



# BAMITA: Bayesian multiple imputation for tensor arrays

Ziren Jiang<sup>1</sup>, Gen Li <sup>2</sup>, Eric F. Lock <sup>1,\*</sup>

<sup>1</sup>Division of Biostatistics and Health Data Science, School of Public Health, University of Minnesota, 2221 University Avenue SE, Minneapolis, MN 55414, United States

<sup>2</sup>Department of Biostatistics, School of Public Health, University of Michigan, 1415 Washington Heights, M4210, Ann Arbor, MI 48109, United States

\*Corresponding author: Division of Biostatistics and Health Data Science, School of Public Health, University of Minnesota, 2221 University Avenue SE, Minneapolis, MN 55414, United States. Email: elock@umn.edu

## SUMMARY

Data increasingly take the form of a multi-way array, or tensor, in several biomedical domains. Such tensors are often incompletely observed. For example, we are motivated by longitudinal microbiome studies in which several timepoints are missing for several subjects. There is a growing literature on missing data imputation for tensors. However, existing methods give a point estimate for missing values without capturing uncertainty. We propose a multiple imputation approach for tensors in a flexible Bayesian framework, that yields realistic simulated values for missing entries and can propagate uncertainty through subsequent analyses. Our model uses efficient and widely applicable conjugate priors for a CANDECOMP/PARAFAC (CP) factorization, with a separable residual covariance structure. This approach is shown to perform well with respect to both imputation accuracy and uncertainty calibration, for scenarios in which either single entries or entire fibers of the tensor are missing. For two microbiome applications, it is shown to accurately capture uncertainty in the full microbiome profile at missing timepoints and used to infer trends in species diversity for the population. Documented R code to perform our multiple imputation approach is available at <https://github.com/lockEF/MultiwayImputation>.

**KEYWORDS:** Bayesian inference; microbiome data; missing data; multiple imputation; multiway data.

## 1. INTRODUCTION

Multiway data, also known as a multidimensional array or a tensor, has become more prevalent in biomedical signal processing and other diverse fields. Tensors can effectively represent real-world datasets that have multiple dimensions or aspects. However, a common challenge associated with multi-way data is missing data, which can occur either on an entry-wise basis (i.e. certain entries have missing data) or a fiber-wise basis (i.e. entire modes have missing data). For example, in this manuscript we consider microbiome studies in which the abundance of microbial taxa are collected longitudinally over several timepoints, for several subjects. The resulting data can be represented as a 3-way tensor: subjects  $\times$  taxa  $\times$  time points; however, the microbiome profiles (abundance across taxa) are entirely missing at several time points for several of the subjects in both studies, resulting in fiber-wise missingness. Our goal is to enable analysis of the full tensor by imputing missing data, while also accurately capturing uncertainty in the imputed fibers. Accounting for this uncertainty is

particularly critical for the validity of subsequent inferences, such as for population trends in species diversity over time.

The Tucker decomposition (Tucker 1966) and CANDECOMP/PARAFAC (CP) decomposition (Carroll and Chang 1970) are two classic formulations to decompose tensor data, and can be viewed as extensions of the singular value decomposition for a matrix. The Tucker decomposition decomposes a given tensor into a core tensor multiplied by matrices corresponding to each mode. The CP decomposition, which can be viewed as a constrained case of the Tucker decomposition, decomposing the tensor into a sum of rank-1 tensors. Kolda and Bader (2009) gives a very comprehensive review of tensor decompositions and their application in various contexts.

Tensor decomposition methods are useful to uncover the underlying structure in a tensor, and have been used to impute missing elements. Acar et al. (2011) describe imputation approaches based on the CP decomposition, and Chen et al. (2013) describe imputation approaches based on a Tucker decomposition. Several related extensions incorporate penalization or smoothing into estimation of the tensor decomposition for missing data imputation (Tan et al. 2013; Yokota et al. 2016; Wu et al. 2018). An alternative approach, first proposed by Liu et al. (2012), builds upon ideas in low-rank matrix completion (Mazumder et al. 2010) by minimizing the tensor trace norm to transform the problem into a convex optimization problem.

The aforementioned imputation approaches are all deterministic, in that they only produce a single-point estimate with no uncertainty for the imputed values. A Bayesian framework is attractive in this context, because it can accommodate the collective uncertainty of the underlying factors that are combined in a low-rank decomposition. There is an extensive literature on low-rank Bayesian modeling for tensors, particularly in the regression context (Hoff 2015; Guhaniyogi et al. 2017; Wang and Xu 2024). However there is little work on Bayesian approaches to multiple imputation for tensors. Chen et al. (2019) proposed a Bayesian tensor decomposition approach that generalizes the Bayesian matrix factorization model of Salakhutdinov and Mnih (2008), incorporating independent and normally distributed errors. However, their implementation only provides a point estimate with no uncertainty and does not account for correlation in the residual error structure.

To our knowledge no existing tensor imputation methods perform multiple imputation, in which multiple values are simulated for the missing entries to reflect their uncertainty. However, this is often critical for applications, to accurately propagate uncertainty through subsequent analyses after imputation. As shown in our examples, imputing missing values with a point estimate and subsequently treating the values as fixed can drastically underestimate uncertainty and lead to inaccurate inference. Thus, we introduce Bayesian Multiple Imputation for Tensor Arrays (BAMITA), a flexible Bayesian model based on the CP factorization with an efficient MCMC sampling algorithm for multi-way data with missing values. Our approach enables valid inference via simulating from the posterior predictive distribution for missing values, accounting for the uncertainty of the imputed elements. Further, our approach incorporates a separable covariance structure on the error terms to account for any correlation structure that may exist on specific tensor modes. We illustrate the advantages of this approach with respect to both imputation accuracy and uncertainty calibration via simulations, and on data from two longitudinal microbiome studies.

## 2. PRELIMINARIES

### 2.1. Notation and background

Denote  $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$  as an  $N$ -dimensional tensor with the length of each dimension being  $I_1, I_2, \dots, I_N$  respectively. Let  $\mathbf{a} \circ \mathbf{b}$  denote the outer product of two vectors  $\mathbf{a}$  and  $\mathbf{b}$ . Denote the Kronecker product of matrices  $\mathbf{A}$  and  $\mathbf{B}$  as  $\mathbf{A} \otimes \mathbf{B}$  and the *Khatri-Rao* product (which is the columnwise Kronecker product) as  $\mathbf{A} \odot \mathbf{B}$ . Let  $\mathbf{X}_{(n)} \in \mathbb{R}^{I_n \times I_1 I_2 \dots I_{n-1} I_{n+1} \dots I_N}$  denote the mode- $n$  matricization of a tensor  $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ , yielding a matrix with  $I_n$  rows and  $\prod_{i \neq n} I_i$  columns. Denote the Moore-Penrose pseudoinverse of matrix  $\mathbf{A}$  as  $\mathbf{A}^\dagger$ .

For our context, elements of the full tensor  $\mathcal{X}$  may be missing. Let  $\mathcal{M} = \{(i_1, \dots, i_N) : \mathcal{X}_{i_1, \dots, i_N} \text{ is missing}\}$  index missing entries,  $\mathcal{X}_{\text{missing}}$  give the latent missing values in the tensor

$\{\mathcal{X}_{i_1, \dots, i_N} : (i_1, \dots, i_N) \in \mathcal{M}\}$ , and  $\mathcal{X}_{\text{observed}}$  give observed values in the tensor. Our goal is to infer  $\mathcal{X}_{\text{missing}}$  given  $\mathcal{X}_{\text{observed}}$ .

### 2.2. Tensor CP decomposition

For a tensor  $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ , the rank  $R$  CP decomposition (Carroll and Chang 1970) factorizes it into a sum of component rank-one tensors that are the outer product of vectors in each dimension,

$$\mathcal{X} \approx \sum_{r=1}^R \mathbf{u}_r^1 \circ \mathbf{u}_r^2 \circ \dots \circ \mathbf{u}_r^N, \tag{1}$$

where  $\mathbf{u}_r^1, \mathbf{u}_r^2, \dots, \mathbf{u}_r^N$  are vectors with length  $I_1, \dots, I_N$  respectively. For  $i = 1, \dots, N$ , define the matrix

$$\mathbf{U}^i = [\mathbf{u}_1^i \quad \mathbf{u}_2^i \quad \dots \quad \mathbf{u}_R^i]. \tag{2}$$

Then, we can express the CP decomposition in the following concise way:

$$\mathcal{X} \approx \llbracket \mathbf{U}^1, \mathbf{U}^2, \dots, \mathbf{U}^N \rrbracket \equiv \sum_{r=1}^R \mathbf{u}_r^1 \circ \mathbf{u}_r^2 \circ \dots \circ \mathbf{u}_r^N.$$

By definition, the norm of each component is not identifiable, i.e.

$$\llbracket \mathbf{X}^1 \Lambda, \mathbf{X}^2, \dots, \mathbf{X}^N \rrbracket = \llbracket \mathbf{X}^1, \mathbf{X}^2 \Lambda, \dots, \mathbf{X}^N \rrbracket = \llbracket \mathbf{X}^1, \mathbf{X}^2, \dots, \mathbf{X}^N \Lambda \rrbracket$$

where  $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_R)$ . Thus, define the following normalized format where the columns of  $\mathbf{U}^1, \mathbf{U}^2, \dots$ , and  $\mathbf{U}^N$  have norm 1 and  $\lambda$  gives the scale factors:

$$\mathcal{X} \approx \llbracket \lambda; \mathbf{U}^1, \mathbf{U}^2, \dots, \mathbf{U}^N \rrbracket \equiv \sum_{r=1}^R \lambda_r \mathbf{u}_r^1 \circ \mathbf{u}_r^2 \circ \dots \circ \mathbf{u}_r^N. \tag{3}$$

We can also write express the factorization in the matrixized form

$$\mathbf{X}_{(i)} = \mathbf{U}^i \Lambda (\mathbf{U}^N \circ \dots \circ \mathbf{U}^{i+1} \circ \mathbf{U}^{i-1} \circ \dots \circ \mathbf{U}^1)^T \tag{4}$$

which will be used in our following derivations.

### 2.3. ALS and EM imputation

The alternating least square (ALS) method aims to compute a CP decomposition with  $R$  components  $\hat{\mathcal{X}} = \sum_{r=1}^R \mathbf{u}_r^1 \circ \mathbf{u}_r^2 \circ \dots \circ \mathbf{u}_r^N$  that best approximates the target tensor  $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$  via minimizing the sum of squared residuals:

$$\min_{\hat{\mathcal{X}}} \|\mathcal{X} - \hat{\mathcal{X}}\|^2.$$

The ALS algorithm iteratively achieves this by solving the least square problem of one component ( $\hat{\mathbf{U}}^1$ , for example) with the other components fixed, which is

$$\begin{aligned} \hat{\mathbf{U}}^1 &= \underset{\mathbf{U}^1}{\text{argmin}} \|\mathcal{X} - \llbracket \lambda; \mathbf{U}^1, \mathbf{U}^2, \dots, \mathbf{U}^N \rrbracket\|_F \\ &= \underset{\mathbf{U}^1}{\text{argmin}} \|\mathbf{X}_{(1)} - \mathbf{U}^1 \Lambda (\mathbf{U}^N \circ \dots \circ \mathbf{U}^2)^T\|_F. \end{aligned}$$

Solving the least square equation, we have

$$\begin{aligned}\hat{\mathbf{U}}^1 &= \mathbf{X}_{(1)}[(\mathbf{U}^N \odot \dots \odot \mathbf{U}^2)^T]^\dagger \\ &= \mathbf{X}_{(1)}(\mathbf{U}^N \odot \dots \odot \mathbf{U}^2)[(\mathbf{U}^N)^T \mathbf{U}^N * \dots * (\mathbf{U}^2)^T \mathbf{U}^2]^\dagger.\end{aligned}$$

Where  $*$  is the Hadamard product which is the elementwise matrix product for two matrices with the same dimensions. We can normalize each column of  $\hat{\mathbf{U}}^1 = [\hat{\mathbf{u}}_1^1 \ \hat{\mathbf{u}}_2^1 \ \dots \ \hat{\mathbf{u}}_R^1]$  to calculate  $\mathbf{U}^1$  and  $\Lambda$  where  $\lambda_r = \|\hat{\mathbf{u}}_r^1\|$  and  $\mathbf{u}_r^1 = \hat{\mathbf{u}}_r^1/\lambda_r$ . Then, we fix  $\mathbf{U}^1, \mathbf{U}^3, \dots, \mathbf{U}^N$  to compute  $\mathbf{U}^2$  in a similar way. In this way,  $\mathbf{U}^1, \dots, \mathbf{U}^N$  are iteratively updated until the algorithm converges (i.e. the calculated matrix is very close to the former matrix) or it attains the maximum number of iterations.

A common frequentist approach to impute missing data for tensors is an expectation-maximization (EM) algorithm, in which missing entries are iteratively updated via ALS or a similar procedure (Acar et al. 2011). Missing data are first initialized, e.g. as  $\mathbf{X}_{i_1, \dots, i_N} = 0$  for  $(i_1, \dots, i_N) \in \mathcal{M}$  if the tensor is centered. Then, the solution to (2.3) is determined for the full tensor  $\mathbf{X}$ , the missing entries are updated as their corresponding elements in  $\hat{\mathbf{X}}$ , and the process is repeated until convergence.

### 3. BAYESIAN TENSOR IMPUTATION ALGORITHM

In this section, we describe BAMITA, a Bayesian approach for the CP decomposition under different assumptions on the variance structure. We first derive the Bayesian MCMC sampling procedure under the assumption that the error (noise) term is independent and normally distributed. We then introduce another model where the error term is assumed to be normally distributed with a separable covariance structure.

#### 3.1. Independent error

##### 3.1.1. The model

First, follow the notation in Section 2, assume the following model based in the CP decomposition with rank  $R$ :

$$\mathbf{X} = \llbracket \mathbf{U}^1, \mathbf{U}^2, \dots, \mathbf{U}^N \rrbracket + \mathcal{E}, \quad (5)$$

where  $\mathbf{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ ,  $\mathbf{U}^1, \dots, \mathbf{U}^N$  defined in (2) with  $R$  components (i.e. rank  $R$ ), and  $\mathcal{E}$  is a tensor of the error terms with the same dimensions. In the rest of this section, we treat  $R$  as fixed. In practice, the number of components  $R$  can be selected through cross-validation. Here, we assume the error terms are independent and identically distributed from a normal distribution with mean 0 and variance  $\sigma^2$ . As our goal is to have a non-informative prior that can be used in a variety of data situations, we consider an (improper) conjugate Jeffreys prior for all  $\mathbf{U}^1, \mathbf{U}^2, \dots, \mathbf{U}^N$  and  $\sigma^2$ , with,

$$p(\mathbf{U}^n) \propto 1$$

for  $n = 1, \dots, N$  and

$$p(\sigma^2) \propto \frac{1}{\sigma^2}.$$

Although the prior we impose is improper, the corresponding posterior will be proper as we will demonstrate that all the conditional distributions for the posterior are proper. Moreover, the model is invariant to the relative scale of the  $\mathbf{U}^i$ ; one can scale the individual factors as in (3), but this will not affect posterior inference for the low-rank structure and missing data, which is our primary focus.

### 3.1.2. Gibbs sampler with full data

Here we give full conditionals for the Gibbs sampling procedure for fully observed data. Note that, given  $\mathbf{U}^2, \dots, \mathbf{U}^N$  and  $\sigma^2$ , the model can be expressed as:

$$\mathbf{X}_{(1)} = \mathbf{U}^1 (\mathbf{U}^N \odot \mathbf{U}^{(N-1)} \odot \dots \odot \mathbf{U}^2)^T + \mathbf{E}_{(1)} \tag{6}$$

where  $\mathbf{X}_{(1)} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$  is the mode-1 matricization of the tensor  $\hat{\mathbf{X}} = \llbracket \mathbf{U}^1, \mathbf{U}^2, \dots, \mathbf{U}^N \rrbracket$  and  $\mathbf{E}_{(1)}$  is the mode-1 matricization of the error tensor  $\mathbf{E}$  with the same dimension as  $\mathbf{X}_{(1)}$ . Let  $\mathbf{A}_{(1)} = (\mathbf{U}^N \odot \mathbf{U}^{(N-1)} \odot \dots \odot \mathbf{U}^2)$  where the subscript 1 indicate that it is the Khatri-Rao product without matrix  $\mathbf{U}^1$ . We partition  $\mathbf{X}_{(1)}$ ,  $\mathbf{U}^1$ , and  $\mathbf{E}_{(1)}$  along their row, i.e.

$$\mathbf{X}_{(1)} = [(\mathbf{x}_{1\cdot}^1)^T \quad (\mathbf{x}_{2\cdot}^1)^T \quad \dots \quad (\mathbf{x}_{I_1\cdot}^1)^T]^T$$

and,

$$\mathbf{U}^1 = [(\mathbf{u}_{1\cdot}^1)^T \quad (\mathbf{u}_{2\cdot}^1)^T \quad \dots \quad (\mathbf{u}_{I_1\cdot}^1)^T]^T$$

and,

$$\mathbf{E}_{(1)} = [(\epsilon_{1\cdot}^1)^T \quad (\epsilon_{2\cdot}^1)^T \quad \dots \quad (\epsilon_{I_1\cdot}^1)^T]^T.$$

Since each element in the error term is assumed to be independent and normally distributed, we have the following Bayesian linear regression model

$$\mathbf{x}_i^1 = \mathbf{u}_i^1 \mathbf{A}_{(1)} + \epsilon_i^1 \tag{7}$$

for  $i = 1, \dots, I_1$  where  $\epsilon_i^1$  be i.i.d. normal distributed with mean 0 and variance  $\sigma^2$ . Then, with the Jeffreys priors  $p(\mathbf{u}_i^1) \propto 1$  and  $p(\sigma^2) \propto \frac{1}{\sigma^2}$ , the posterior distribution for  $\mathbf{u}_i^1$  is a normal distribution with mean  $(\mathbf{A}_{(1)}^T \mathbf{A}_{(1)})^{-1} \mathbf{A}_{(1)}^T \mathbf{x}_i^1$  and covariance matrix  $\sigma^2 (\mathbf{A}_{(1)}^T \mathbf{A}_{(1)})^{-1}$ . Therefore, we have the following Gibbs sampling algorithm for fully observed data:

- Given  $\mathbf{U}^2, \dots, \mathbf{U}^N$  and  $\sigma^2$ , draw

$$\mathbf{U}^1 = [(\mathbf{u}_{1\cdot}^1)^T \quad (\mathbf{u}_{2\cdot}^1)^T \quad \dots \quad (\mathbf{u}_{I_1\cdot}^1)^T]^T$$

with

$$\mathbf{u}_i^1 \sim N((\mathbf{A}_{(1)}^T \mathbf{A}_{(1)})^{-1} \mathbf{A}_{(1)}^T \mathbf{x}_i^1, \sigma^2 (\mathbf{A}_{(1)}^T \mathbf{A}_{(1)})^{-1})$$

for  $i = 1, \dots, I_1$  where  $\mathbf{A}_{(1)} = (\mathbf{U}^N \odot \mathbf{U}^{(N-1)} \odot \dots \odot \mathbf{U}^2)$ .

- Draw  $\mathbf{U}^2, \dots, \mathbf{U}^N$  similarly:

$$\mathbf{U}^n = [(\mathbf{u}_{1\cdot}^n)^T \quad (\mathbf{u}_{2\cdot}^n)^T \quad \dots \quad (\mathbf{u}_{I_n\cdot}^n)^T]^T$$

with

$$\mathbf{u}_i^n \sim N((\mathbf{A}_{(n)}^T \mathbf{A}_{(n)})^{-1} \mathbf{A}_{(n)}^T \mathbf{x}_i^n, \sigma^2 (\mathbf{A}_{(n)}^T \mathbf{A}_{(n)})^{-1})$$

for  $i = 1, \dots, I_n$  where  $\mathbf{A}_{(n)} = (\mathbf{U}^N \odot \dots \odot \mathbf{U}^{(n+1)} \odot \mathbf{U}^{(n-1)} \odot \dots \odot \mathbf{U}^1)$ .

- Given  $\mathbf{U}^1, \dots, \mathbf{U}^N$ , draw  $\sigma^2$  via its inverse-gamma full conditional  $IG(\alpha = (I_1 \cdot \dots \cdot I_N)/2, \beta = (\|\mathbf{X} - \hat{\mathbf{X}}\|_F)/2)$  where  $\hat{\mathbf{X}}$  is calculated using the simulated  $\mathbf{U}^1, \dots, \mathbf{U}^N$  from the previous steps, and  $\|\cdot\|_F$  be the Frobenius norm.

### 3.1.3. Missing data imputation

We propose to impute the missing entries  $\mathbf{X}_{i_1, \dots, i_N}$ ,  $(i_1, \dots, i_N) \in \mathcal{M}$  with simulated values from their posterior predictive distribution  $p(\text{Vec}(\mathbf{X}_{\text{missing}}) | \text{Vec}(\mathbf{X}_{\text{observed}}))$  for each MCMC iteration. Note that, if we assume the error terms are independently distributed, the conditional distribution will simply be determined by the posterior distribution for the error variance and low-rank structure for the missing entries. The missing entries will be imputed according to their posterior sampling. Given a partially observed tensor  $\mathbf{X}$  and number of modes  $R$ , our Bayesian multiple imputation algorithm for normal i.i.d error is given in Algorithm 1.

---

#### Algorithm 1 Bayesian multiple imputation with normal i.i.d error

---

- 1: Impute the missing elements as  $\mathbf{X}_{i_1, \dots, i_N} = 0$ ,  $(i_1, \dots, i_N) \in \mathcal{M}$ .
  - 2: Set the initial value of  $(\hat{\mathbf{U}}^1)^0, \dots, (\hat{\mathbf{U}}^N)^0$  as either randomly generated number or the result of the frequentist EM algorithm. Set the initial value of  $(\hat{\sigma}^2)^0 \sim \text{IG}(\alpha = (I_1 \cdot \dots \cdot I_N)/2, \beta = (\|\mathbf{X} - \hat{\mathbf{X}}\|_F)/2)$  where  $\hat{\mathbf{X}}$  is calculated using the initial value of  $(\hat{\mathbf{U}}^1)^0, \dots, (\hat{\mathbf{U}}^N)^0$ .
  - 3: **for**  $r = 1, \dots, B$ -th MCMC iteration **do**
  - 4:   Sample  $(\hat{\mathbf{U}}^1)^r, \dots, (\hat{\mathbf{U}}^N)^r$ , and  $(\hat{\sigma}^2)^r$  with the previous posterior distributions.
  - 5:   Calculate the underlying structure of  $\mathbf{X}$  as  $\tilde{\mathbf{X}}^r = \llbracket (\mathbf{U}^1)^r, \dots, (\mathbf{U}^N)^r \rrbracket$
  - 6:   Calculate  $\hat{\mathbf{X}}^r$  where missing elements  $\mathbf{X}_{i_1, \dots, i_N}$ ,  $(i_1, \dots, i_N) \in \mathcal{M}$  are imputed following  $N(\tilde{\mathbf{X}}_{i_1, \dots, i_N}^r, (\hat{\sigma}^2)^r)$
  - 7: **end for**
  - 8: Impute the missing elements  $\mathbf{X}_{i_1, \dots, i_N}$ ,  $(i_1, \dots, i_N) \in \mathcal{M}$  as mean value of  $\{\tilde{\mathbf{X}}_{i_1, \dots, i_N}^b, \dots, \tilde{\mathbf{X}}_{i_1, \dots, i_N}^B\}$  where  $b$  be the burn-in value of the Gibbs sampler.
- 

In practice, the rank  $R$  can be determined by cross-validation. In Algorithm 2, we describe a general cross-validation procedure for choosing  $R$  using the mean squared error (MSE) of the imputed held-out elements.

---

#### Algorithm 2 Selecting number of the component with cross validation

---

- 1: Randomly divide the observed elements into  $K$  equal folds with index  $\mathcal{F}_1, \dots, \mathcal{F}_K$ .
  - 2: **for**  $r = 1, \dots, R$  number of components **do**
  - 3:   **for**  $k = 1, \dots, K$ -th fold **do**
  - 4:     Hold out the elements in the  $k$ -th fold  $\mathbf{X}_{i_1, \dots, i_N} = 0$ ,  $(i_1, \dots, i_N) \in \mathcal{F}_k$  as missing.
  - 5:     Run the corresponding Bayesian imputation algorithm (Algorithm 1 for independent error) and get the imputed value  $\hat{\mathbf{X}}_{i_1, \dots, i_N}$ ,  $(i_1, \dots, i_N) \in \mathcal{F}_k$  for the held-out data.
  - 6:     Calculate the mean squared error (MSE) for the imputation over the held-out data. 
$$\delta_k = \sum_{(i_1, \dots, i_N) \in \mathcal{F}_k} (\hat{\mathbf{X}}_{i_1, \dots, i_N} - \mathbf{X}_{i_1, \dots, i_N})^2.$$
  - 7:   **end for**
  - 8:   Calculate the average MSE with number of component equals  $r$  as  $\bar{\delta}^r = \frac{1}{K} \sum_{k=1}^K \delta_k$ .
  - 9: **end for**
  - 10: Select the optimal number of components with the smallest average MSE.
- 

## 3.2. Correlated error

### 3.2.1. The model

For many tensor applications, data have a residual correlation along some of the modes that is not efficiently captured by the low-rank decomposition. For example, in our longitudinal infant gut

microbiome data, observations are assumed to be correlated across time intervals for the same participants, and abundances are correlated across the taxa. Therefore, instead of assuming the errors are independent, here we assume the errors are normally distributed but correlated. For simplicity, we assume the error tensor in (5) has the separable covariance structure described by Hoff (2011):

$$\mathcal{E} \sim N_{I_1 \times I_2 \times \dots \times I_N}(\mathbf{0}, \Sigma_1, \dots, \Sigma_N), \tag{8}$$

where the tensor normal distribution  $N_{I_1 \times I_2 \times \dots \times I_N}(\mathbf{0}, \Sigma_1, \dots, \Sigma_N)$  is defined according to its  $n$ -th mode unfolding  $\mathcal{E}_{(n)}$  which follows a matrix normal distribution:

$$\mathcal{E}_{(n)} \sim N_{I_n \times I_{-n}}(\mathbf{0}, \Sigma_n, \Sigma_{-n})$$

where  $I_{-n} = \prod_{i=1, \dots, N; i \neq n} I_i$  and  $\Sigma_{-n} = \Sigma_1 \otimes \Sigma_2 \otimes \dots \Sigma_{n-1} \otimes \Sigma_{n+1} \otimes \dots \Sigma_N$ . The vectorization of the tensor then follows a multivariate normal distribution:

$$\text{Vec}(\mathcal{E}) \sim N_{I_1 \cdot I_2 \dots \cdot I_N}(\mathbf{0}, \Sigma_N \otimes \Sigma_{N-1} \dots \otimes \Sigma_1).$$

Compared with estimating the entire unrestricted covariance matrix  $\text{Cov}(\text{Vec}(\mathcal{E}))$  which can be unrealistic due to its potentially high dimension, using a separable covariance structure provides a more stable and parsimonious way of modeling the correlation along different modes. Moreover, it is also more interpretable for the covariance matrices  $\Sigma_i$  for each mode, e.g. a covariance for time points and a covariance for taxa.

Similar to the independent error, we use a non-informative prior to make our algorithm broadly accommodating, with the uniform prior on each  $\mathbf{U}^1, \dots, \mathbf{U}^N$

$$\mathbf{U}^n \propto 1$$

and inverse-Wishart prior on  $\Sigma_1, \dots, \Sigma_N$ ,

$$P(\Sigma_n) \sim \text{IW}(\text{Diag}(I_n), (I_n + 2)).$$

In practice, the set of inverse-Wishart priors for each dimension depends on the specific scientific context of the application. In our implementation, the first dimension is considered the sample dimension and samples are independent. Thus, we set  $\Sigma_1$  to the identity and use a inverse-Wishart prior on the other dimensions.

### 3.2.2. Bayesian Gibbs sampler for full data

Similar to the independent case, our model facilitates a conjugate MCMC sampling procedure for fully observed data:

- Given  $\mathbf{U}^2, \dots, \mathbf{U}^N$  and  $\Sigma_2, \dots, \Sigma_N$ , sample  $\Sigma_1$  with

$$\Sigma_1 | \mathbf{U}^1, \mathbf{U}^2, \dots, \mathbf{U}^N, \Sigma_2, \dots, \Sigma_N \sim \text{IW}(S_n, I_n + 2 + I_{-n})$$

where  $S_n = \text{Diag}(I_n) + ((\tilde{\mathbf{A}}^T \tilde{\mathbf{A}})^{-1} \tilde{\mathbf{A}}^T \tilde{\mathbf{X}}_{(1)} \tilde{\mathbf{A}} - \tilde{\mathbf{X}}_{(1)})^T ((\tilde{\mathbf{A}}^T \tilde{\mathbf{A}})^{-1} \tilde{\mathbf{A}}^T \tilde{\mathbf{X}}_{(1)} \tilde{\mathbf{A}} - \tilde{\mathbf{X}}_{(1)})$ ,  $\tilde{\mathbf{A}} = \mathbf{A} \Sigma_{-1}^{-1/2}$ , and  $\tilde{\mathbf{X}}_{(1)} = \mathbf{X}_{(1)} \Sigma_{-1}^{-1/2}$ .

- Given  $\Sigma_1$  and  $\mathbf{U}^2, \dots, \mathbf{U}^N$  and  $\Sigma_2, \dots, \Sigma_N$ , sample  $\mathbf{U}^1$  with

$$\mathbf{U}^1 | \mathbf{U}^2, \dots, \mathbf{U}^N, \Sigma_1, \dots, \Sigma_N \sim \text{MatrixNormal}((\tilde{\mathbf{A}}^T \tilde{\mathbf{A}})^{-1} \tilde{\mathbf{A}}^T \tilde{\mathbf{X}}_{(1)}, (\tilde{\mathbf{A}}^T \tilde{\mathbf{A}})^{-1}, \Sigma_1)$$

- Sample  $\mathbf{U}^2, \Sigma_2; \dots; \mathbf{U}^N, \Sigma_N$  in a similar way.

### 3.2.3. Missing data imputation

The multiple imputation algorithm for  $\mathcal{X}_{\text{missing}}$  given  $\mathcal{X}_{\text{observed}}$  proceeds somewhat differently for a separable covariance structure than that under independence in Algorithm 1, because the observed values provide additional information to inform missing elements beyond the low rank signal. Let  $\boldsymbol{\mu}_m$  and  $\boldsymbol{\mu}_o$  be the vectorized mean of the missing and observed entries, respectively, given by the low-rank structure (2). Under the multivariate normal error term with separable covariance structure, the conditional posterior predictive distribution of the missing entries will be

$$p(\text{Vec}(\mathcal{X}_{\text{missing}} | \text{Vec}(\mathcal{X}_{\text{observed}})) \sim N(\boldsymbol{\mu}_{\text{missing}|\text{observed}}, \boldsymbol{\Sigma}_{\text{missing}|\text{observed}}) \quad (9)$$

where

$$\boldsymbol{\mu}_{\text{missing}|\text{observed}} = \boldsymbol{\mu}_m + \boldsymbol{\Sigma}_{12}(\boldsymbol{\Sigma}_{22})^{-1}(\mathcal{X}_{\text{observed}} - \boldsymbol{\mu}_o)$$

and

$$\boldsymbol{\Sigma}_{\text{missing}|\text{observed}} = \boldsymbol{\Sigma}_{11} - \boldsymbol{\Sigma}_{12}(\boldsymbol{\Sigma}_{22})^{-1}\boldsymbol{\Sigma}_{21}.$$

Here,  $\boldsymbol{\Sigma}_{11}$ ,  $\boldsymbol{\Sigma}_{12}$ ,  $\boldsymbol{\Sigma}_{21}$ ,  $\boldsymbol{\Sigma}_{22}$  partition the full covariance matrix of the vectorized tensor according to the missing indices (i.e.  $\boldsymbol{\Sigma}_{11} = \text{Cov}(\boldsymbol{\mu}_m, \boldsymbol{\mu}_m)$  and  $\boldsymbol{\Sigma}_{21} = \text{Cov}(\boldsymbol{\mu}_o, \boldsymbol{\mu}_m)$ , etc).

---

#### Algorithm 3 Bayesian multiple imputation with correlated error

---

- 1: Impute the missing elements in  $\mathcal{X}$  as 0.
  - 2: Set the initial value of  $(\hat{\mathbf{U}}^1)^0, \dots, (\hat{\mathbf{U}}^N)^0$  as either randomly generated number or the result of the frequentist EM algorithm. Set the initial value of  $(\hat{\boldsymbol{\Sigma}}_n)^0$  as standard diagonal matrix with dimension  $I_n$  for  $n = 1, \dots, N$ .
  - 3: **for**  $r = 1, \dots, B$ -th MCMC iteration **do**
  - 4:   Sample  $(\hat{\mathbf{U}}^1)^r, (\hat{\boldsymbol{\Sigma}}_1)^r; \dots; (\hat{\mathbf{U}}^N)^r, (\hat{\boldsymbol{\Sigma}}_N)^r$  with the previous posterior distributions.
  - 5:   Calculate the underlying structure of  $\mathcal{X}$  as  $\tilde{\mathcal{X}}^r = \llbracket (\mathbf{U}^1)^r, \dots, (\mathbf{U}^N)^r \rrbracket$
  - 6:   Partition the mean  $\text{Vec}(\tilde{\mathcal{X}}^r)$  into  $\text{Vec}(\tilde{\mathcal{X}}^r)_o$  and  $\text{Vec}(\tilde{\mathcal{X}}^r)_m$ . Partition  $\tilde{\boldsymbol{\Sigma}}^r = \hat{\boldsymbol{\Sigma}}_N^r \otimes \hat{\boldsymbol{\Sigma}}_{N-1}^r \dots \otimes \hat{\boldsymbol{\Sigma}}_1^r$  into  $\boldsymbol{\Sigma}_{11}^r, \boldsymbol{\Sigma}_{12}^r, \boldsymbol{\Sigma}_{21}^r$ , and  $\boldsymbol{\Sigma}_{22}^r$  according to the corresponding missingness of the elements.
  - 7:   Simulate  $\hat{\mathcal{X}}_{\text{missing}}^r$  according to the posterior predictive distribution described in 9 and get  $\hat{\mathcal{X}}^r$  where the missing entries are imputed with  $\hat{\mathcal{X}}_{\text{missing}}^r$ .
  - 8: **end for**
  - 9: Impute the missing entries  $\mathcal{X}_{i_1 \dots i_N}$  as mean value of  $\{\tilde{\mathcal{X}}_{i_1 \dots i_N}^b, \dots, \tilde{\mathcal{X}}_{i_1 \dots i_N}^B\}$  where  $b$  be the burn-in value of the Gibbs sampler.
- 

Given a observed tensor  $\mathcal{X}_{\text{observed}}$  and number of modes  $N$ , we now describe an algorithm to simulate from the posterior predictive distribution,  $p(\mathcal{X}_{\text{missing}} | \mathcal{X}_{\text{observed}})$ , for missing entries. Our Bayesian multiple imputation algorithm for error term with separable covariance structure is described in Algorithm 3.

Note that, in our implementation, to be consistent with our data application, the first dimension is considered the sample dimension and samples are independent. Thus, we set  $\boldsymbol{\Sigma}_1$  to the identity and use an inverse-Wishart prior on  $\boldsymbol{\Sigma}_2, \dots, \boldsymbol{\Sigma}_N$ . In practice, whether or not the full covariance is modeled for a given dimension can depend on the context of the application.

## 4. SIMULATION

To evaluate the performance of our proposed Bayesian tensor imputation methods, we conducted a series of simulation experiments. Since we propose two Bayesian imputation methods for multiway data, we aim to use the simulation experiments to evaluate:

1. The performance of the Bayesian independent imputation algorithm in terms of rank selection under the cross-validation.
2. The performance of the Bayesian multiple imputation algorithms with independent or correlated error in terms of the imputed data entries.
3. The performance of the Bayesian multiple imputation algorithms with independent or correlated error in terms of fiber-wise imputation and inferring uncertainty for functions of the fibers.

For each simulation condition, we run 100 experiments with two independent MCMC chains. All the results are calculated as the median of the converged experiments of the Bayesian algorithm, where the convergence is evaluated using a composite version of the scale reduction factor defined as:

$$\text{SRF} = \frac{2 \sum_{i=1}^{n_1+n_2} (X_i - \bar{X})^2 / (n_1 + n_2 - 1)}{\sum_{i=1}^{n_1} (X_{1i} - \bar{X}_1)^2 / (n_1 - 1) + \sum_{i=1}^{n_2} (X_{2i} - \bar{X}_2)^2 / (n_2 - 1)}$$

where  $\{X_{11}, X_{12}, \dots, X_{1n_1}\}$  and  $\{X_{21}, X_{22}, \dots, X_{2n_2}\}$  be the posterior samples of the two separate chains and  $\{X_1, X_2, \dots, X_{n_1+n_2}\}$  is the combination of those chains.

#### 4.1. Simulation study 1: rank selection and independent error

For our first study, simulated data  $\mathcal{X}$  of dimension  $I_1 \times I_2 \times I_3$  is generated following (5), where  $[\mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \mathbf{U}^{(3)}]$  has a rank-3 underlying structure and  $\mathcal{E}$  has independent normal error with mean 0 and variance 1. Each  $\mathbf{U}^{(i)}$ ,  $i = 1, 2, 3$  is a matrix of dimension  $I_i \times 3$  whose elements are drawn from a standard normal distribution. To demonstrate the performance of our algorithm in different scales of dimensions, we considered three scenarios:  $(10 \times 10 \times 10)$ ,  $(20 \times 20 \times 20)$ , and a higher-dimensional imbalanced scenario  $(10 \times 100 \times 1000)$ . Different missing patterns and proportions of the tensor elements are also considered. For an entry-wise missing scenario, each element of the tensor is randomly set to be missing with the probability of 0.2, 0.5, or 0.7 (yielding missing proportions of 20%, 50%, or 70%, respectively). For a fiber-wise missing scenario, the elements for the entire fiber of the third mode are randomly set to be missing with the corresponding probability.

We run our algorithms with the assigned number of components (rank) equal to 1, 2, 3, 4, or 5 (the true rank is 3) and select the rank according to cross-validation. The validation set is composed of the elements randomly drawn from the observed elements (i.e. not the elements which are set to be missing) of the simulated tensor data. The rank is selected based on the mean squared error (MSE) over the validation set (25% of all the observed elements). For the fiber-wise missing condition, the validation set is also randomly assigned for the entire fiber to be consistent. We evaluate the performance of our Bayesian independent imputation algorithm based on the mean squared error (MSE) and the coverage rate of the 95% credible interval with the true rank and the selected rank. The upper and lower bounds for the 95% credible interval for each element are calculated using the 0.025 and 0.975 quantiles of the MCMC samples. The MSE and coverage are calculated as median values over all the missing elements.

We compare the performance of our Bayesian multiple imputation with independent error against the frequentist EM algorithm with the true rank described in Section 2.3. Results are presented in Table 1. The MSE for the true rank and the selected rank are very close across conditions, which indicates the reasonable performance of rank selection via cross-validation. In fact, the cross-validation selects the true rank in most of the experiments except for the scenarios of  $10 \times 10 \times 10$  dimension and 70% missingness.

Coverage rates for credible intervals are all approximately 95%, so uncertainty is correctly inferred. Moreover, our Bayesian independent imputation method performs better than the EM algorithm in terms of the MSE in most of the scenarios. When the dimension of the tensor is  $(10 \times 100 \times 1000)$ , the two algorithms have similar performance. This matches our expectation since when the total sample size is large enough (for the element-wise missing), the results solved

**Table 1.** Simulation results for study 1: rank selection and independent error.

Missing pattern	Tensor dimension	Missing proportion	BAMITA independent				EM algorithm	True	Selected	Overall
			True rank		Selected rank		True rank	Converged	Converged	Converged
			MSE	Coverage (%)	MSE	Coverage (%)	MSE	%	%	%
Entry missing	(10 × 10 × 10)	20%	0.399	<b>94.8</b>	<b>0.327</b>	94.7	0.363	99	100	99
		50%	0.566	<b>95.2</b>	<b>0.508</b>	92.6	0.718	96	100	96
		70%	1.323	<b>93.0</b>	<b>0.869</b>	86.0	1.136	90	91	85
	(20 × 20 × 20)	20%	0.269	94.4	<b>0.267</b>	<b>94.5</b>	0.266	99	100	99
		50%	0.306	94.5	<b>0.302</b>	<b>94.7</b>	0.326	94	99	94
		70%	0.357	94.0	<b>0.335</b>	<b>95.0</b>	0.522	76	100	76
	(10 × 100 × 1000)	20%	<b>0.257</b>	94.4	<b>0.257</b>	<b>94.4</b>	0.259	97	97	97
		50%	<b>0.271</b>	94.4	<b>0.271</b>	<b>94.4</b>	0.272	98	98	98
		70%	<b>0.256</b>	94.3	<b>0.256</b>	<b>94.3</b>	0.287	92	92	91
Fiber missing	(10 × 10 × 10)	20%	<b>0.375</b>	<b>95.3</b>	<b>0.375</b>	<b>95.3</b>	0.438	100	100	100
		50%	0.523	<b>93.1</b>	<b>0.478</b>	92.8	0.876	88	89	85
		70%	0.771	<b>87.3</b>	<b>0.641</b>	83.8	1.128	34	43	23
	(20 × 20 × 20)	20%	<b>0.285</b>	94.8	<b>0.285</b>	94.8	0.279	100	100	100
		50%	0.300	94.5	<b>0.298</b>	<b>94.8</b>	0.371	94	100	94
		70%	0.335	90.9	<b>0.335</b>	<b>93.1</b>	0.720	65	81	65
	(10 × 100 × 1000)	20%	<b>0.251</b>	<b>94.4</b>	<b>0.251</b>	<b>94.4</b>	0.265	100	100	100
		50%	<b>0.254</b>	<b>93.2</b>	<b>0.254</b>	<b>93.2</b>	0.346	98	98	97
		70%	0.260	<b>87.6</b>	<b>0.259</b>	<b>87.6</b>	0.570	52	48	47

Results are presented as median values over all the missing elements. The best performance in each setting was marked in boldface.

by the frequentist EM algorithm are generally similar to the posterior mean for the Bayesian model with flat priors and independent error. MCMC convergence for the fixed number of iterations is generally achieved, but less so for scenarios with 70% missing fibers.

#### 4.2. Simulation study 2: tensor imputation and correlated error

The second simulation study evaluates the performance of our Bayesian imputation algorithm with correlated error, in terms of the imputed entries or fibers in the tensor. Similar to the first simulation, we generate  $\mathcal{X}$  according to (5) with  $\llbracket \mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \mathbf{U}^{(3)} \rrbracket$  as rank-3 underlying structure. To better mimic the condition of our application data, we simulate  $\mathcal{X}$  with dimensions  $10 \times 10 \times 10$ ,  $20 \times 20 \times 20$ , and  $65 \times 168 \times 6$ . The error terms  $\mathcal{E}$  are generated with separable covariance structure:

$$\mathcal{E} \sim N_{I_1 \times I_2 \times I_3}(\mathbf{0}, \Sigma_1, (\Sigma_2 \Sigma_2)^T, (\Sigma_3 \Sigma_3^T))$$

where  $\Sigma_i, i = 2, 3$  are the  $I_2 \times I_2$  and  $I_3 \times I_3$  covariance matrix with diagonal elements being 0.9 and other elements randomly set to be either 0.3 or  $-0.3$  with equal probability. Following our applications,  $\Sigma_1$  is set to be the  $I_1 \times I_1$  diagonal matrix (i.e. there is no correlation structure among the first dimension) with diagonal elements being 0.5. The different missing patterns and proportions of the tensor elements are the same as in the first simulation study. In this simulation study the rank is set to be 3 and the performance of each algorithm is directly evaluated over the missing elements.

The missing value is imputed using the posterior mean, and we calculate the mean squared error (MSE) and the coverage rate for the 95% credible interval of the missing elements for the two Bayesian methods. For the frequentist EM method, we calculate the MSE for the missing elements as a comparison. We also compare with the missForest method (Stekhoven and Bühlmann 2012) applied to the matricized data along the first mode, as a general approach that does not assume low-rank tensor structure.

The results can be seen in Table 2. For the 70% fiber-wise missing with a dimension of  $(10 \times 10 \times 10)$ , we do not include the results as only 4% of the total experiments converged. We can see that the Bayesian correlated imputation algorithm outperforms the Bayesian independent imputation algorithm and the frequentist EM algorithm in most of the scenarios in terms of the MSE except for the 70% fiber-wise missing case with dimension  $(65 \times 168 \times 6)$ . The coverage of the correlated algorithm is also comparable to the independent algorithm, and both have coverage close to the nominal rate of 95%.

#### 4.3. Simulation study 3: imputation for function of a fiber

In the third simulation study, we examine whether an arbitrary function of the entire fiber can be reasonably captured when the data is generated with a correlation structure. This is motivated by our data applications where, instead of focusing on the imputed data entries, we are interested in the alpha diversity calculated as a function of the entire mode-2 fiber for each subject and each time point (see section 5 for more detail). We argue that although the MSE of each imputed element is mostly adopted as the metric for evaluating the imputation performance, the “structure” of the imputed data slice (i.e. fiber) is also of interest for downstream analyses of the tensor.

The data generating mechanism is similar to that of the second simulation study except that the covariance matrix  $\Sigma_1$  and  $\Sigma_3$  are now identity matrices with dimensions  $I_2$  and  $I_3$  respectively, and  $\Sigma_2$  is a matrix with diagonal elements being 1 and other non-diagonal elements being 0.15. Under the data-generating mechanism, we do not expect the Bayesian correlated algorithm to outperform the independent algorithm in terms of the point-wise imputation error since the other modes (modes 1 and 3) are uncorrelated. However, with the additional model of the covariance structure for mode 2, our correlated algorithm may better capture the variation of the function of the imputed mode 2 fiber. To evaluate this, for each imputed fiber  $\hat{\mathcal{X}}[i, \cdot, k]$ , we calculate the linear predictor  $\beta_b \hat{\mathcal{X}}[i, \cdot, k]$  with randomly generated coefficient  $\beta_b, b = 1, \dots, 100$ . Then we evaluate the mean

**Table 2.** Simulation results for study 2: tensor imputation with correlated error.

Missing pattern	Tensor dimension	Missing proportion	BAMITA Correlated		BAMITA Indep	missForest	EM	Corr	Inde	Overall	
			MSE Coverage(%)	MSE low-rank Coverage(%)	MSE Coverage(%)	MSE	Algorithm MSE	prop (%)	prop (%)	prop (%)	
Entry missing	(10 × 10 × 10)	20%	0.009	0.001	0.067	0.617	0.076	100	99	99	
			97.4	98.2	94.4		0.076	100	99	99	
		50%	0.033	0.007	0.064	0.914	0.110	100	92	92	
			95.8	96.3	94.8		0.076	100	99	99	
			70%	0.102	0.044	0.113	–	0.415	99	85	84
			93.6	93.7	94.6		0.076	100	99	99	
	(20 × 20 × 20)	20%	0.007	0.001	0.105	0.371	0.105	100	100	100	
			97.2	98.3	94.6		0.076	100	99	99	
		50%	0.033	0.002	0.107	0.561	0.107	100	100	100	
			95.9	97.3	94.8		0.076	100	99	99	
			70%	<b>0.083</b>	0.009	0.117	0.886	0.151	100	82	82
			94.4	95.2	94.7		0.076	100	99	99	
(65 × 168 × 6)	20%	0.050	0.001	0.273	0.364	0.287	95	98	93		
		95.8	97.3	94.3		0.076	100	99	99		
	50%	0.138	0.005	0.275	0.397	0.277	90	93	84		
		95.2	96.5	94.3		0.076	100	99	99		
		70%	0.232	0.013	0.293	0.494	0.300	77	94	72	
		94.6	95.9	94.3		0.076	100	99	99		
Fiber missing	(10 × 10 × 10)	20%	0.025	0.002	0.065	0.655	0.098	98	97	95	
			94.5	98.4	94.7		0.076	100	99	99	
		50%	0.085	0.017	0.083	0.988	0.339	75	64	54	
			92.7	95.0	94.8		0.076	100	99	99	
			70%	–	–	–	–	–	9	16	4
			–	–	–						
	(20 × 20 × 20)	20%	0.027	0.001	0.108	0.387	0.108	100	100	100	
			93.2	98.3	94.8		0.076	100	99	99	
		50%	0.073	0.003	0.112	0.595	0.145	96	88	85	
			91.4	96.9	94.8		0.076	100	99	99	
			70%	0.130	0.020	0.145	0.967	0.507	80	36	32
			90.2	92.4	94.4		0.076	100	99	99	
(65 × 168 × 6)	20%	0.109	0.002	0.286	0.421	0.283	82	99	81		
		92.3	98.1	94.7		0.076	100	99	99		
	50%	0.316	0.082	0.309	0.515	0.437	82	87	71		
		91.9	96.0	94.4		0.076	100	99	99		
		70%	0.798	0.538	0.555	0.904	0.761	59	49	33	
		90.0	90.0	93.1		0.076	100	99	99		

Results are presented as median values over all the missing elements.

MSE and coverage (only for fiber-wise missing) for the linear predictors over the 100 randomly generated coefficients  $\beta_b$ . The results are displayed as “MSE(Fiber)” in Table 3.

The results are shown in Table 3, from which we can see that, although the Bayesian independent algorithm has a relatively good coverage rate for each imputed element, the mean coverage for the random linear combination of each imputed fiber is not ideal, especially in the high-dimension cases. The Bayesian correlated algorithm, on the other hand, has substantially better coverage performance in terms of the imputed fiber.

## 5. INFANT GUT MICROBIOME APPLICATION

The gut houses a rich array of microbial organisms, presenting a diverse landscape. This dynamic ecosystem holds potential as significant indicators for both digestive and broader health conditions. We consider a longitudinal study of the gut microbiome on 52 infants in the neonatal intensive care unit (NICU), where stool samples were collected over the first 3 mo of life (Cong et al. 2017). Using 16s rRNA sequencing technology, we obtained microbiome data, which we aggregated to the genus level, yielding 152 distinct genera. Employing standard preprocessing techniques, we addressed zero values by introducing pseudo counts, transformed the data into compositional profiles, and applied the centered log-ratio (clr) transformation. Data were aggregated every 5 consecutive days, yielding 30 time intervals. Consequently, we obtained a tensor data array with dimensions of  $52 \times 152 \times 30$ . However, due to the unavailability of samples from every infant at every time point, the tensor array exhibits a fiber-wise missing structure, with approximately 71% of the samples being absent. Our objective is to employ BAMITA to address missing values and assess the dynamic diversity in the microbiome over this population. Before applying our algorithm to this dataset, we first check the suitability of the normality assumption and the separable covariance assumption. Results are presented in Section S2, and suggest that the two assumptions are reasonable for these data. The microbiome diversity is often measured using alpha diversity which refers to diversity on a local scale (Thukral 2017). Here, we compute diversity using the Shannon-Wiener index (Shannon 1948). For a given fiber of the data  $X_{i,k}$ , the Shannon alpha diversity is calculated as  $-\sum_{j=1}^{152} p_j \log(p_j)$  where  $p_j = \frac{\exp(\text{Clr}X_{i,j,k})}{\sum_{j=1}^{152} \exp(\text{Clr}X_{i,j,k})}$  is the proportion of element  $j$  in the sample.

We impute the missing fibers of the microbiome tensor array with the two proposed Bayesian algorithms and the EM algorithm approach. For the Bayesian multiple imputation algorithm with separable covariance structure error, we assume independence in the first dimension (the 52 infants) but account for correlation across time and genera. The performance of the different algorithms is evaluated through cross-validation. For each of 200 simulation experiments, we randomly hold an additional 25% of the observed fibers out and calculate the mean squared error of the imputed values and the true observed value. Shannon diversity is also computed for the missing fibers at each MCMC iteration, and they are compared to the observed diversity measures for the validation set. For the Bayesian imputation with covariance structure (BAMITA Correlated), we also evaluate the MSE for the estimated low-rank structure, i.e.  $[(\mathbf{U}^1), \dots, (\mathbf{U}^N)]$ , which represent the imputation without adjustment of covariance structure.

The results are summarized in Table 4. The low-rank Bayesian imputation approach with correlation structure performs substantially better than other approaches with respect to MSE for both imputed values in the fiber and Shannon diversity. This suggests that the data have a strong low-rank structure, with substantial correlation in the residual covariance. Moreover, while coverage rates are appropriate for the fiber-wise entries under both models, coverage for Shannon diversity is much higher for the correlated model. This illustrates the notion that accurately inferring uncertainty in the marginal distribution for the entries of an array does not imply uncertainty will be accurately inferred for multivariate functions that are used in downstream analysis. The MSE for the low-rank structure in the correlated algorithm is similar to the MSE of the independent algorithm, indicating that adjusting the covariance structure helps the imputation performance.

**Table 3.** Simulation results for study 3: imputation for a function of a fiber.

Missing pattern	Tensor dimension	Missing proportion	BAMITA Correlated		BAMITA Independent		EM Algorithm True Rank	Converged Proportion	Converged Proportion	Converged Proportion	
			MSE (Imputation) Coverage(%)	MSE (Fiber) Coverage(%)	MSE (Imputation) Coverage(%)	MSE (Fiber) Coverage(%)	MSE (Imputation) Coverage(%)	Correlated	Independent	Overall	
Entry missing ( $10 \times 10 \times 10$ )	20%	50%	0.409	-	<b>0.314</b>	-	0.322	100	100	100	
			88.0	<b>78.0</b>	<b>95.0</b>	53.0	-	-	-	-	
			0.396	-	<b>0.327</b>	-	0.423	100	97	97	
	50%	70%	91.3	<b>90.0</b>	<b>95.0</b>	72.0	-	-	-	-	
			0.596	-	<b>0.466</b>	-	0.776	98	99	97	
			90.7	<b>89</b>	<b>94.9</b>	83.5	-	-	-	-	
	$(20 \times 20 \times 20)$	20%	50%	0.266	-	<b>0.260</b>	-	0.265	100	100	100
				92.2	<b>91.2</b>	<b>94.8</b>	36.2	-	-	-	-
				0.284	-	<b>0.274</b>	-	0.282	100	100	100
70%	50%	92.0	<b>91.2</b>	<b>94.9</b>	53.2	-	-	-	-		
		0.278	-	<b>0.272</b>	-	0.326	100	86	86		
		93.0	<b>91.0</b>	<b>94.9</b>	62.3	-	-	-	-		
$(65 \times 168 \times 6)$	20%	50%	0.287	-	<b>0.260</b>	-	0.270	100	99	99	
			91.9	<b>91.8</b>	<b>94.8</b>	13.8	-	-	-	-	
			<b>0.268</b>	-	0.271	-	0.299	100	98	98	
70%	50%	92.9	<b>92.6</b>	<b>94.8</b>	22.3	-	-	-	-		
		<b>0.250</b>	-	0.278	-	0.305	100	96	96		
		93.6	<b>92.1</b>	<b>94.7</b>	26.0	-	-	-	-		

*continued*

**Table 3. Continued**

Missing pattern	Tensor dimension	Missing proportion	BAMITA Correlated		BAMITA Independent		EM Algorithm True Rank	Converged Proportion	Converged Proportion	Converged Proportion
			MSE (Imputation) Coverage(%)	MSE (Fiber) Coverage(%)	MSE (Imputation) Coverage(%)	MSE (Fiber) Coverage(%)	MSE (Imputation) -	Correlated -	Independent -	Overall -
Fiber missing	(10 × 10 × 10)	20%	0.404 90.0	0.767 <b>90.2</b>	<b>0.311</b> <b>94.8</b>	<b>0.617</b> 83.0	0.332 -	99	100	99
		50%	0.875 89.5	1.122 <b>87.1</b>	<b>0.514</b> <b>94.6</b>	<b>0.915</b> 82.8	0.612 -	78	92	73
		70%	<b>2.609</b> 86.6	5.153 83.8	4.041 <b>93.2</b>	<b>3.510</b> <b>84.2</b>	0.964 -	9	39	7
	(20 × 20 × 20)	20%	0.285 93.0	0.668 <b>90.3</b>	<b>0.266</b> <b>94.9</b>	<b>0.616</b> 67.9	0.273 -	100	100	100
		50%	0.320 92.3	0.766 <b>87.4</b>	<b>0.285</b> <b>94.8</b>	<b>0.701</b> 68.4	0.367 -	100	91	91
		70%	0.430 91.8	0.837 <b>86.1</b>	<b>0.369</b> <b>94.6</b>	<b>0.769</b> 69.6	0.645 -	93	61	59
	(65 × 168 × 6)	20%	0.261 94.6	<b>0.940</b> <b>80.0</b>	<b>0.259</b> <b>94.8</b>	0.947 29.4	0.288 -	100	99	99
		50%	0.465 92.9	0.990 <b>76.3</b>	<b>0.337</b> <b>94.6</b>	<b>0.958</b> 31.0	0.431 -	89	85	77
		70%	0.869 90.7	1.005 <b>71.8</b>	<b>0.540</b> <b>93.2</b>	<b>0.979</b> 32.0	0.686 -	32	43	15

Results are presented as median values over all the missing elements. The best performance in each setting is marked in boldface.

**Table 4.** Fiber-wise imputation results under cross-validation for the neonatal ClrX application.

Number of components	BAMITA Correlated			BAMITA Independent		EM Algorithm
	Imputation MSE (Coverage)	Low-rank MSE	Shannon MSE (Coverage)	Imputation MSE (Coverage)	Shannon MSE (Coverage)	Imputation MSE
	1	0.340 (95.4)	0.583	0.744 (84.3)	0.560 (94.1)	1.834 (55.3)
2	0.345 (95.3)	0.570	0.698 (84.4)	0.533 (94.3)	1.471 (55.0)	0.557
3	0.353 (95.1)	0.561	0.665 (84.2)	0.518 (94.5)	1.035 (63.0)	0.554
4	0.365 (95.0)	0.547	0.635 (81.9)	0.505 (94.5)	0.924 (63.3)	0.543
5	0.384 (94.8)	0.552	0.611 (82.0)	0.505 (94.3)	0.903 (61.7)	0.555
6	0.391 (94.6)	0.547	0.608 (81.8)	0.497 (94.2)	0.874 (60.0)	0.554
7	0.423 (94.3)	0.576	0.608 (80.9)	0.496 (94.1)	0.827 (59.7)	0.555
8	0.439 (94.2)	0.578	0.605 (80.8)	0.499 (94.1)	0.794 (59.1)	0.554

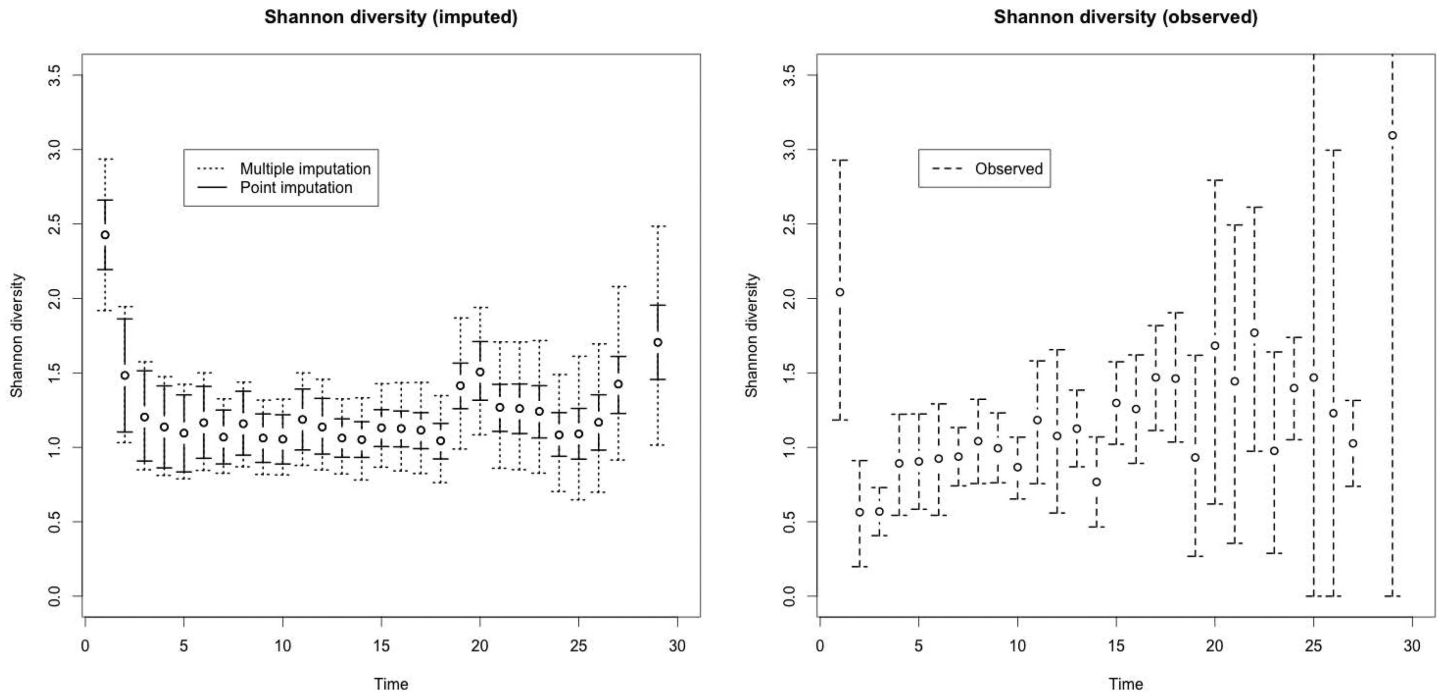
'Imputation MSE' gives relative MSE for held-out values in the tensor, 'low-rank MSE' gives relative MSE when imputing via the low-rank term only in the correlated model, and 'Shannon MSE' gives MSE for imputed Shannon entropy for held-out fibers. Coverage rates for 95% credible intervals are shown in parentheses.

We apply the rank-1 correlated model to the full data, with no validation set, to infer trends in microbiome diversity over time. We consider three approaches to generate uncertainty bounds for the mean diversity at each time point. For approach 1., we impute the missing data at their posterior mean, treat the values as fixed, and create 95% confidence intervals using the classical frequentist approach via a t-distribution. For approach 2., we use only the observed data at each timepoint, and create a 95% interval via a t-distribution. Note the t-interval for a mean is equivalent to a Bayesian credible interval with a uniform prior on the mean and log-uniform prior on the variance of the values. Thus, for approach 3. we use the posterior samples to propagate uncertainty from the imputation step and generate credible intervals for the full model. That is, let  $\text{mean}(\alpha_{k,t})$  and  $\text{sd}(\alpha_{k,t})$  be the sample mean and standard deviation for diversity at timepoint  $k$  and MCMC iteration  $t$ , including observed values and imputed values simulated from the posterior. Then, we simulate a value for the population mean  $\bar{\alpha}_{k,t}$  via

$$\bar{\alpha}_{k,t} = \text{mean}(\alpha_{k,t}) + \text{sd}(\alpha_{k,t})T_{51}$$

where  $T_{51}$  is a t-distributed random variable with 51 degrees of freedom. Intervals obtained via the quantiles of the  $\bar{\alpha}_{k,t}$  will then properly account for both variability in the imputed values and sampling variability.

The resulting trends for Shannon diversity, with 95% confidence/credible bounds, are shown in Fig. 1. The confidence bounds generated using the point-imputed data are substantially more narrow than the bounds generated using multiple imputation in some instances, illustrating the danger of underestimating uncertainty if imputed data are treated as fixed and known. In contrast, the bounds generated using only the observed data are sporadic and very wide in some cases, illustrating the disadvantages of completely ignoring missing data. The multiple imputation approach is a principled compromise between these two extremes, and show a pattern in which diversity decreases over the first few days of life in the NICU and then stays relatively constant.



**Figure 1.** Estimates for Shannon diversity over time for the neonatal infant ClrX data. The left panel gives the mean under imputation and credible intervals generated using either point or multiple imputation, the right panel gives the mean and credible interval for observed data only.

Section S1 presents another application to longitudinal microbiome data for an experiment in mice. This application similarly illustrates the advantages of the correlated BAMITA method for imputation accuracy and uncertainty propagation.

## 6. DISCUSSION

Our results demonstrate the advantages of accounting for residual covariance and uncertainty when imputing missing values in tensor data. While the motivating application for the development of BAMITA was longitudinal microbiome data, the model is generally applicable to a wide variety of application scenarios. Aspects of the model and sampling algorithm may be modified or extended, e.g. to capture spatiotemporal structure in relevant modes (Yokota et al. 2016; Guan 2024) rather than a general covariance. Moreover, the assumption of normality for the residual error may be relaxed. For example, a tensor modeling approach is often effective for multi-condition RNA-Seq gene expression data (Hore et al. 2016), which may have a Poisson or negative binomial distribution.

In our data applications we selected the rank using cross-validation (Algorithm 2). Thus, subsequent analyses with the selected rank for the same data are prone to post-selection inference. However, as just one parameter (the rank) is empirically estimated, over fitting is not a major concern. For example, in our simulation, coverage rates are appropriate when the rank is selected through cross-validation. Nevertheless, a possible extension is to infer the number of components  $R$  (i.e. rank) as a parameter within the Bayesian model to account for its uncertainty, using reversible jump MCMC (Frühwirth-Schnatter et al. 2024) or other approaches. However, this would increase the computational complexity of the approach.

To make our algorithm applicable in various data application scenarios, we adopted the flat prior for the underlying parameters  $\mathbf{U}$  which is independent of the data scale. However, under the specific data application scenarios, informative Gaussian priors can also be considered in our algorithm as a conjugate prior. One future topic is to extend our algorithm with the informative Gaussian priors on the underlying structure elements  $\mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \dots, \mathbf{U}^{(N)}$ .

In our real data applications, we use cross-validation to decide whether to adopt the Bayesian multiple imputation algorithm with the separable covariance structure or the independent structure, using mean squared error (MSE) of the imputed elements. However, alternative Bayesian model selection criteria, such as the deviance information criterion (DIC) (Spiegelhalter et al. 2002), could also be used to select the most appropriate model.

Computing time is often a bottleneck to fully Bayesian inference for high-dimensional data. We have carefully specified our models to facilitate efficient Gibbs sampling in high-dimensions. However, the model with independence error allows for a much more efficient algorithm; our largest dataset, described in Section 5, took several hours to run under the correlated model and under 10 min with independent error. Thus, this is a trade-off to modeling the residual covariance in higher dimensions.

## ACKNOWLEDGMENTS

The authors thank the anonymous reviewers for their valuable suggestions.

## SUPPLEMENTARY MATERIAL

Supplementary material is available at *Biostatistics Journal* online.

## FUNDING

This work was supported in part by National Institutes of Health grants R01-HG010731 and R01-GM130622.

## CONFLICT OF INTEREST

None declared.

## REFERENCES

- Acar E, Dunlavy DM, Kolda TG, Morup M. 2011. Scalable tensor factorizations for incomplete data. *Chemom Intell Lab Syst.* 106:41–56.
- Carroll JD, Chang J-J. 1970. Analysis of individual differences in multidimensional scaling via an n-way generalization of “eckart-young” decomposition. *Psychometrika.* 35:283–319.
- Chen X, He Z, Sun L. 2019. A Bayesian tensor decomposition approach for spatiotemporal traffic data imputation. *Transport Res C Emerging Technol.* 98:73–84.
- Chen Y-L, Hsu C-T, Liao H-YM. 2013. Simultaneous tensor decomposition and completion using factor priors. *IEEE Trans Pattern Anal Mach Intell.* 36:577–591.
- Cong X, Judge M, Xu W, Diallo A, Janton S, Brownell EA, Maas K, Graf, J. 2017. Influence of feeding type on gut microbiome development in hospitalized preterm infants. *Nursing Res.* 66:123–133.
- Frühwirth-Schnatter S, Hosszejni D, Lopes HF. 2024. Sparse Bayesian factor analysis when the number of factors is unknown. *Bayesian Anal.* 1:1–31.
- Guan L. 2024. Smooth and probabilistic parafac model with auxiliary covariates. *J Comput Graph Stat.* 33:538–550.
- Guhaniyogi R, Qamar S, Dunson DB. 2017. Bayesian tensor regression. *J Mach Learn Res.* 18:1–31.
- Hoff PD. 2011. Separable covariance arrays via the Tucker product, with applications to multivariate relational data. *Bayesian Anal.* 6:179–196.
- Hoff PD. 2015. Multilinear tensor regression for longitudinal relational data. *Annals Appl Stat.* 9:1169.
- Hore V, Vinuela A, Buil A, Knight J, McCarthy MI, Small K, Marchini J. 2016. Tensor decomposition for multiple-tissue gene expression experiments. *Nat Genet.* 48:1094–1100.
- Kolda TG, Bader BW. 2009. Tensor decompositions and applications. *SIAM Rev.* 51:455–500.
- Liu J, Musialski P, Wonka P, Ye J. 2012. Tensor completion for estimating missing values in visual data. *IEEE Trans Pattern Anal Mach Intell.* 35:208–220.
- Mazumder R, Hastie T, Tibshirani R. 2010. Spectral regularization algorithms for learning large incomplete matrices. *J Mach Learn Res.* 11:2287–2322.
- Salakhutdinov R, Mnih A. 2008. Bayesian probabilistic matrix factorization using Markov chain Monte Carlo. In: *Proceedings of the 25th International Conference on Machine Learning, Helsinki, Finland.* New York NY, USA: Association for Computing Machinery. p. 880–887.
- Shannon CE. 1948. A mathematical theory of communication. *Bell Syst Techn J.* 27:379–423.
- Spiegelhalter DJ, Best NG, Carlin BP, Van Der Linde A. 2002. Bayesian measures of model complexity and fit. *J R Stat Soc Ser B (Stat Methodol).* 64:583–639.
- Stekhoven DJ, Bühlmann P. 2012. Missforest—non-parametric missing value imputation for mixed-type data. *Bioinformatics.* 28:112–118.
- Tan H, Feng G, Feng J, Wang W, Zhang Y-J, Li F. 2013. A tensor-based method for missing traffic data completion. *Trans Res C Emerg Technol.* 28:15–27.
- Thukral AK. 2017. A review on measurement of alpha diversity in biology. *Agric Res J.* 54:1–10.
- Tucker LR. 1966. Some mathematical notes on three-mode factor analysis. *Psychometrika.* 31:279–311.
- Wang K, Xu Y. 2024. Bayesian tensor-on-tensor regression with efficient computation. *Stat Its Interface.* 17:199.
- Wu Y, Tan H, Li Y, Zhang J, Chen X. 2018. A fused CP factorization method for incomplete tensors. *IEEE Trans Neural Netw Learn Syst.* 30:751–764.
- Yokota T, Zhao Q, Cichocki A. 2016. Smooth parafac decomposition for tensor completion. *IEEE Trans Signal Process.* 64:5423–5436.