

Addressing Missing Data in Untargeted Metabolomics: Identifying Missingness Type and Evaluating the Impact of Imputation Methods on Experimental Replication

Trenton J. Davis^{1,2,†}, Tarek R. Firzli^{3,†}, Emily A Higgins Keppler^{1,2,†}, Matthew Richardson^{4,5,*}, and Heather D. Bean^{1,2,*}

¹ School of Life Sciences, Arizona State University, Tempe, AZ, USA

² Center for Fundamental and Applied Microbiomics, The Biodesign Institute, Tempe, AZ, USA

³ School of Medicine, University of Nevada, Reno, NV, USA

⁴ Department of Respiratory Sciences, College of Life Sciences, University of Leicester, Leicester, UK

⁵ NIHR Biomedical Research Centre (Respiratory Theme), Institute for Lung Health, Leicester, UK

†These authors contributed equally to this manuscript.

*Corresponding authors: Matthew Richardson (mr251@leicester.ac.uk); Heather D. Bean (Heather D. Bean@asu.edu)

ABSTRACT: (Missing data is a significant issue in metabolomics that is often neglected when conducting data pre-processing, particularly when it comes to imputation. This can have serious implications for downstream statistical analyses and lead to misleading or uninterpretable inferences. In this study, we aim to identify the primary types of missingness that affect untargeted metabolomics data and compare strategies for imputation using two real-world comprehensive two-dimensional gas chromatography (GC×GC) data sets. We also present these goals in the context of experimental replication whereby imputation is conducted in a within-replicate-based fashion—the first description and evaluation of this strategy—and introduce an R package **MetabImpute** to carry out these analyses. Our results conclude that, in these two data sets, missingness was most likely of the missing at-random (MAR) and missing not-at-random (MNAR) types as opposed to missing completely at-random (MCAR). Gibbs sampler imputation and Random Forest gave the best results when imputing MAR and MNAR compared against single-value imputation (zero, minimum, mean, median, and half-minimum) and other more sophisticated approaches (Bayesian principal components analysis and quantile regression imputation for left-censored data). When samples are replicated, within-replicate imputation approaches led to an increase in the reproducibility of peak quantification compared to imputation that ignores replication, suggesting that imputing with respect to replication may preserve potentially important features in downstream analyses for biomarker discovery.

INTRODUCTION

Untargeted metabolomics analyses geared toward biomarker discovery are at the forefront of the clinical metabolomics field. Mass spectrometry hyphenated with chromatography, including gas and liquid chromatography (GC and LC, respectively), are commonly used methods for investigating the metabolism of various systems, including *in vitro* microbial and tissue cultures, and *ex vivo* samples such as blood, urine, and breath. The use of multi-dimensional chromatography (*e.g.*, GC×GC and LC×LC) is a significant advancement over traditional one-dimensional chromatography in that it allows for the characterization of a much larger number of metabolites in these highly complex samples.¹⁻³ Especially in untargeted analyses using GC×GC or LC×LC, it is common to detect several thousands of compounds. This increase in chromatographic capacity, however, presents challenges in data pre-processing and downstream statistical analyses.² In addition to the high-dimensional nature of the acquired data, which typically yields many more features than samples, high instances of missing values—upwards of 80% for some features—is a major problem. Missing data severely limits even the most robust tools available in the statistical toolkit, creating uncertainty around population estimators of location (means, medians) and dispersion (variances). This decreases statistical power and can lead to inconclusive or misleading inferences. Subsequently, more sophisticated multivariate and statistical learning approaches are also handicapped.

Missing values in metabolomics data sets arise from a variety of technical and biological sources. Data points can be absent because the compounds are truly not present in the sample, because of instrumental limitations (*e.g.*, detection limits), or as a result of ineffective processing of the acquired data. These drivers can be broadly classified into three missingness mechanisms: missing completely at-random (MCAR), missing at-random (MAR), and missing not-at-random (MNAR).^{4, 5} Data points are MCAR if they are missing due to wholly random events or errors, are not related to whether data are present or absent, and occur over the entirety of the data distribution.⁴ For example, in volatile metabolomics, data missing due to issues with sample preparation techniques such as incomplete derivatization or sub-optimal extraction (*e.g.*, adsorption or desorption), or an instrumental issue can be deemed MCAR. Missingness of the MAR type are data points whose missingness is dependent or associated with other observed values.⁴ This can arise as a result of suboptimal data processing steps related to peak alignment and/or deconvolution, or from the choice of a sampling modality or chromatographic column that is selective for some chemical characteristics (*e.g.*, polarity) over others. Finally, MNAR is due to the variable itself, for example data that are missing due to factors such as limits of quantification or detection (LOQ or LOD, respectively).^{4, 6}

One of the most common strategies for handling missing data is a combination of removing compounds that have high proportions of missingness and replacing the remaining missing data points with zero. An application of this approach is the “80% rule” whereby features that have missing data points in more than 20% of samples (*i.e.* “frequency of observation” < 80%) are removed from the data set, after which the remaining missing values are replaced with zeros.⁷ A variation of the 80% rule aimed at reducing the possibility of losing potentially important features is the “modified 80% rule”, which restricts the rule to within sample groups or classes.⁸ While this can reduce the sparsity of the data set, a disadvantage of imputing missing data points to zero (or to a similarly small value such as half of the global minimum) is skewing the distribution or underestimating measures of variance,⁹ despite the assumption that the missingness might be due to LOD and the distribution is left-censored. As such, it is important that more careful thought be given to addressing missingness in metabolomics data, and the implications on pre-processing and downstream statistical analyses, especially as they relate to missingness type.

Characterizing missingness type is challenging, primarily due to the difficulty of identifying the mechanisms that led to the missingness; this creates problems because many imputation solutions require the missingness type be known.¹⁰ While it may be possible to test for a necessary condition of MCAR explicitly, this is not the case for MAR.¹⁰ Testing for MAR versus MNAR is also difficult unless there is a strong *a priori* assumption for MNAR.¹¹ When testing for MAR vs. MCAR, it is important to note that because MAR is a condition of MCAR, tests of MCAR are thus only testing a necessary condition of that missingness type and should only be applied if an assumption of MAR holds *a priori*.¹⁰ In addition, missingness type falls into two general categories of “ignorable” and “non-ignorable”.¹⁰ MCAR and MAR can be considered ignorable in that virtually any imputation method can be used without serious negative consequences. There are numerous approaches for testing the condition of MCAR and MAR against the MNAR alternative, including Little’s test (for MCAR) and likelihood ratio tests (MAR vs. MNAR, though only if there is a strong *a priori* assumption for MNAR).¹⁰ MNAR is considered non-ignorable, and the imputation approach used should be carefully considered. Addressing missingness of this type is difficult in part due to the ambiguity surrounding its driver(s). Consequently, tests for MNAR are much more limited and generally not confident in evaluating the MNAR assumption.¹⁰

In an effort to better understand missingness mechanisms and their implications for pre-processing in untargeted metabolomics data, we explore two real-world GC×GC data sets. We have three goals for this study. First, we aim to gain some level of insight into the missingness mechanism(s) of multi-dimensional metabolomics data.

Second, based on inferences surrounding missingness type, we compare the performance of nine widely known methods for imputing missing data: zero, mean, median, minimum, half-minimum (HM), Random Forest (RF), Bayesian principal component analysis (BPCA), quantile regression imputation of left-censored data (QRILC), and Gibbs sampler imputation (GSImp). Third, we introduce an imputation approach that takes into account the availability of replicate measurements in the data set and the effect of imputation on reproducibility. Analyzing samples in replicate can provide valuable information about missing values, specifically to the question of whether they are “truly” missing, or rather missing due to stochastic events. This approach, along with strategies to assist in evaluating missingness in metabolomics data sets, are wrapped into an R package we created called **Meta-bImpute**. The strategy of imputing in the context of replication has significant implications in data pre-processing and subsequent statistical analysis, and, to our knowledge, this study is the first to describe a replication-based imputation approach and is the first systematic examination of the effect of imputation on reproducibility.

METHODS

Metabolomics data sets. Two real-world metabolomics data sets were used for this study and are described previously.^{12, 13} The first data set—hereafter referred to as the fungal data—is comprised of six strains each of the fungi *Coccidioides immitis* and *C. posadasii*, each grown in biological triplicate in either the mycelia or spherule life cycle, and three uninoculated liquid media controls per life cycle, for a total of 78 samples. The second data set—hereafter referred to as the bacterial data—contains 81 *Pseudomonas aeruginosa* cystic fibrosis chronic lung infection isolates analyzed in biological triplicate, including uninoculated liquid media controls, for a total of 258 samples. For both experiments, headspace volatile compounds were extracted using solid-phase microextraction (SPME) and analyzed by two-dimensional gas chromatography coupled to time-of-flight mass spectrometry (GC×GC-TOFMS). After data cleanup comprised of artifact and contaminant removal, a total of 767 chromatographic features were identified for the fungal data set and 979 chromatographic features for the bacterial data set.

Evaluation of missingness mechanism. Pairwise correlations (Pearson and Spearman) between each feature’s missingness vector (where present values are replaced with 1 and missing values with 0) and all other features in the data set were calculated. Features were considered MCAR if no significant correlations were found between any pair of features (Benjamini-Hochberg-corrected $P \leq .05$); features were grouped as MAR or MNAR

otherwise. Additionally, the strength of the pairwise correlations regardless of significance was used to assess missingness in which moderate or stronger correlations (≥ 0.5) were suggestive of MCAR. The Kolmogorov-Smirnov (KS) goodness-of-fit test and the Cucconi test with Benjamini-Hochberg correction were applied to the remaining features to determine whether their distributions were left-truncated, indicative of possible MNAR missingness.¹⁴ Features with left-truncated distributions when compared to left censored normal distributions were characterized as MNAR, and MAR otherwise. Features with greater than 60% missingness were excluded from analysis.

Simulation and missing value generation. Simulated complete data sets of identical dimensionality were created by generating matrices based on the estimated covariance structure of the incomplete fungal and bacterial data sets using an approach described previously.¹⁵ For both simulated data sets we utilized the approach detailed by Kokla *et al.* to simulate MAR, mixed (1:1) MAR/MNAR, and MNAR missingness mechanisms for missing proportions from 5% to 50%.¹⁵ This process involves sequentially choosing two features at random and removing a select proportion of the largest values in one of these features in the case of MAR. For MNAR, we consider only left-truncated missingness. Here, single variables have a random proportion of their lowest values removed, and this process is repeated until the desired missingness proportion is reached.

Missing value imputation. Nine common imputation methods were evaluated (descriptions in Supplementary Information). Minimum, half-minimum (HM), mean, and median imputation methods are “single-value” imputation methods that impute missing values of a feature to the global values indicated by the respective methods. For zero imputation, missing values were replaced with zeros. For the MAR, MNAR, and mixed MAR/MNAR simulated data sets, zero imputation was excluded because simulated data were drawn from a normal distribution based on covariance matrices. Because this distribution is continuous with no clear lower bound, there is no obvious numeric choice that is analogous to zero. For analyses involving the fungal and bacterial data sets, zero imputation was utilized as the basis of comparison between imputation methods. Four additional algorithmic approaches—Bayesian principal component analysis (BPCA)¹⁶ Random Forest (RF),¹⁷ quantile regression imputation of left-censored data (QRILC),¹⁸ and Gibbs sampler imputation (GSImp)¹⁹—were considered. Additionally, for every method above, an approach that considers replication was devised. First, for samples in which there are less than 50% present values within the set of replicates, the present values within the replicate group are permuted to zero and the missing values are imputed to zero (*i.e.* all values in the replicate group are set to zero). For the data sets in this study, which

have three replicates, the lone present value was permuted to zero and the two missing values were imputed to zero. Second, for samples in which there are more than 50% present values within the set of replicates, the present values are left unchanged while missing values are imputed using one of the nine imputation approaches described above: for minimum (RepMin), half-minimum (RepHM), mean (RepMean), and median (RepMedian), using the present values within the replicate samples; for BPCA (RepBPCA), RF (RepRF), QRILC (RepQRILC), and GSImp (RepGSImp), using the present values across the entire data set. For zero (RepZero), the missing values within the replicate are imputed with zeros.

Performance evaluation. Normalized root-mean-square error (NRMSE) and PCA-Procrustes analysis (PCA-P).^{15, 20} Intraclass correlation coefficients (ICC) were used as a measure of reproducibility for each feature.²¹ ICCs were classified as: excellent $\geq .90$, good 0.75 - 0.89, moderate 0.50 - 0.74, and poor < 0.50 .²² Overall change in ICC was measured for an imputation method by taking the mean change in ICC over n number of features (F) versus two baselines, the original data (with missing values) and the zero-imputed data:

$$\Delta ICC_{imputed} = \frac{\sum_F (ICC_{imputed_F} - ICC_{baseline_F})}{n}$$

RESULTS AND DISCUSSION

Evaluation of data set missingness types. We first set out to identify the missingness profiles of two metabolomics data sets, one from an analysis of fungal volatile metabolites and the other from an analysis of bacterial volatile metabolites. Because the acquisition of metabolomics data involves many confounding factors that can affect the presence or absence of a metabolite, such as instrumental LOD/LOQ, numerous method parameters, and sampling methodologies, we posit that missingness of the completely at-random type is unlikely. Using Pearson and Spearman correlation analyses between missing and present volatiles, it was found that every feature was significantly correlated ($P \leq .05$) to the missingness of at least one other feature for both data sets (pairwise correlations not shown). Thus, there were likely not any features that could be categorized as MCAR using this method. By exclusion, we determined that our data was comprised of either MAR, MNAR, or a mixture of the two. Furthermore, we assumed that the most likely mechanism of left-censored MNAR is due to below-limits of detection, a commonly presumed cause of missingness in metabolomics⁶. The Kolmogorov-Smirnov (KS) test, which tests for a left-truncated distribution, was used to investigate this. The missingness profiles of both data sets were evaluated by the proportion of features classified as each missingness type (relative to the total number of features) and by the contribution of individual missing values (relative to the total number of missing

values). Results suggested a large number of volatile features possessing left-truncated distributions, and we characterized these to be predominantly MNAR. The largest proportion of features across the entirety of both data sets were considered MAR, followed by features that could not be evaluated due to high degrees of missingness ($\geq 60\%$). In both data sets features that were excluded from analysis of missingness type because they were detected in $< 60\%$ of samples had the highest contribution to overall missingness followed by those with MAR (Figure 1).

Volatiles that were excluded have more than 60% of their data missing; thus, it is reasonable to conclude that these features contributed largely to the missingness. Although it is impossible to know with certainty the mechanism of missingness for these features, pairwise correlation analysis suggested that there was a low degree of MCAR when using only significant correlations (classified as MCAR if Benjamini-Hochberg-corrected $P < .05$ for all pairwise correlations). When assessing MCAR using correlation coefficients regardless of significance, inferences are less clear as strong correlations (≥ 0.7) suggest the presence of small proportions of MCAR (16% and 19% for the fungal and bacterial data sets, respectively). Between MAR and MNAR, however, there is more ambiguity. This is largely due to the fact that the KS test assumes a normal distribution and is testing whether or not the distributions of the volatile features are approximately normal (Gaussian) versus left-censored. Density plots of these data sets show that many features' distributions are skewed to the right, consistent with data that is missing due to LOD/LOQ (Figure S1). Cullen and Frey plots, which can be used to infer the distributions of these features, show that normal distributions are not common in either data set (Figure 2). Because of this, it is likely that models dependent on normal distributions poorly fit the data and give spurious inferences about missingness. To bolster the results from the KS test, we applied the non-parametric Cucconi test which compares location and scale between two distributions. Results from the Cucconi test differed from those of the KS test in that they indicated higher proportions of MAR data (Figure 1). These conflicting results speak to the difficulty in characterizing distributions and identifying missingness mechanisms.

While it was found that a majority of volatiles appear to be MAR or MNAR, there are limitations to these approaches, primarily due to the difficulty in identifying the true distribution of metabolomic features. Distribution analysis showed that there is a wide variation in potential "true" distributions of features, and this analysis is itself limited by high degrees of missingness. Often, assumptions must be made by researchers regarding the likeliest missingness mechanism in order to properly select an imputation method; however, clear guidelines do not

exist to select appropriate methods. It is important to note that the choice of an imputation method should not be taken lightly, as biases may be introduced into the data. These relationships can skew statistical results and subsequent inferences. Though we did not explore other methods here, inferences about the missingness type might be made by examining frequencies of observation of features within groups, such as sample replicates or sample classes. Examining the signal-to-noise ratio of features can also aid in inspecting missingness type. Information gleaned from these strategies can inform what imputation approaches may be the most appropriate, as well as how to impute with respect to the experimental design (*i.e.* imputation across the data set as a whole, or imputation within sample classes or replicates).

Imputation evaluation of simulated complete data. To evaluate the impacts of imputation on the data sets, simulated complete data matrices based on the covariance structure of the real data were generated for each data set. Matrices with missingness levels ranging from 5% to 50% were also generated with intervals of 5% in three types of missingness mechanisms: MAR, MNAR, and mixed (1:1) MAR/MNAR. Imputation methods were applied and normalized root-mean-square error (NRMSE) and PCA-Procrustes (PCA-P) analyses were performed to compare the imputed matrices to the complete matrices, where smaller values indicate better relative performance. In the MAR mechanism, we find that RF imputation outperformed other imputation methods in both performance metrics for the MAR mechanism of missingness and GSImp outperformed other imputation methods for the MNAR when analyzed by NRMSE (Figure 3), confirming the findings of other studies.^{15, 19} The PCA-P analysis of the MNAR data showed that HM and QRILC perform equally well as GSImp, even at high levels of missingness. In the mixed missingness mechanism approach combining MAR and MNAR, we find RF is the highest performing imputation method when evaluated by NRMSE, while PCA-P scores suggest that at low proportions of missingness (up to 20%), all but mean and median imputation methods perform well. However, as the proportion of missingness increases, BPCA begins to perform poorly, and while HM, Minimum, GSImp, QRILC, and RF all continue to perform quite well, GSImp was superior at the highest levels of missingness. GSImp may be particularly robust as an imputation method as it employs elastic net regularization to overcome the problem of high dimensionality (p features $\gg n$ observations) and better accommodates unique correlation structures. Notably, two of the worst performing imputation methods by NRMSE score—HM and QRILC—performed well by PCA-P (Figure 3). The discrepancy in performance of QRILC between these two metrics may be due to the assumption of normality required for QRILC. These results also indicate that different imputation methods could be more useful depending on what downstream analysis is used, a conclusion that previous research also sug-

gests.²³ For example, RF may be a good choice for imputation when the data are dominated by MAR or a mixture of MAR and MNAR, which are the two most likely scenarios for the bacterial and fungal data (Figure 1). This analysis of imputation methods, however, does not address the degree to which imputation affects sample reproducibility when replicate data are available, as they are for both the fungal and bacterial data.

Effects of imputation on replication. The samples in the bacterial and fungal data sets were analyzed in biological replicates of three, and therefore these correlated samples contain additional information that can be leveraged to improve the reproducibility of the samples via imputation of missing values. A simple measure of reproducibility is the intraclass correlation coefficient (ICC). Applying ICC thresholds have been used by us and others as a filter during data cleanup and post-processing to simplify the data structure and aid in the identification of important features. This is accomplished by removing those features that have low reproducibility while retaining potentially important rare features.^{12, 13, 24-26} We measured the reproducibility of the metabolomes within these triplicates before and after imputation using the nine standard (non-replicate) imputation methods, as well as the same methods adapted for replicate data. The replicate imputation approaches assign all values for a feature in a replicate group to zero if the within-group missingness exceeds 50% for that feature (where a group is defined as a set of biological replicates); alternatively, if within-group missingness is less than 50%, the missing values are imputed based on one of the standard imputation approaches. This “majority-minority” imputation approach treats peaks that are detected in a majority of sample replicates as more likely to be a product of the metabolism being studied, and peaks that are detected in a minority of sample replicates as more likely to be present due to chance. Using the mean of the change in ICC between the original data or the zero-imputed data (as a baseline), and the standard and replicate imputation methods, we found that all of the standard imputation methods decreased feature ICC compared to the original data containing missing values, while replicate imputation methods increased ICC across all features for both the data sets when compared to the original data (Figure 4a) or standard zero imputation (Figure 4b). Further, the replicate imputation method outperformed the equivalent standard imputation method in all cases, when measured by ICC (Figure 4).

We analyzed the driving factors of the increase in ICC observed when employing within-replicate imputation and found that a large proportion of features possessed excellent ICC after imputation of missing values, whereas with standard imputation methods there was a higher proportion of features with poor ICC (Figure 5). Taking a closer look, we observed large numbers of features with great increases in ICC in the majority of replicate imputation methods when compared to both zero-imputed data as a baseline, and no imputation. For example, Re-

pHM imputation results in 346 and 464 volatiles shifting from poor ICC to excellent ICC in the fungal and bacterial data sets, respectively (Figure 6; Table S1). We hypothesized that this upward shift in feature ICCs was primarily due to the permutation of replicate groups to zero in the case that only a single value out of three was present. To explore this hypothesis, the proportion of replicate groups that were permuted to zero of all the groups present was plotted against the change in ICC of all imputation methods compared to the zero imputed data (as a baseline) for each feature. Scatter plots of these results show that features in which there are a large number of replicates permuted entirely to zero tend to have large positive ICC changes (Figure S2). Pearson correlations of the ICC changes to proportion of features permuted to zero in the fungal and bacterial data ranged from 0.78-0.89 ($P = 0+$) and 0.62-0.83 ($P = 0+$), respectively, depending on the within-replicate imputation method used (Table S2), suggesting that permuting features of a replicate group to zero if the majority of values are missing is an important mechanism for increasing ICC.

MetabImpute for exploring missingness and performing replicate imputation. The replicate imputation methods described in this study can be adapted for any number of replicates, and the threshold for the percentage of missing values that are imputed may also be modified. Based on the analyses performed in this study, we developed an R package called **MetabImpute** to explore various aspects of metabolomics data sets, especially those with technical or biological replicates. Included in the package are tools to evaluate missingness mechanisms, variable distributions, to impute using all the methods described in this paper and finally to evaluate the effects of imputation on ICC. Table S3 lists the package dependencies and adapted/modified packages utilized in **MetabImpute**.

CONCLUSIONS

Handling missing data is a significant issue in metabolomics data post-processing, and can significantly affect downstream analysis. Methods do exist to assess mechanisms of missingness; however, many rely on inherent assumptions about the missingness mechanisms themselves, and others cannot be utilized in metabolomics data that commonly contain significantly more features than samples. Additionally, certain methods of excluding features with high missingness, such as the 80% rule, introduce the risk of excluding potentially important features (*e.g.*, biomarkers).

In this study, we took two real-world GC×GC volatile metabolomics data sets that were collected in biological triplicates and evaluated their missingness profiles and the mechanisms that underlie them. We concluded that

out of the features we were able to analyze, the missingness profiles in both the data sets were likely not MCAR, but a mixture of MAR and MNAR, with a higher proportion of the former. A large number of volatiles also possessed high degrees of missingness, and followed many different distributions. Using simulated data generated from the covariance matrices of our two data sets, eight imputation methods were examined for accuracy in the setting of MAR, mixed MAR/MNAR and MNAR missingness using NRMSE and PCA-P analysis. RF imputation appeared to outperform other imputation methods with MAR data, while GSImp outperformed other imputation methods with MNAR data. In mixed missingness data, results were less clear, however RF outperformed other methods by the NRMSE measure. Several methods performed well by the PCA-P measure when there are low proportions of missingness, while GSImp outperformed other methods when missingness increased, even up to 50%. This suggests that certain

data sets where there is a high level of missingness and a mixture of MAR and MNAR are suspected, GSImp may be an optimal imputation method, though RF is a good choice when the mechanism of missingness is not well defined, but likely contains a significant proportion of MAR.

In the setting of data with biological or technical replicates, the effects of imputation have not been studied previously. We introduced and evaluated a methodology described as replicate imputation, where features that have a high percentage of missingness within the replicates are permuted to zero, and missing value imputation is performed on the remaining features using the present values within the sets of replicates. Uniformly, utilizing the replicate imputation method led to higher overall ICC, which may preserve relationships within replicate groups and more features for downstream analysis that could be important biomarker discovery. We conveniently present these missingness evaluation and imputation strategies in an R package called **MetabImpute** such that explorations into missingness in metabolomics data sets can be easily incorporated into a data pre-processing pipeline.

ASSOCIATED CONTENT

Supporting Information

See accompanying PDF file.

Data Availability

The data sets included in this study (chemical feature peak areas and retention time information) are available at the NIH Common Fund's National Metabolomics Data Repository (NMDR) website, the Metabolomics Workbench, at www.metabolomicsworkbench.org, where they have been assigned: Project ID PR001064 and Study ID ST001659 (<https://doi.org/10.21228/M85H6W>) (fungal data); and Project ID PR000970 and Study ID ST001414 (<https://doi.org/10.21228/M89Q4F>) (bacterial data). The **MetabImpute** R package is available at <https://github.com/BeanLabASU/MetabImpute>.

Notes

The authors declare no competing financial interests.

ACKNOWLEDGMENTS

Funding for the collection of the metabolomics data was provided by the Cystic Fibrosis Foundation (Hill17P0) and Cystic Fibrosis Foundation Therapeutics (Hill18A0-CI), the National Institutes of Health (R56HL139846), and the Arizona Biomedical Research Centre (ADHS18-198861). Bacterial isolates were obtained from the Cystic Fibrosis Isolate Core at Seattle Children's Center for Global Infectious Disease Research, funded by the National Institutes of Health (P30 DK089507) and the Cystic Fibrosis Foundation (HOFFMA20Y2-OUT).

REFERENCES

1. Patrushev, Y. V., Advantages of two-dimensional gas chromatography. *Kinet. Catal.* **2015**, *56* (3), 386-393.
2. Higgins Keppler, E. A.; Jenkins, C. L.; Davis, T. J.; Bean, H. D., Advances in the application of comprehensive two-dimensional gas chromatography in metabolomics. *Trends Analyt. Chem.* **2018**, *109*, 275-286.
3. Jandera, P., Comprehensive two-dimensional liquid chromatography — practical impacts of theoretical considerations. A review. *Cent. Eur. J. Chem* **2012**, *10* (3), 844-875.
4. Rubin, D. B., Inference and missing data. *Biometrika* **1976**, *63* (3), 581-592.
5. Little, R. J. A.; Rubin, D. B., *Statistical analysis with missing data*. 2 ed.; John Wiley & Sons: 2002.
6. Do, K. T.; Wahl, S.; Raffler, J.; Molnos, S.; Laimighofer, M.; Adamski, J.; Suhre, K.; Strauch, K.; Peters, A.; Gieger, C.; Langenberg, C.; Stewart, I. D.; Theis, F. J.; Grallert, H.; Kastenmüller, G.; Krumsiek, J., Characterization of missing values in untargeted MS-based metabolomics data and evaluation of missing data handling strategies. *Metabolomics* **2018**, *14* (10), 128.
7. Smilde, A. K.; van der Werf, M. J.; Bijlsma, S.; van der Werff-van der Vat, B. J. C.; Jellema, R. H., Fusion of mass spectrometry-based metabolomics data. *Anal. Chem.* **2005**, *77* (20), 6729-6736.

8. Yang, J.; Zhao, X.; Lu, X.; Lin, X.; Xu, G., A data preprocessing strategy for metabolomics to reduce the mask effect in data analysis. *Front. Mol. Biosci.* **2015**, *2* (4).
9. Gelman, A.; Hill, J., Missing-data Imputation. In *Data analysis using regression and multilevel/hierarchical models*, Cambridge university press: 2006; pp 529-542.
10. Rhoads, C. H., Problems with tests of the missingness mechanism in quantitative policy studies. *Stat Politics Policy* **2012**, *3* (1).
11. Molenberghs, G.; Beunckens, C.; Sotito, C.; Kenward, M. G., Every missingness not at random model has a missingness at random counterpart with equal fit. *J. R. Stat. Soc. Series B Stat. Methodol.* **2008**, *70* (2), 371-388.
12. Higgins Keppler, E. A.; Mead, H. L.; Barker, B. M.; Bean, H. D., Life Cycle Dominates the Volatilome Character of Dimorphic Fungus *Coccidioides* spp. *mSphere* **2021**, *6* (2), e00040-21.
13. Davis, T. J.; Karanjia, A. V.; Bhebbhe, C. N.; West, S. B.; Richardson, M.; Bean, H. D., *Pseudomonas aeruginosa* Volatilome Characteristics and Adaptations in Chronic Cystic Fibrosis Lung Infections. *mSphere* **2020**, *5* (5), e00843-20.
14. Chernobai, A.; Rachev, S. T.; Fabozzi, F. J., Composite Goodness-of-Fit Tests for Left-Truncated Loss Samples. In *Handbook of Financial Econometrics and Statistics*, Lee, C. F.; Lee, J. C., Eds. Springer: New York, NY, 2015; pp 575-596.
15. Kokla, M.; Virtanen, J.; Kolehmainen, M.; Paananen, J.; Hanhineva, K., Random forest-based imputation outperforms other methods for imputing LC-MS metabolomics data: a comparative study. *BMC Bioinformatics* **2019**, *20* (1), 492.
16. Stacklies, W.; Redestig, H.; Scholz, M.; Walther, D.; Selbig, J., pcaMethods—a bioconductor package providing PCA methods for incomplete data. *Bioinformatics* **2007**, *23* (9), 1164-1167.
17. Tang, F.; Ishwaran, H., Random forest missing data algorithms. *Stat. Anal. Data Min.* **2017**, *10* (6), 363-377.
18. Lazar, C. *imputeLCMD: A collection of methods for left-censored missing data imputation*, 2; 2015.
19. Wei, R.; Wang, J.; Jia, E.; Chen, T.; Ni, Y.; Jia, W., GSimp: A Gibbs sampler based left-censored missing value imputation approach for metabolomics studies. *PLoS Comput. Biol.* **2018**, *14* (1), e1005973.

20. Miller-Atkins, G.; Acevedo-Moreno, L.-A.; Grove, D.; Dweik, R. A.; Tonelli, A. R.; Brown, J. M.; Allende, D. S.; Aucejo, F.; Rotroff, D. M., Breath Metabolomics Provides an Accurate and Noninvasive Approach for Screening Cirrhosis, Primary, and Secondary Liver Tumors. *Hepatol. Commun.* **2020**, *4* (7), 1041-1055.
21. Wolak, M. E.; Fairbairn, D. J.; Paulsen, Y. R., Guidelines for estimating repeatability. *Methods Ecol. Evol.* **2012**, *3* (1), 129-137.
22. Koo, T. K.; Li, M. Y., A Guideline of Selecting and Reporting Intraclass Correlation Coefficients for Reliability Research. *J Chiropr Med* **2016**, *15* (2), 155-63.
23. Di Guida, R.; Engel, J.; Allwood, J. W.; Weber, R. J. M.; Jones, M. R.; Sommer, U.; Viant, M. R.; Dunn, W. B., Non-targeted UHPLC-MS metabolomic data processing methods: a comparative investigation of normalisation, missing value imputation, transformation and scaling. *Metabolomics* **2016**, *12* (5), 93.
24. Bean, H. D.; Rees, C. A.; Hill, J. E., Comparative analysis of the volatile metabolomes of *Pseudomonas aeruginosa* clinical isolates. *J Breath Res* **2016**, *10* (4), 047102.
25. Rees, C. A.; Nordick, K. V.; Franchina, F. A.; Lewis, A. E.; Hirsch, E. B.; Hill, J. E., Volatile metabolic diversity of *Klebsiella pneumoniae* in nutrient-replete conditions. *Metabolomics* **2017**, *13* (2), 18.
26. Mead, H. L.; Roe, C. C.; Higgins Keppler, E. A.; Caballero Van Dyke, M. C.; Laux, K. L.; Funke, A. L.; Miller, K. J.; Bean, H. D.; Sahl, J. W.; Barker, B. M., Defining Critical Genes During Spherule Remodeling and Endospore Development in the Fungal Pathogen, *Coccidioides posadasii*. *Front. Genet.* **2020**, *11* (483). (Word Style "TF_References_Section"). References are placed at the end of the manuscript. Authors are responsible for the accuracy and completeness of all references. Examples of the recommended formats for the various reference types can be found at <http://pubs.acs.org/page/4authors/index.html>. Detailed information on reference style can be found in The ACS Style Guide, available from Oxford Press.

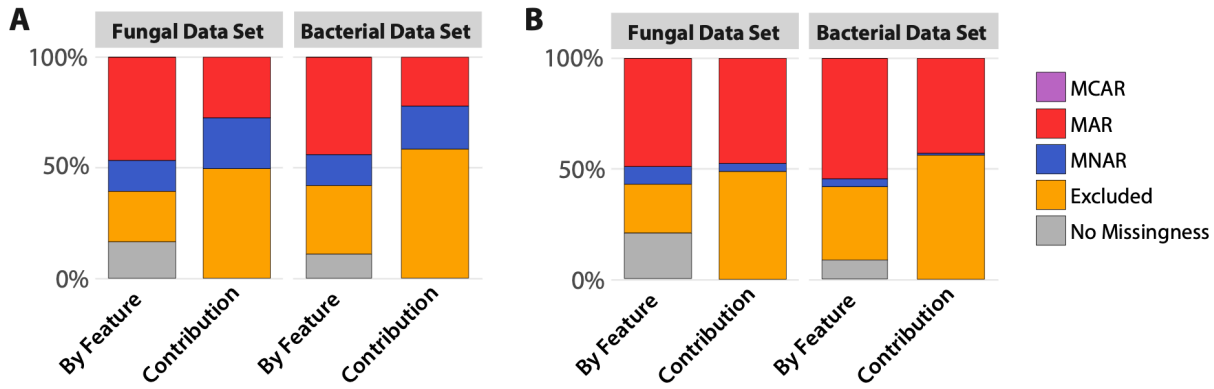


Figure 1. Percent missingness type by volatile feature (number of features of each missingness type divided by total number of features), and percent contribution of each volatile feature's missingness to the overall missingness profile (number of missing values within each feature for each missingness type divided by total number of missing values). (a) Classification using Pearson correlation and KS test. (b) Classification using Pearson correlation and Cucconci test. Excluded refers to features which could not be evaluated due to high degrees of missingness ($\geq 60\%$). MCAR was not detected in either data set.

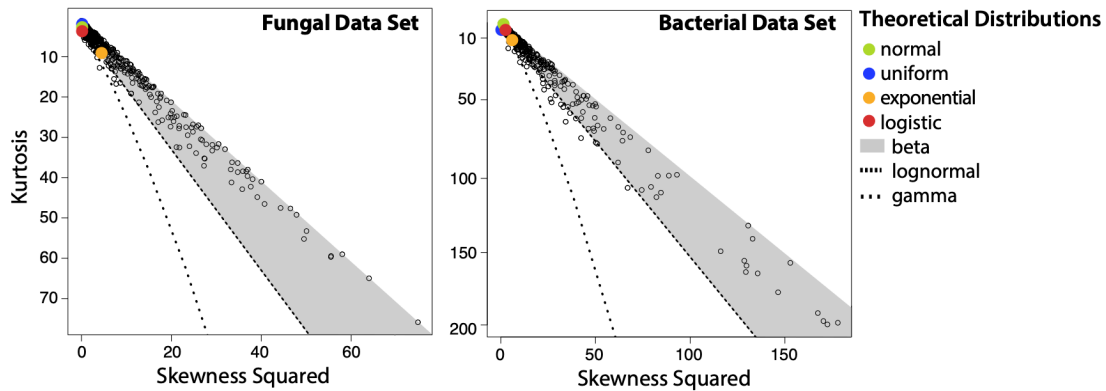


Figure 2. Cullen and Frey plots of estimated distributions of all variables in each data set.

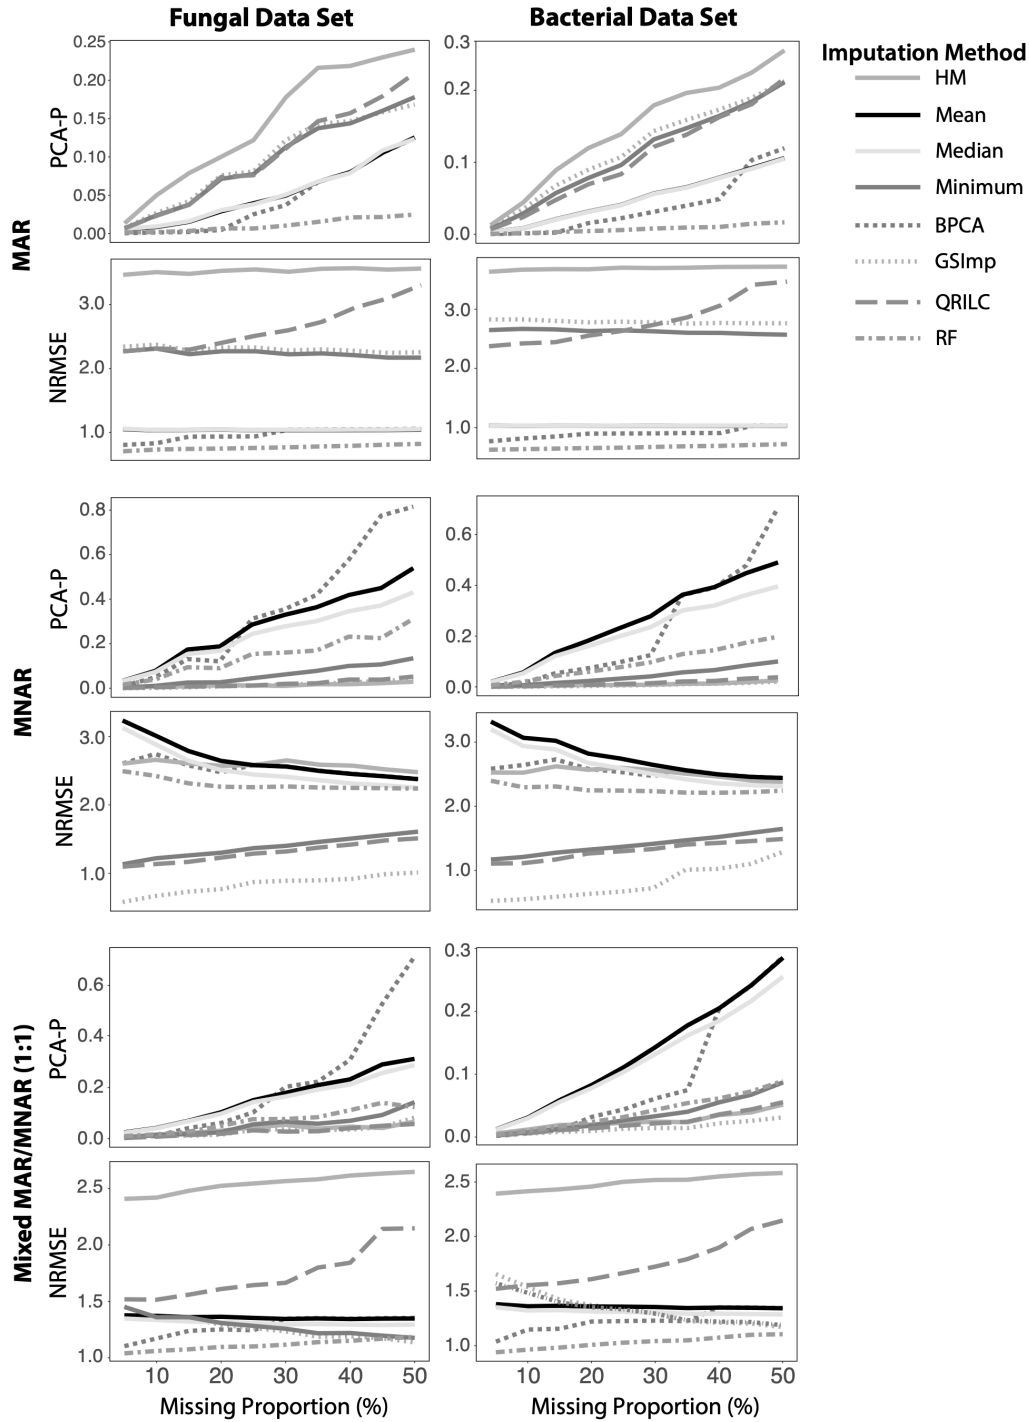


Figure 3. Performance of imputation methods applied to various proportions of MAR, MNAR, and mixed (1:1) MAR:MNAR missingness simulated from the fungal and bacterial data sets, where smaller values indicate better relative performance. HM = half-minimum, BPCA = Bayesian principal components analysis, GSImp = Gibbs sampler imputation, QRILC = quantile regression imputation of left-censored data, RF = Random Forest.

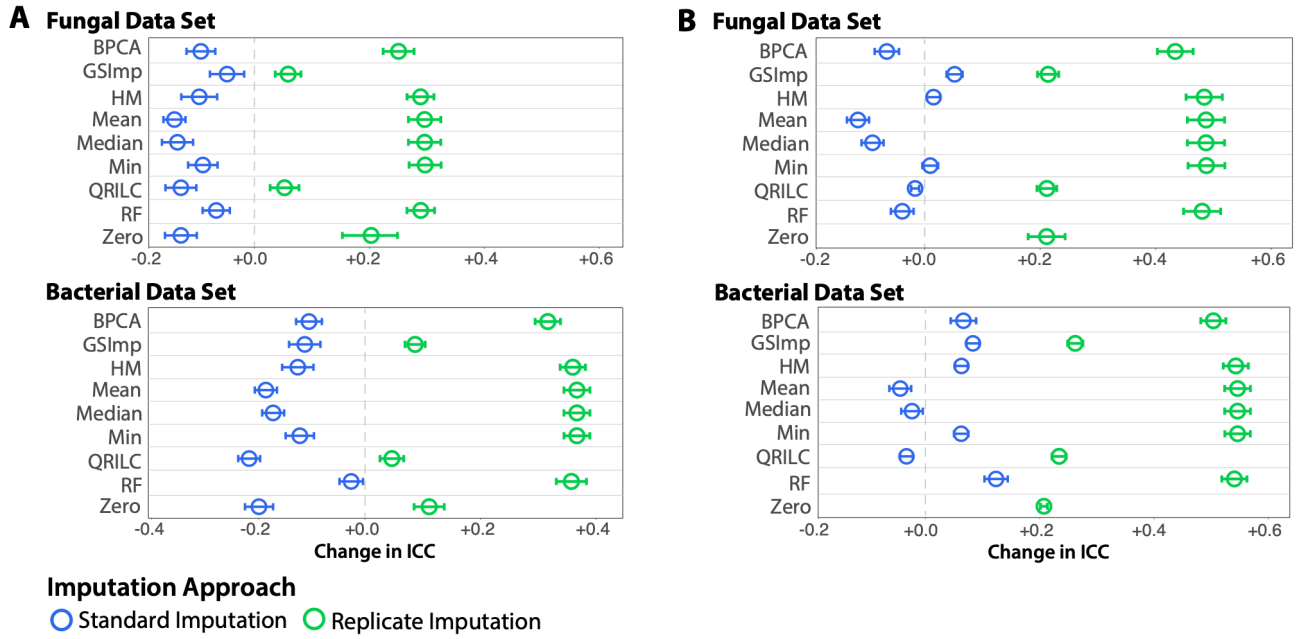


Figure 4. Mean change in intraclass correlation coefficients (ICC) across all individual features between imputation methods compared to (a) original un-imputed data with missing values, and (b) zero-imputed data as a baseline using standard imputation (blue) or replicate imputation (green) methods. Whiskers represent the non-simultaneous 95% confidence intervals. HM = half-minimum, BPCA = Bayesian PCA, GSImp = Gibbs sampler imputation, QRILC = quantile regression imputation of left-censored data, RF = Random Forest.




Figure 5. Proportion of features in each ICC category following imputation. HM = half-minimum, BPCA = Bayesian PCA, GSImp = Gibbs sampler imputation, QRILC = quantile regression imputation of left-censored data, RF = Random Forest, Rep = replicate version of imputation. Excellent: ≥ 0.90 ; Good: 0.75–0.89; Moderate: 0.50–0.74; Poor: < 0.50 .

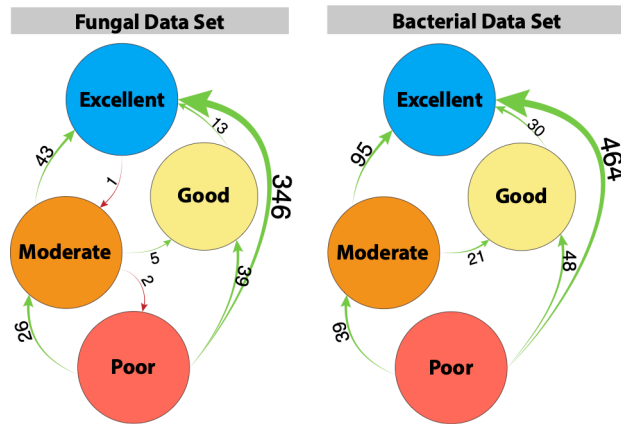


Figure 6. Network plot of ICC shifts of volatile features following within-replicate half-minimum (Re-pHM) imputation. Excellent: ≥ 0.90 ; Good: 0.75-0.89; Moderate: 0.50-0.74; Poor: < 0.50 .