \leq \lambda_d$. W.l.o.g. we prove that

$$\mathbf{y}_d = \boldsymbol{\Sigma}_{d(1,\dots,d-1)} \boldsymbol{\Sigma}_{(1,\dots,d-1)}^{-1} \mathbf{y}_{(1,\dots,d-1)} \iff \sum_{i=1}^d \frac{u_i v_{di}}{\sqrt{\lambda_i}} = 0,$$

denoting $\boldsymbol{\Sigma}(Y)$ simply $\boldsymbol{\Sigma}$ for the centered Y , and an arbitrary point \mathbf{y} . Using matrix algebra,

$$\begin{aligned}
 \sum_{i=1}^d \sqrt{\lambda_i} v_{di} &= \sum_{i=1}^d \frac{v_{di}^2}{\lambda_i} \sum_{i=1}^d \lambda_i v_{di} v_{2i}, \dots, \frac{v_{di} v_{(d-1)i}}{\lambda_i} \sum_{i=1}^d \lambda_i v_{(d-1)i} \\
 &\times \sum_{i=1}^d \frac{1}{\lambda_i} \sqrt{\lambda_i} v_{1i}, \sum_{i=1}^d \frac{1}{\lambda_i} \sqrt{\lambda_i} v_{2i}, \dots, \sum_{i=1}^d \frac{1}{\lambda_i} \sqrt{\lambda_i} v_{(d-1)i}.
 \end{aligned}$$

After reordering the terms, one obtains

$$\sum_{i=1}^d \frac{1}{\lambda_i} \sqrt{\lambda_i} \sum_{j=1}^d \frac{v_{dj}}{\lambda_j} \sum_{k=1}^d v_{ki} v_{kj} = 0.$$

Due to the orthogonality of V , $d^2 - d$ terms from the two outer sum signs are zero. Gathering non-zero terms, i.e., those with $i = j$ only, we have that

$$\sum_{i=1}^d \frac{1}{\lambda_i} \sqrt{\lambda_i} \frac{v_{di}}{\lambda_i} = \sum_{i=1}^d \frac{u_i v_{di}}{\sqrt{\lambda_i}} = 0.$$

$i=1$ i $i=1$ i




Figure 11: Illustration of the derivation of (4).

Q

Derivation of (4): The integrated quantity is the conditional depth density that can be obtained from the joint one by the volume transformation (denoting $d_M(z, \mu)$ the Mahalanobis distance between a point of depth z and μ):

$$f_{D((X|X_{obs}=x_{obs})|X)}(z) = f_{D(X|X)}(z) \cdot C \cdot T_{down} \cdot d_M(z, \mu)^\Sigma \cdot T_{up} \cdot d_M(z, \mu^*)^\Sigma \times T_{angle} \cdot d_M(z, \mu)^\Sigma \cdot d_M(z, \mu^*)^\Sigma$$

Any constant C is ignored as it is unimportant in the drawing. The three terms below correspond to descaling the density to dimension one (downscaling), re-scaling it to the dimension of the missing values (upscaling), and the linear transformation from dimension d to dimension $|miss|$ =number of missing coordinates of a point (angle transformation):

$$\begin{aligned} T_{down} \cdot d_M(z, \mu)^\Sigma &= d_M^{1-d}(z, \mu) = \frac{1}{d_M^{d-1}(z, \mu)} \\ T_{up} \cdot d_M(z, \mu^*)^\Sigma &= d_M^{|\text{miss}(x)|-1}(z, \mu^*) \\ &= \frac{d_M^2(z, \mu) - d_M^2(D(\mu^*|X), \mu)^\Sigma}{d_M^{|\text{miss}(x)|-1}} \\ T_{angle} \cdot d_M(z, \mu)^\Sigma \cdot d_M(z, \mu^*)^\Sigma &= \frac{1}{\sin \theta} = \frac{1}{\frac{d_M(z, \mu^*)}{d_M(z, \mu)}} \\ &= \frac{d_M(z, \mu)}{d_M^2(z, \mu) - d_M^2(D(\mu^*|X), \mu)^\Sigma} \end{aligned}$$

T_{down} and T_{up} are illustrated in Figure 11 (left); for T_{angle} see Figure 11 (right). Setting $d_M(z, \mu) = d_M(z)$ to shorten notation gives (4).