OXFORD

PEPerMINT: peptide abundance imputation in mass spectrometry-based proteomics using graph neural networks

Tobias Pietz[1,‡], Sukrit Gupta [ID][1,2,‡], **Christoph N. Schlaffner** [ID][1,3], **Saima Ahmed** [ID][3],
Hanno Steen [ID][3], **Bernhard Y. Renard** [ID][1,4], **Katharina Baum** [ID][1,4,5,*]

[1]Hasso Plattner Institute, Digital Engineering Faculty, University of Potsdam, Potsdam, 14482, Germany
[2]Department of Computer Science and Engineering, Indian Institute of Technology, Ropar, Rupnagar, 140001, India
[3]Department of Pathology, Boston Children's Hospital and Harvard Medical School, Boston, MA, 02115, United States
[4]Windreich Department for Artificial Intelligence and Human Health and Hasso Plattner Institute at Mount Sinai, Icahn School of Medicine at Mount Sinai, New York City, NY, 10029, United States
[5]Department of Mathematics and Computer Science, Free University Berlin, Berlin, 14195, Germany

*Corresponding author. Department of Mathematics and Computer Science, Free University Berlin, 14195 Berlin, Germany.
E-mail: katharina.baum@fu-berlin.de (K.B.)

‡Shared authorship.

Abstract

Motivation: Accurate quantitative information about protein abundance is crucial for understanding a biological system and its dynamics. Protein abundance is commonly estimated using label-free, bottom-up mass spectrometry (MS) protocols. Here, proteins are digested into peptides before quantification via MS. However, missing peptide abundance values, which can make up more than 50% of all abundance values, are a common issue. They result in missing protein abundance values, which then hinder accurate and reliable downstream analyses.

Results: To impute missing abundance values, we propose PEPerMINT, a graph neural network model working directly on the peptide level that flexibly takes both peptide-to-protein relationships in a graph format as well as amino acid sequence information into account. We benchmark our method against 11 common imputation methods on 6 diverse datasets, including cell lines, tissue, and plasma samples. We observe that PEPerMINT consistently outperforms other imputation methods. Its prediction performance remains high for varying degrees of missingness, different evaluation approaches, and differential expression prediction. As an additional novel feature, PEPerMINT provides meaningful uncertainty estimates and allows for tailoring imputation to the user's needs based on the reliability of imputed values.

Availability and implementation: The code is available at https://github.com/DILiS-lab/pepermint.

1 Introduction

Proteins are the main acting molecules in cells. The characterization of their quantity in different biological contexts plays a fundamental role in understanding cellular function and regulation in disease (Anderson and Anderson 1998, Buccitelli and Selbach 2020). Methods based on label-free mass spectrometry (MS) are commonly used for high-throughput quantification of protein abundance in biological samples (Rozanova *et al.* 2021). In MS-based bottom-up proteomics, proteins are enzymatically digested into peptides before subjecting them to a mass spectrometer. Individual peptides are then commonly identified by matching their spectra to corresponding databases (Sinitcyn *et al.* 2018). With data-dependent acquisition (DDA), only the top most abundant peptides within a given analysis time window are individually fragmented and used for identification and quantification. In contrast, data-independent acquisition (DIA) fragments all peptides within a given time and mass window. The higher sensitivity of DIA has increased its use in recent years (Guo and Aebersold 2023). Finally, several aggregation methods exist to infer protein abundance by computationally aggregating the measured peptide abundance values into protein abundances (Cox *et al.* 2014, Fischer and Renard 2016) to allow downstream analysis on the protein level.

With label-free MS, peptide abundance measurements exhibit a high number of missing values (e.g. 22.1%–68.8% for the datasets used in this article). These might either be due to peptides with an abundance below the detection limit, often referred to as missing not at random (MNAR), or due to random errors and stochastic fluctuations in the measurement process, often referred to as missing completely at random (MCAR) (Karpievitch *et al.* 2012, Lazar *et al.* 2016, Välikangas *et al.* 2017). While performing peptide-to-protein aggregation, these missing values can propagate to the protein level and ultimately hamper downstream analyses (Liu and Dongre 2021, Kong *et al.* 2022). Therefore, different methods for imputing missing values following different paradigms—relying on single (e.g. minimal) values, leveraging local similarities or global structure—have been suggested and benchmarked (Lazar *et al.* 2016, Liu and Dongre 2021, Shen *et al.* 2022) (see overview in Table 1). Basic methods, such as average, k-nearest neighbors (KNNs), iterative singular value decomposition (ISVD), principal component analysis (PCA), or random forest (RF), which are applicable

Table 1. Overview of imputation methods used for our benchmark.

Method	Benchmarked by	Supported by
Single value		
MinDet	[1, 2]	[2, 3, 4, 5]
MinProb	[1, 2, 6, 7]	[2, 3, 4, 5]
Median	[2, 8]	[2]
Local similarity		
KNN	[1, 2, 6, 7, 8, 9]	[2, 3, 4, 5, 10]
RF [11]	[2, 7, 8]	[2, 3, 4, 5]
MICE [12]	[2, 8]	[2, 10, 12]
Global structure		
ISVD [13]	[1, 2, 6, 7]	[2, 15]
BPCA [14]	[2, 6, 7, 8, 9]	[2, 3, 4, 5, 15]
DAE [8]	[8]	[8]
VAE [8]	[8]	[8]
CF [8]	[8]	[8]

We capture basic methods, more complex ones, and imputation methods based on deep neural networks representing all three generally considered categories of imputation methods. Further selection criteria were their appearance in proteomics imputation benchmark studies and their availability in terms of support in open-source software packages. [1] Lazar *et al.* (2016), [2] Wang *et al.* (2020), [3] Feng *et al.* (2023), [4] Gatto and Lilley (2012), [5] Rainer *et al.* (2022), [6] Liu and Dongre (2021), [7] Jin *et al.* (2021), [8] Webel *et al.* (2024), [9] Webb-Robertson *et al.* (2015), [10] Pedregosa *et al.* (2011), [11] Stekhoven and Bühlmann (2012), [12] Van Buuren and Groothuis-Oudshoorn (2011), [13] Troyanskaya *et al.* (2001), [14] Bishop (1998), [15] Stacklies *et al.* (2007).

beyond proteomics have been especially widely adopted (Oba *et al.* 2003, Webb-Robertson *et al.* 2015). More complex extensions, mostly dedicated to protein imputation, based on mixture models or matrix factorization, have been suggested (Shen *et al.* 2022, Kong *et al.* 2023). In addition, adaptations of basic methods such as RF or linear regression models have been proposed to include additional features like mRNA measurements (Asensio *et al.* 2022).

Deep learning (DL) was suggested for missing value imputation in omics datasets. Arisdakessian *et al.* (2019) introduced a basic neural network of several fully connected layers and a single dropout layer for imputing single-cell RNA sequencing datasets. Webel *et al.* (2024) proposed the application of denoising autoencoders (DAE) and variational autoencoders (VAE) for the imputation of MS-based proteomics datasets. Barzine *et al.* (2020) used a neural network and mRNA expression values along with context information from GO terms and UniProt keywords to predict missing protein abundance values.

While additional information, such as mRNA measurements, can improve imputation performance, obtaining mRNA data requires costly additional wet-lab experiments or might even be infeasible (e.g. for plasma samples). Furthermore, there are multiple proteomics-specific features beyond GO terms or UniProt keywords that can provide additional information and context to machine learning models for learning patterns across similar proteins or peptides that current imputation methods fail to exploit. In particular, similarities in physical properties of the measured molecules, such as peptide mass, sequence length, and charge state, are helpful for peptide-to-protein aggregation (Fischer and Renard 2016). Also, amino acid sequence information is available for missing proteins and peptides, and language models pre-trained on amino acid sequences have shown good performance on a variety of protein-related tasks (Elnaggar *et al.* 2021). The embeddings derived from these pre-trained language models encode the valuable biophysical properties of the underlying protein or peptide but currently

remain unused as features for imputation. In addition, neither of the described DL-based imputation methods considers the particular relationships between proteins and peptides. Peptides originating from the same protein are expected to have strongly correlated abundances, a relationship that can be exploited to improve imputation performance and that also enables leveraging abundance information on non-unique peptides. Moreover, while multiple contributions have focused on imputing values in high missingness scenarios (Barzine *et al.* 2020, Kong *et al.* 2023), little attention has been paid to the inherent uncertainty coming with such imputations. So far, most imputation methods have not been designed with uncertainty in mind, resulting in uncertainty estimates for imputed values either being not available or obtained via multiple imputations (Li *et al.* 2015). Nevertheless, uncertainty estimates are of high value as they can enhance the trust in imputation results, and also help users filter out uncertain imputations.

We here address these gaps and propose a new DL-based model for imputation in proteomics datasets that exploit additional proteomics features in the form of amino acid sequences and peptide–protein relationships. As graph neural network (GNN) models have shown considerable success in modeling complex relationships between molecules and learning from biological and omics data (Ioannidis *et al.* 2019, Li *et al.* 2021, Witzke *et al.* 2023), our method relies on a GNN architecture. What is more, while most proteomics imputation methods still impute the protein level, our model acts directly on the peptide level, a strategy shown to yield improved imputation results (Lazar *et al.* 2016). Furthermore, our DL architecture enables uncertainty estimates for imputed values at low computational overhead to provide the user with a valuable tool for imputation prediction diagnostics. We systematically benchmark our novel method against 11 imputation methods from different categories across 6 representative datasets with different ground-truth mechanisms using 3 evaluation metrics (see overview in Fig. 1). Furthermore, we showcase its uncertainty quantification capabilities.

2 Materials and methods

We introduce PEPerMINT (PEPtide Mass spectrometry Imputation NeTwork), a method combining abundance values and information from amino acid sequences and protein–peptide relations to impute missing values on the peptide level. For its implementation and systematic benchmarking, we use our novel open-source PyProteoNet framework (see Supplementary data).

2.1 PEPerMINT imputation

For our PEPerMINT imputation model, we propose a neural network architecture combining a learnable transformation of abundance values, a GNN operating on the peptide graph, as well as amino acid sequence embeddings derived from a transformer-based language model (see Fig. 1A for a visual overview).

2.1.1 Input features

We assume a proteomics dataset with abundance values for n (potentially non-unique) peptides measured across s samples given as $n \times s$ matrix $\mathbf{A}$ where the elements of $\mathbf{A}$ either represent logarithmized (natural logarithm) and standardized
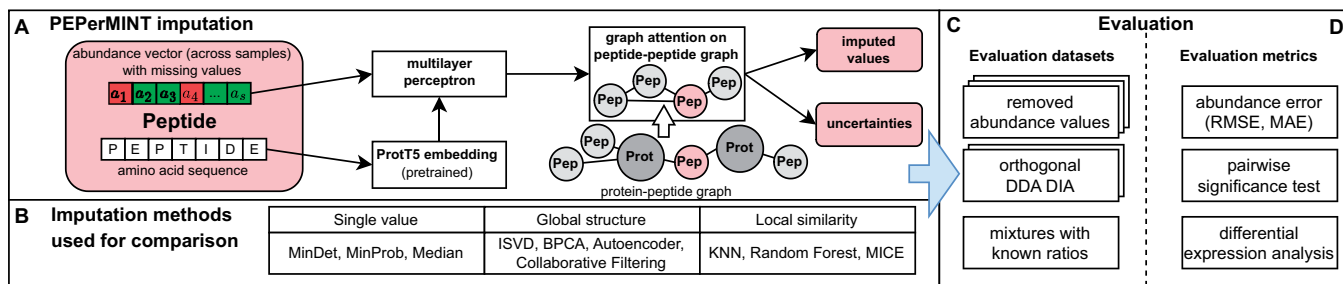
Figure 1. Overview of our PEPerMINT imputation method and our benchmarking framework. (A) Our PEPerMINT imputation model combines both peptide sequence information and abundance values across samples into a latent representation. Structural information is included via a peptide–peptide graph using a graph attention layer. (B) PEPerMINT is compared to 11 published imputation methods from 3 different categories. (C and D) We perform a systematic evaluation on six diverse datasets with ground truth derived from three different mechanisms with respect to three different evaluation metrics (see Materials and methods for details).


Figure 2. Simplified representation of the architecture of PEPerMINT with input feature representations (gray) and learnable (multilayer) transformations (blue). See Supplementary Fig. S1 for a detailed visualization.

(zero mean, unit variance) abundance values or missing values. Missing values are ignored for logarithmization and standardization. We address the problem of predicting abundance values for the missing values. PEPerMINT takes two inputs: the abundance matrix A and an $n \times 1024$ sequence embedding matrix S. S is precomputed from the peptide amino acid sequences using the ProtT5 language model, which has previously shown good performance generating protein embeddings from sequence strings for tasks like predicting protein secondary structure (Elnaggar *et al.* 2021). This allows PEPerMINT to account for abundance values from non-missing samples as well as different biophysical peptide properties encoded in the sequence embeddings (Elnaggar *et al.* 2021).

2.1.2 Peptide–peptide graph

The digestion of proteins into peptides for MS-based quantification results in the characteristic protein–peptide structure of MS-based datasets that can be described by a bipartite graph (Muth *et al.* 2016) where each peptide is assigned to one or more proteins. This structure can provide valuable information for the imputation of missing values since peptides belonging to the same protein are expected to show similar abundance profiles across samples. We transform this structure into a peptide-only graph $G = (V, E)$ whereby peptides are nodes $\in V$ that have an edge $\in E$ between them if they belong to the same protein. Therefore, in G, all peptides belonging to the same protein are fully connected, and all peptides from proteins with shared peptides form a connected component (see Fig. 1A middle). We provide G as input to PEPerMINT.

2.1.3 Neural network architecture

Figure 2 shows a simplified representation of PEPerMINT's architecture. PEPerMINT scales down the sequence embeddings of each peptide by applying a learnable transformation $f_{\Theta_{\text{Sequence}}} : \mathbb{R}^{n \times 1024} \rightarrow \mathbb{R}^{n \times 16}$. This aims to balance the size of abundance and sequence-based information. Next, for each peptide, we concatenate the sequence embedding and the vector containing peptide abundances across samples (abundance vector) and apply another learnable non-linear transformation to create a latent representation $f_{\Theta_{\text{Latent}}} : \mathbb{R}^{n \times (s+16)} \rightarrow \mathbb{R}^{n \times 128}$. To account for the protein–peptide structure of the dataset represented by the peptide–peptide graph G we use an attention-based GNN consisting of a single GATv2 (Brody *et al.* 2021) layer with 64 heads with each head outputting a vector of shape $\lfloor \frac{s}{2} \rfloor$. To keep the peptide-specific information from our latent representation, we add a skip connection bypassing the GNN. We add another learnable transformation on the concatenated output of the skip connection and the GNN output $f_{\Theta_{\text{Final}}} : \mathbb{R}^{n \times (64 \lfloor \frac{s}{2} \rfloor + 128)} \rightarrow \mathbb{R}^{n \times 128}$.

2.1.4 Uncertainty prediction of imputed values

To allow the estimation of uncertainty for imputed values, abundance values are predicted in a Bayesian setting. At the same time, this allows our model to better adapt to differing amounts of measurement noise for individual peptides (heteroscedastic noise) (Anderle *et al.* 2004, Bakalarski *et al.* 2008). Therefore, instead of single abundance values, mean and variance values of Gaussian abundance distributions are predicted (Bishop 1994) by two separate output heads (Stirn *et al.* 2023) ($f_{\Theta_\mu} : \mathbb{R}^{n \times 128} \rightarrow \mathbb{R}^{n \times s}$ and $f_{\Theta_{\sigma^2}} : \mathbb{R}^{n \times 128} \rightarrow \mathbb{R}^{n \times s}$).

2.1.5 Training scheme and self-supervised learning

We create a test set for each dataset by masking 10% of its abundance values uniformly at random (setting them to missing). However, for DDA/DIA datasets, the test set is given by all missing DDA abundance values that have a corresponding non-missing DIA value. From the remainder of the peptides (after picking the test set), we pick 10% of non-missing values uniformly at random as the validation set and mask them. On the resulting dataset, training is performed in a self-supervised manner. Similar to the training of denoising AE

as, e.g. done by Webel *et al.* (2024), for each training step, we mask a fraction γ of non-missing values and compute the loss over them. The fraction γ is sampled randomly with samples uniformly distributed over the $[5\%, 15\%)$ interval to improve model generalization.

To improve the training performance in the Bayesian setting, the model is trained in two rounds. First, we only train the f_{Θ_μ} head (see Fig. 2) with mean squared error (MSE) loss before tuning the mean μ and variance σ^2 together within a second training run using both output heads $(f_{\Theta_\mu}, f_{\Theta_{\sigma^2}})$ and Gaussian negative log-likelihood loss (Sluijterman *et al.* 2023). For both training rounds, we employ early stopping with respect to the MSE computed on the holdout validation set after each epoch. We define one epoch as consisting of 500 randomly masked datasets.

2.2 Imputation methods used for comparison

To evaluate PEPerMINT, we compare it against a broad, representative selection of 11 methods from the literature that are commonly used for imputation and have appeared frequently in other proteomics benchmarks (see Table 1 and Supplementary data for details).

i) *Single-value methods*: These methods either impute missing abundance values with the same single value or, for each missing value, randomly draw a value from a predetermined distribution. We evaluated MinDet (using the 0.01 quantile of non-missing values within each sample), MinProb (drawing from a normal distribution around the 0.01 quantile within each sample), and Median (peptide-wise across samples) as commonly used methods of this class.

ii) *Local similarity methods*: These methods assume that missing values of a peptide can be predicted from the abundance values of similar peptides. We selected KNN imputation as it is simple and widely used (Webb-Robertson *et al.* 2015, Lazar *et al.* 2016, Wang *et al.* 2020). We also included imputation based on a Bayesian ridge regression model as suggested by the MICE (Van Buuren and Groothuis-Oudshoorn 2011) imputation framework and an RF-based imputation (Stekhoven and Bühlmann 2012) as a commonly used method with good performance reported previously (Jin *et al.* 2021).

iii) *Global structure methods*: These methods assume that proteomics datasets contain redundant information and can, thus, be well described by a low-dimensional representation, which is leveraged for inferring missing values. We use Bayesian principal component analysis (BPCA) (Oba *et al.* 2003) and ISVD (Troyanskaya *et al.* 2001) as the most frequently used representatives. To include DL methods, two AE-based methods (variational AE and denoising AE) and a method based on collaborative filtering (CF), all recently proposed in (Webel *et al.* 2024), were considered.

2.3 Datasets

We use six benchmark datasets with the goal of spanning a variety of biological backgrounds, varying degrees of complexity (blood plasma, cell lines, tumor tissue) with diverse dataset sizes (between <500 to $>13\,000$ proteins and $<2\,600$ to $>100\,000$ peptides), and differing percentages of missingness on the peptide level (22.8%–68.1%) for evaluation with respect to three different types of ground truth (see overview in Table 2, and further details in the Supplementary data).

The first three benchmark datasets (A1–A3) do not contain explicit ground-truth values. Therefore, we mask abundances, as commonly done in the literature (Lazar *et al.* 2016, Jin *et al.* 2021), using the measured abundance of masked values as ground truth.

In addition, we use two datasets (B1–B2) acquired in DDA mode with orthogonal ground truth acquired in DIA mode. The more accurate DIA measurements contain fewer missing values, which allows the evaluation of imputation methods on genuinely missing values in the DDA data. To make the DIA and DDA data comparable, all DIA abundance values are scaled to have the same mean as the corresponding DDA abundance values.

For the evaluation of differential expression (DE), we use a dataset (labeled C) of protein mixtures with known (spiked-in) ratios from different organisms serving as ground truth. Similar datasets have been used in the literature to evaluate methods for peptide-to-protein aggregation (Cox *et al.* 2014) and imputation (Webb-Robertson *et al.* 2015, Jin *et al.* 2021).

2.4 Evaluation metrics

For abundance-based evaluation, we use the root mean squared error (RMSE) on all missing values that have non-missing ground-truth values (masked values or values with orthogonal DIA measurements) similar to earlier evaluations of imputation methods (Webb-Robertson *et al.* 2015, Barzine *et al.* 2020, Jin *et al.* 2021). To allow variance estimation, we compute the RMSE sample-wise. As an additional abundance-based evaluation, we compare imputation methods with pairwise significance tests using a Bonferroni-corrected one-sided (paired) Wilcoxon signed-rank test. For every pair of imputation methods, the test compares the two absolute errors of imputed values for each peptide and dataset sample.

In addition, we evaluate imputation methods for the correct identification of differentially expressed peptides. For each peptide, the corresponding sample abundance values

Table 2. Overview of benchmark datasets and their characteristics including availability via ProteomeXchange, dataset ground-truth category (masked: removed abundance values, DDA/DIA, mixture: mixture of known ratios), number of samples (S), number of biological samples (BS), technical replicates per biological sample (TR/BS), number of proteins and peptides, percentage of peptide-level missing values.

Name, Reference	Identifier	Category	#S	#BS	#TR/BS	#Peptides	#Proteins	Missingness (%)
A1: prostate cancer (Chianese *et al.* 2022)	PXD029525	Masked	18	6	3	57 770	6292	54.0
A2: Crohn's fibrosis (Vieujean *et al.* 2021)	PXD022214	Masked	13	13	1	37 158	4481	33.7
A3: breast cancer (Petrosyan *et al.* 2022)	PXD035857	Masked	15	15	1	103 608	13 627	68.8
B1: HEK293-*E.coli* (Dowell *et al.* 2021)	PXD018408	DDA/DIA	16	2	8	16 400	3045	25.2
B2: HIV blood (Ahmed *et al.* 2022)	PXD047528	DDA/DIA	15	15	1	2535	435	41.8
C: HeLa-*E.coli* (Cox *et al.* 2014)	PXD000279	Masked + mixture	6	2	3	50 260	6683	22.1

between groups of replicates of biological samples with different spike-in ratios are compared using a Benjamini–Hochberg corrected Welch's *t*-test. Depending on the significance threshold, different peptides are detected as differentially expressed. Those are compared to a known ground truth of differentially expressed peptides (all spiked-in peptides in our mixture dataset C) to compute true positives and false positives. We assess the performance over varying significance thresholds via a receiver operating characteristics (ROC) curve (see more details in Supplementary data, Section D).

3 Results

We evaluated the performance of our PEPerMINT peptide imputation methods on six proteomics datasets with various biological backgrounds and missingness characteristics. Then, we also compared it against a broad, representative selection of 11 widely used imputation methods. A comparison of the runtime of all presented imputation methods can be found in Supplementary Fig. S8.

3.1 Abundance-based evaluation

We first performed an abundance-based evaluation via the sample-wise RMSE for four datasets with artificially introduced missing values and two DDA/DIA datasets with ground-truth values acquired using the DIA (Fig. 3A). Particularly, we find that our PEPerMINT imputation method gives the best performance across all evaluated datasets, outperforming the second best-performing method (BPCA) by up to 20% on the breast cancer dataset. Out of the other evaluated methods, RF, BPCA, MICE, and CF also show good results. Interestingly, imputing missing values with the peptide-wise median gives better results than the more complex KNN imputation and methods based on autoencoders (DAE, VAE). ISVD, MinDET, and MinProb imputation were found to be generally worse, except for the good performance of MinDet and MinProb on the HIV

blood dataset. We obtain similar results for dataset-wise mean absolute error as an alternative metric (see Supplementary data).

Further, as predicting the missing abundance of a peptide could be hampered if measured only in a few samples, we investigated whether imputation performance depends on the degree of missingness per peptide. Therefore, we stratified the evaluated peptides by their fraction of missing values across samples (Fig. 3B) for the prostate cancer dataset; similar results for other datasets, see Supplementary data). Again, PEPerMINT outperforms other imputation methods for any fraction of missing values. The biggest advantage over competitor methods is observed on peptides with high fractions of missing values. It can also be noted that PEPerMINT, RF, BPCA, CF, DAE, and especially ISVD imputation show improved performance when the fraction of missing values decreases. In contrast, MinProb and MinDet perform better with high fractions of missing values.

Figure 4 shows the results for statistically comparing imputation methods using a Wilcoxon signed-rank test (see Materials and methods). Our PEPerMINT method performs significantly better than all other evaluated methods on the majority of benchmark datasets. It should be noted that in contrast to the sample-wise RMSE results shown in Fig. 3, the Wilcoxon test compares individual imputed values without averaging the error per dataset sample. The good performance of PEPerMINT also holds when stratifying peptides for missingness across samples (see Supplementary Fig. S4).

3.2 Evaluation of DE prediction

DE analysis is a common downstream analysis task performed on MS-based proteomics datasets. Therefore, we compared our proposed method with the other imputation methods with respect to the performance of DE analysis on the imputed dataset. For evaluation, we used the ground-truth protein ratios that can be inferred from the species-specific mixture rates. We restricted the evaluation to peptides that can uniquely be assigned to one species. The


Figure 3. (A) Sample-wise RMSE of all evaluated imputation methods on all six benchmark datasets with 95% confidence interval error bars (bootstrapped). (B) Results for the prostate cancer (A1) dataset stratified by their fraction of missing values over the samples (see Supplementary for stratified results of other datasets). Our newly proposed PEPerMINT imputation outperforms other methods on all datasets, irrespective of the missingness fraction.

ROC curve in Fig. 5 shows that our method is performing better than the other methods, with the highest area under the ROC curve (AUC). The precision–recall curve (Supplementary Fig. S5) supports this result.

3.3 Predicted uncertainty of imputed values

Our PEPerMINT approach allows the out-of-the-box prediction of uncertainty for imputed values, helping users obtain a quantitative estimate of their trustworthiness. To evaluate the usefulness of this computed uncertainty, we compared the imputed values against their ground truth colored by their predicted uncertainty (Fig. 6A). The imputed values with the lowest uncertainty (dark blue) tend to show better predictions (low error) and high abundances. The latter fits with the characteristics of data acquired via MS because high abundance values commonly are proportionally less influenced by measurement noise and are, therefore, assumed to

be more reliable (Renard *et al.* 2012). In addition, we find that removing imputed values with high predicted uncertainty from the evaluation generally improves imputation quality (Fig. 6B). Furthermore, we observe that filtering out imputations with substantial uncertainty but keeping those with low uncertainty, can again massively increase the accuracy of downstream analysis. Re-using the experimental setup of the DE analysis in Fig. 5, we find that by filtering at a predicted uncertainty threshold of 0.2 in Fig. 6C, we can obtain an AUC of 0.84, compared to an AUC of 0.78 for the PEPerMINT imputation alone (versus 0.68 without imputation, see Supplementary data). This further validates the quality and benefit of the uncertainty predictions given by our method.

4 Discussion and conclusion

Overall, PEPerMINT results in superior performance compared to other benchmarked methods across datasets, missingness levels of the peptides, and evaluation metrics. In addition, PEPerMINT provides a handle to the problem of imputation quality by predicting uncertainties for imputed values, with evident improvement potential for downstream analyses. This also distinguishes PEPerMINT from most other imputation methods, which commonly cannot result in confidence statements. We showed that PEPerMINT's uncertainty estimates are highly correlated with imputation error, thereby aptly guiding users on when to rely on or filter out the imputed values. Further, filtering out imputed peptide abundance values with high predicted uncertainty eventually decisively improved the performance of the DE prediction task.

From the diverse characteristics of our benchmark datasets, it can be derived that PEPerMINT's high performance is not limited to a specific dataset size or fraction of missing values. Furthermore, our benchmark comprises datasets with and without technical replicates, i.e. samples with very high similarity. Thus, imputation could be considered easier when relying on technical replicates. However, PEPerMINT's performance seems unaffected by this factor and even outperforms other methods by the largest margin on the breast
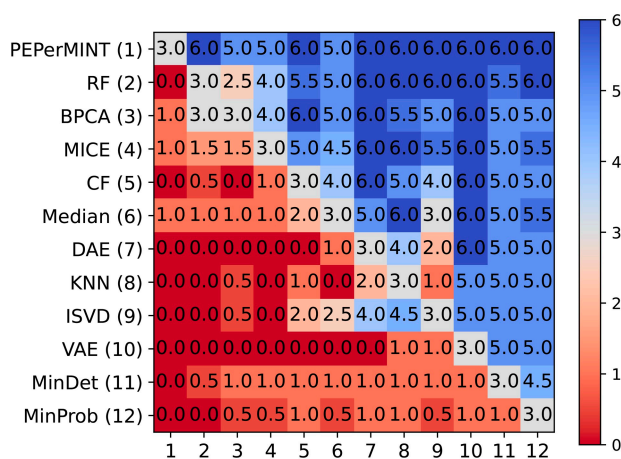

Figure 4. Pairwise comparison of different imputation methods by statistical significance test results (see Materials and methods). Colors encode on how many of the six evaluated datasets the imputation method given in the row performs significantly better (5% significance level) than the imputation method given in the column (insignificant test results increase the count by 0.5). Blue cells indicate the method given in the row outperforms the imputation method given by the column in the majority of cases.
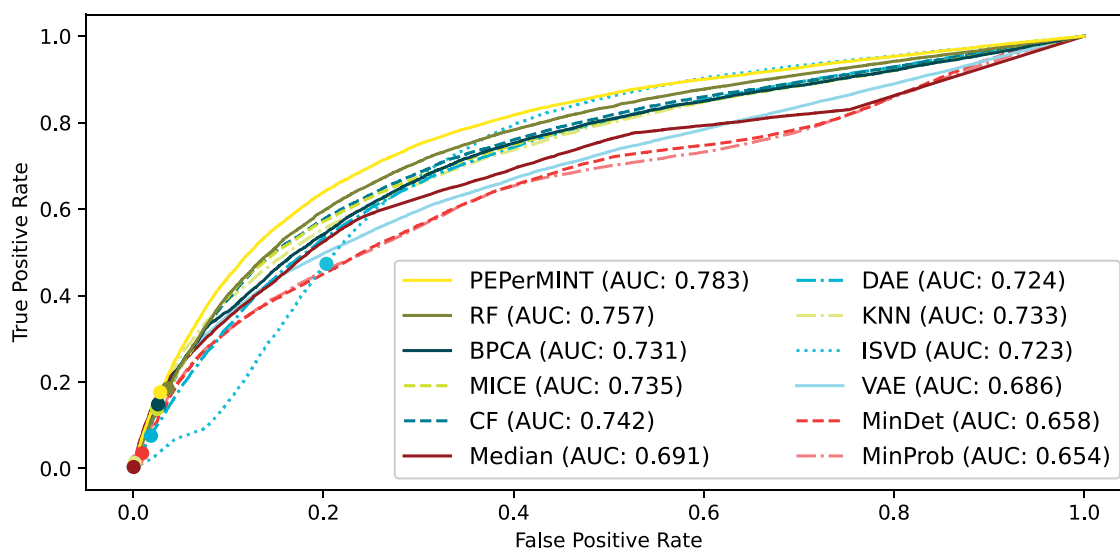

Figure 5. ROC curve for the performance of DE analysis of peptides on the HeLa-*E.coli* dataset imputed with different imputation methods with 5% FDR thresholds marked (dots). Our PEPerMINT imputation method (yellow) outperforms other methods, having the largest area under the curve (AUC).
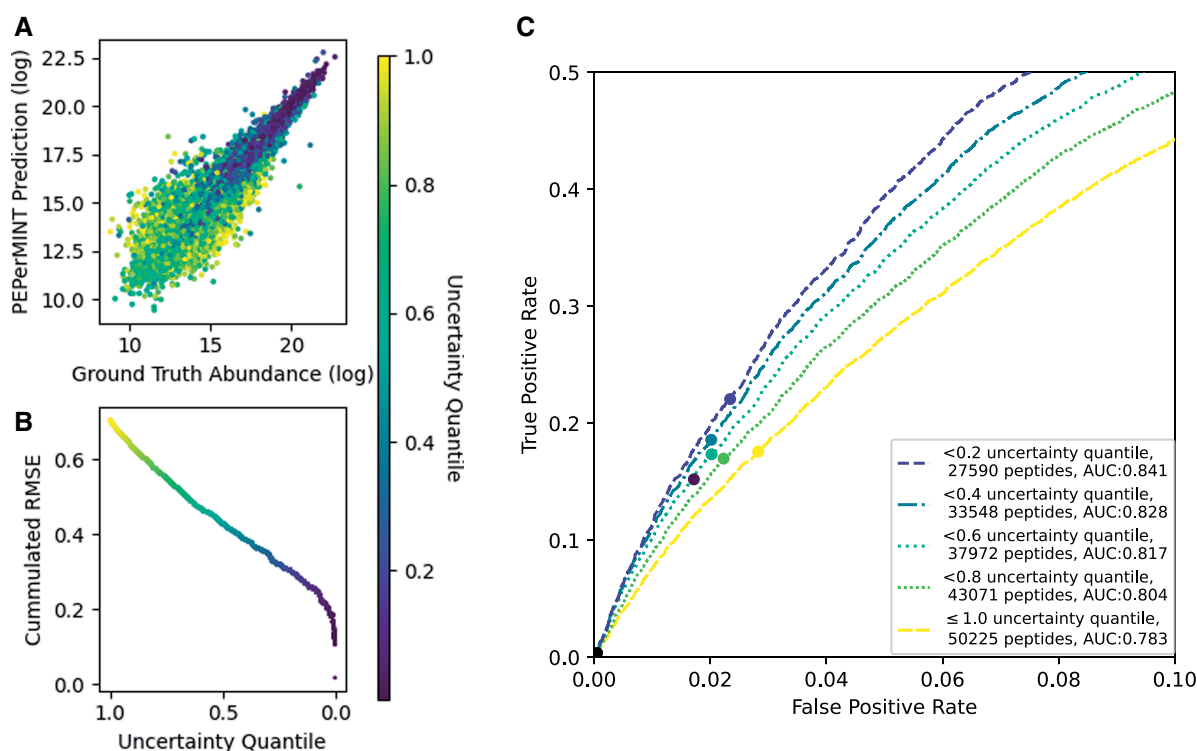
Figure 6. PEPerMINT's predicted uncertainty for the imputed values for the HeLa-*E.coli* dataset (see Supplementary for other datasets). (A) Imputed abundance values versus ground truth colored by predicted uncertainty (low: dark blue, high: yellow). (B) Imputed values ordered by their predicted uncertainty with RMSE computed over different uncertainty quantiles (Iversen *et al.* 2023). (C) ROC curve zoomed in at low FDR values for the performance of DE analysis for imputing values only up to a predicted uncertainty threshold (see Supplementary data for details). The yellow ROC curve (≤1.0 uncertainty) is identical to the yellow PEPerMINT ROC curve from Fig. 5. Filtering out imputed values with high predicted uncertainties decisively improves DE analysis performance after imputation.

cancer dataset, which is devoid of technical replicates. This further hints at PEPerMINT actually learning biologically relevant patterns instead of merely averaging across technical replicates.

Comparing our different categories of datasets, PEPerMINT performs best on our masked benchmark datasets that, by design, only exhibit MCAR missing values. This can be explained by our self-supervised training scheme, which also masks uniformly at random and aligns well with MCAR missing values. Nevertheless, PEPerMINT still shows very good results on DDA benchmark datasets with DIA ground-truth values that can be assumed to contain both MCAR and MNAR missing values. Further, PEPerMINT also performs well on the HIV blood dataset, in which a high fraction of missing values is due to lowly abundant peptides (MNAR) as the blood plasma proteome is well studied with missing values rarely occurring. Its increased percentage of MNAR compared to the other datasets could be the cause for the different ranking of imputation method performances on this dataset, e.g. very good performance of MinDet and MinProb imputation that replace missing values with low abundance values.

For peptides with a high percentage of missing values, PEPerMINT compares especially well against other well-performing methods such as BPCA or RF imputation. This can be explained by PEPerMINT's ability to exploit additional information (amino acid sequence, abundance of peptides belonging to the same protein) to obtain context about a peptide's properties, even if little abundance information is available for the peptide itself. Indeed, using ablation studies,

we find both additional information layers to provide at least some performance benefit to PEPerMINT (see Supplementary Fig. S9). As our method also allows the flexible integration of other information layers both in tabular and graph form, it could be readily extended to improve proteomics imputation even further. Depending on the analysis goal, sample size, missingness as well as desirable and achievable confidence levels of imputations, it has to be decided on a case-by-case basis whether imputation is appropriate for a specific dataset. For explorative studies of small sample sizes with high confidence imputations, it may be advisable to even impute in cases of high missingness ($>50\%$), while this may neither be necessary nor desirable in confirmatory studies with large sample sizes.

A limitation of our method when compared with other imputation methods is its higher runtime (see Supplementary Fig. S8). However, the fastest-running methods like Median or MinProb also perform worse than more complex methods with longer runtimes like BPCA or RF. When executed on a GPU, PEPerMINT shows a runtime similar to or faster than that of BPCA imputation. Of note, all considered imputation methods finish within minutes, which is well acceptable for MS-based proteomics analysis workflows.

Further benchmarking criteria (Harris *et al.* 2023) and methods for proteomic imputation relying on DL and ensembling (Weiping *et al.* 2021) or statistical models that take the protein–peptide structure into account (Etourneau *et al.* 2023) are emerging. They can be used for future exploration and for potential extensions of PEPerMINT, our GNN-based method working directly on the peptide level that flexibly takes both peptide-to-protein relationships and

amino acid sequence information into account to improve the prediction of missing abundance values.

Supplementary data

Supplementary data are available at *Bioinformatics* online.

Conflict of interest

No competing interests are declared.

Funding

Data and code availability

All datasets, with the exception of the HIV blood dataset, were downloaded from the links provided in their original publications. The repositories containing the code and data files used for all experiments can be found under https://github.com/DILiS-lab/pepermint and https://zenodo.org/records/11216899.

References

Ahmed S, Viode A, van Zalm P *et al.* Using plasma proteomics to investigate viral infections of the central nervous system including patients with HIV-associated neurocognitive disorders. *J Neurovirol* 2022; **28**:341–54.

Anderle M, Roy S, Lin H *et al.* Quantifying reproducibility for differential proteomics: noise analysis for protein liquid chromatography-mass spectrometry of human serum. *Bioinformatics* 2004; **20**:3575–82.

Anderson NL, Anderson NG. Proteome and proteomics: new technologies, new concepts, and new words. *Electrophoresis* 1998; **19**:1853–61.

Arisdakessian C, Poirion O, Yunits B *et al.* DeepImpute: an accurate, fast, and scalable deep neural network method to impute single-cell RNA-seq data. *Genome Biol* 2019; **20**:211–4.

Asensio JO, Verheijen M, Caiment F Predicting missing proteomics values using machine learning: filling the gap using transcriptomics and other biological features. *Comput Struct Biotechnol J* 2022; **20**:2057–69.

Bakalarski CE, Elias JE, Villén J *et al.* The impact of peptide abundance and dynamic range on stable-isotope-based quantitative proteomic analyses. *J Proteome Res* 2008; **7**:4756–65.

Barzine MP, Freivalds K, Wright JC *et al.* Using deep learning to extrapolate protein expression measurements. *Proteomics* 2020; **20**:e2000009.

Bishop C. *Mixture Density Networks*. Aston University, Birmingham, 1994.

Bishop C. Bayesian PCA. *Adv Neural Inf Process Syst* 1998; **11**: 382–388.

Brody S, Alon U, Yahav E. How attentive are graph attention networks? In: *International Conference on Learning Representations, virtual conference*. OpenReview 2022, https://openreview.net/forum?id=F72ximsx7C1.

Buccitelli C, Selbach M. mRNAs, proteins and the emerging principles of gene expression control. *Nat Rev Genet* 2020; **21**:630–44.

Chianese U, Papulino C, Passaro E *et al.* Histone lysine demethylase inhibition reprograms prostate cancer metabolism and mechanics. *Mol Metab* 2022; **64**:101561.

Cox J, Hein MY, Luber CA *et al.* Accurate proteome-wide label-free quantification by delayed normalization and maximal peptide ratio extraction, termed MaxLFQ. *Mol Cell Proteomics* 2014; **13**:2513–26.

Dowell JA, Wright LJ, Armstrong EA *et al.* Benchmarking quantitative performance in label-free proteomics. *ACS Omega* 2021; **6**:2494–504.

Elnaggar A, Heinzinger M, Dallago C *et al.* ProtTrans: toward understanding the language of life through self-supervised learning. *IEEE Trans Pattern Anal Mach Intell* 2021; **44**:7112–27.

Etourneau L, Fancello L, Wieczorek S *et al.* A new take on missing value imputation for bottom-up label-free LC-MS/MS proteomics. bioRxiv. 2023. https://www.biorxiv.org/content/10.1101/2023.11.09.566355v1 (13 March 2024, date last accessed).

Feng Z, Fang P, Zheng H *et al.* DEP2: an upgraded comprehensive analysis toolkit for quantitative proteomics data. *Bioinformatics* 2023; **39**:btad526.

Fischer M, Renard BY. iPQF: a new peptide-to-protein summarization method using peptide spectra characteristics to improve protein quantification. *Bioinformatics* 2016; **32**:1040–7.

Gatto L, Lilley KS. MSnbase-an R/bioconductor package for isobaric tagged mass spectrometry data visualization, processing and quantitation. *Bioinformatics* 2012; **28**:288–9.

Guo T, Aebersold R. Recent advances of data-independent acquisition mass spectrometry-based proteomics. *Proteomics* 2023; **23**:2200011.

Harris L, Fondrie WE, Oh S *et al.* Evaluating proteomics imputation methods with improved criteria. *J Proteome Res* 2023; **22**:3427–38.

Ioannidis VN, Marques AG, Giannakis GB. Graph neural networks for predicting protein functions. In: *IEEE CAMSAP*, New York, NY, USA, pp. 221–5. 2019.

Iversen P, Witzke S, Baum K, *et al.* Identifying drivers of predictive uncertainty using variance feature attribution. *arXiv*. 2023. https://arxiv.org/abs/2312.07252v1 (13 March 2024, date last accessed).

Jin L, Bi Y, Hu C *et al.* A comparative study of evaluating missing value imputation methods in label-free proteomics. *Sci Rep* 2021; **11**:1760.

Karpievitch YV, Dabney AR, Smith RD *et al.* Normalization and missing value imputation for label-free LC-MS analysis. *BMC Bioinformatics* 2012; **13** (Suppl 16):S5.

Kong W, Hui HWH, Peng H *et al.* Dealing with missing values in proteomics data. *Proteomics* 2022; **22**:e2200092.

Kong W, Wong BJH, Hui HWH, *et al.* ProJect: a powerful mixed-model missing value imputation method. *Brief Bioinform* 2023; **24**:bbad233.

Lazar C, Gatto L, Ferro M *et al.* Accounting for the multiple natures of missing values in label-free quantitative proteomics data sets to compare imputation strategies. *J Proteome Res* 2016; **15**:1116–25.

Li P, Stuart EA, Allison DB *et al.* Multiple imputation: a flexible tool for handling missing data. *JAMA* 2015; **314**:1966–7.

Li R, Yuan X, Radfar M, *et al.* Graph signal processing, graph neural network and graph learning on biological data: a systematic review. *IEEE Rev Biomed Eng* 2021; 16:109–135.

Liu M, Dongre A. Proper imputation of missing values in proteomics datasets for differential expression analysis. *Brief Bioinform* 2021; **22**:bbaa112.

Muth T, Renard BY, Martens L *et al.* Metaproteomic data analysis at a glance: advances in computational microbial community proteomics. *Expert Rev Proteomics* 2016; **13**:757–69.

Oba S, Sato M-A, Takemasa I *et al.* A Bayesian missing value estimation method for gene expression profile data. *Bioinformatics* 2003; **19**:2088–96.

Pedregosa F, Varoquaux G, Gramfort A *et al.* Scikit-learn: machine learning in python. *J Mach Learn Res* 2011; **12**:2825–30.

Petrosyan V, Dobrolecki LE, LaPlante EL *et al.* Immunologically 'cold' triple negative breast cancers engraft at a higher rate in patient derived xenografts. *NPJ Breast Cancer* 2022; **8**:104.

Rainer J, Vicini A, Salzer L *et al.* A modular and expandable ecosystem for metabolomics data annotation in R. *Metabolites*, 2022; **12**:173.

Renard BY, Xu B, Kirchner M *et al.* Overcoming species boundaries in peptide identification with Bayesian information criterion-driven error-tolerant peptide search (BICEPS). *Mol Cell Proteomics* 2012; **11**:M111–014167.

Rozanova S, Barkovits K, Nikolov M *et al. Quantitative Mass Spectrometry-Based Proteomics: An overview*, pp. 85–116. New York, NY: Springer, 2021.

Shen M, Chang Y-T, Wu C-T *et al.* Comparative assessment and novel strategy on methods for imputing proteomics data. *Sci Rep* 2022; **12**:1067.

Sinitcyn P, Rudolph JD, Cox J *et al.* Computational methods for understanding mass spectrometry-based shotgun proteomics data. *Annu Rev Biomed Data Sci* 2018;**1**:207–34.

Sluijterman L, Cator E, Heskes T. Optimal training of mean variance estimation neural networks. *Neurocomputing* 2024;**597**:127929.

Stacklies W, Redestig H, Scholz M *et al.* pcaMethods—a bioconductor package providing PCA methods for incomplete data. *Bioinformatics* 2007;**23**:1164–7.

Stekhoven DJ, Bühlmann P. MissForest-non-parametric missing value imputation for mixed-type data. *Bioinformatics* 2012;**28**:112–8.

Stirn A, Wessels H, Schertzer M, *et al.* Faithful heteroscedastic regression with neural networks. In: *AISTATS Proceedings*, volume 206 of PMLR, pp. 5593–613. 2023.

Troyanskaya O, Cantor M, Sherlock G *et al.* Missing value estimation methods for DNA microarrays. *Bioinformatics* 2001;**17**:520–5.

Van Buuren S, Groothuis-Oudshoorn K. Mice: multivariate imputation by chained equations in R. *J Stat Soft* 2011;**45**:1–67.

Vieujean S, Hu S, Bequet E *et al.* Potential role of epithelial endoplasmic reticulum stress and anterior gradient protein 2 homologue in Crohn's disease fibrosis. *J Crohns Colitis* 2021;**15**:1737–50.

Välikangas T, Suomi T, Elo LL *et al.* A comprehensive evaluation of popular proteomics software workflows for label-free proteome quantification and imputation. *Brief Bioinform* 2017;**19**:1344–55.

Wang S, Li W, Hu L *et al.* NAguideR: performing and prioritizing missing value imputations for consistent bottom-up proteomic analyses. *Nucleic Acids Res* 2020;**48**:e83.

Webb-Robertson B-JM, Wiberg HK, Matzke MM *et al.* Review, evaluation, and discussion of the challenges of missing value imputation for mass spectrometry-based label-free global proteomics. *J Proteome Res* 2015;**14**:1993–2001.

Webel H, Niu L, Nielsen AB, *et al.* Imputation of label-free quantitative mass spectrometry-based proteomics using self supervised deep learning. *Nature Communications* 2024;**15**:1.

Weiping M, Sunkyu K, Shrabanti C, *et al.* DreamAI: algorithm for the imputation of proteomics data. *bioRxiv* 2021, https://www.biorxiv.org/content/10.1101/2020.07.21.214205v2 (3 March 2024, date last accessed).

Witzke S, Danz N, Baum K, *et al.* Mobility data improve forecasting of covid-19 incidence trends using graph neural networks. In: *epiDAMIK 6.0 at KDD*, ACM, New York, NY, USA. 2023.