

Systems biology

Learning and Imputation for Mass-spec Bias Reduction (LIMBR)

Alexander M. Crowell ^{1,*}, Casey S. Greene ², Jennifer J. Loros³ and Jay C. Dunlap^{1,*}

¹Department of Molecular and Systems Biology, Geisel School of Medicine at Dartmouth, Hanover, NH 03755, USA, ²Department of Systems Pharmacology and Translational Therapeutics, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104, USA and ³Department of Biochemistry and Cell Biology, Geisel School of Medicine at Dartmouth, Hanover, NH 03755, USA

*To whom correspondence should be addressed.

Associate Editor: Jonathan Wren

Received on April 14, 2018; revised on September 14, 2018; editorial decision on September 19, 2018; accepted on September 23, 2018

Abstract

Motivation: Decreasing costs are making it feasible to perform time series proteomics and genomics experiments with more replicates and higher resolution than ever before. With more replicates and time points, proteome and genome-wide patterns of expression are more readily discernible. These larger experiments require more batches exacerbating batch effects and increasing the number of bias trends. In the case of proteomics, where methods frequently result in missing data this increasing scale is also decreasing the number of peptides observed in all samples. The sources of batch effects and missing data are incompletely understood necessitating novel techniques.

Results: Here we show that by exploiting the structure of time series experiments, it is possible to accurately and reproducibly model and remove batch effects. We implement Learning and Imputation for Mass-spec Bias Reduction (LIMBR) software, which builds on previous block-based models of batch effects and includes features specific to time series and circadian studies. To aid in the analysis of time series proteomics experiments, which are often plagued with missing data points, we also integrate an imputation system. By building LIMBR for imputation and time series tailored bias modeling into one straightforward software package, we expect that the quality and ease of large-scale proteomics and genomics time series experiments will be significantly increased.

Availability and implementation: Python code and documentation is available for download at <https://github.com/aleccrowell/LIMBR> and LIMBR can be downloaded and installed with dependencies using 'pip install limbr'.

Contact: alexander.m.crowell@gmail.com or jay.c.dunlap@dartmouth.edu

Supplementary information: [Supplementary data](#) are available at *Bioinformatics* online.

1 Introduction

In recent years, the scope of whole proteome mass spectrometry (MS) experiments has expanded drastically (Chick *et al.*, 2016; Weekes *et al.*, 2014). In principle, increased sample size can make the expression patterns identified by these experiments more robust. However, increasing sample sizes exacerbate many of the technical limitations of MS. The issues of missing data and batch effects that

are common to this experimental platform hamper analyses of dozens of samples (Karpievitch *et al.*, 2012), and this can be particularly problematic when sequential datasets are collected with the goal of discerning trends, for instance in the field of circadian rhythms.

The first challenge presented by larger scale experiments is the failure to observe peptides in all samples. The abundance of a peptide can go unmeasured in a sample because of low abundance,

random sampling or as a function of its sequence (Mallick *et al.*, 2007; Tabb *et al.*, 2003). Many circadian analysis methods, including the best Jonckheere–Terpstra–Kendall based methods, do not support missing data, forcing researchers to remove peptides with any unobserved values or use inferior tests of rhythmicity (Hughes *et al.*, 2010; Hutchison *et al.*, 2015; Mauvoisin *et al.*, 2014; Robles *et al.*, 2014, 2017; Wang *et al.*, 2017a). Because the missingness of peptides can depend on factors other than abundance and random sampling (e.g. sub-cellular localization or degree of structure), the analysis may in addition be biased. Missingness as a function of abundance can also bias analyses because only the most abundant and least dynamically expressed peptides will be present in a complete-case analysis (Gelman and Hill, 2007). If we imagine a dynamically expressed peptide with a modest average expression level, such a peptide will have lower levels of expression than either a higher expressing dynamic peptide or a constitutively expressed peptide with the same mean expression. Such peptides are more likely to be dropped from analyses. The prevalence of missing data also leads to the counterintuitive situation that as the number of replicates increases, fewer and fewer peptides are measured in all samples and available for analysis with best-in-class methods. It is therefore critical to recover at least those peptides missing in only a small fraction of observations before analysis.

Batch effects are the second challenge that larger scale experiments exacerbate. A systematic study of variance in iTRAQ experiments showed that measurements of abundance in most proteins can be expected to exhibit 10–20% relative errors, going as high as 40% for low weight proteins with few detected peptides (Hultin-Rosenberg *et al.*, 2013). Most modern experimental designs include pooled control samples; however, such pooled controls account only for variability introduced by MS runs and not for bias introduced by sample handling, which has been shown to contribute similar amounts of variability (Piehowski *et al.*, 2013). Even for the best Jonckheere–Terpstra–Kendall based methods, the addition of 10–20% error meaningfully diminishes the ability to classify circadian expression patterns (Hutchison *et al.*, 2015). Batch effects are a greater problem as the number of samples increases, first because the number of batches must increase accordingly making relative comparisons in the expression of a given peptide more challenging, and second, because more peptides will be influenced by at least one batch effect. It is therefore critical to be able to model and remove bias trends when analyzing large-scale proteomics data.

Here we provide a toolset developed in response to these issues, Learning and Imputation for Mass-spec Bias Reduction (LIMBR). LIMBR employs a K nearest neighbors (KNN) based imputation strategy that is both non-parametric and has been shown to be highly effective in the context of proteomics data (Wang *et al.*, 2017b). LIMBR's bias trend modeling procedure is based on surrogate variable analysis (SVA), a proven bias modeling algorithm initially devised for micro-array data (Leek, 2007; Leek *et al.*, 2012; Leek and Storey, 2007) that has been used extensively and successfully in many systems (Benjamin *et al.*, 2017; Lopez *et al.*, 2014; Parsana *et al.*, 2017; Tsang *et al.*, 2014; Wang *et al.*, 2016) and has also spawned many adaptations (Chakraborty *et al.*, 2013; Karpievitch *et al.*, 2009; Parker *et al.*, 2014). LIMBR improves SVA with optimizations specific to large-scale MS experimental designs, particularly in its handling of general and circadian time courses. We report extensive tests on both simulated and experimental data, documenting that LIMBR accurately removes bias trends without requiring information about sample preparation and handling. LIMBR requires that sample pre-processing and MS run batches be randomized, so that bias trends are separable from biological signal.

2 Approach

The need for LIMBR software arose out of technical challenges presented by a very large circadian proteomics time course conducted in *Neurospora crassa*. In brief, cultures of the fungus were synchronized to the same circadian time and released into a circadian free-run; samples were then collected at regular intervals and iTRAQ based MS was performed. This dataset has measurements for 48h with a 2h resolution. Two genotypes are measured and all assays are performed with three replicates for a total of 144 experimental samples. With pooled controls, 192 total MS samples were assayed.

Because of the number of replicates and total conditions, only 24% of the total assayed peptides were detected in all samples. Approximately one-third of peptides were missing in almost all samples, one-third were present in the majority of samples, and the remainder were roughly uniformly distributed across intermediate levels of missingness (Fig. 1a). For downstream analysis methods that perform complete-case analysis, 76% of detected peptides would have been discarded. An illustration of the decrease in the number of peptides considered in a complete-case analysis as the size of the experiment is increased is shown in Figure 1b.

Bias trends that correspond to batch effects were also evident in this large experiment. We sorted samples by time point (Fig. 1c) and did not observe evident patterns associated with time, which we would have expected because of the nature of circadian regulation in which expression exhibits gradual cyclical changes. For a dataset without batch effects, we would observe strong agreement between replicates and a larger checkerboard pattern, as day-phase protein expression should have been similar and generally opposed to night-phase expression. Instead, arranging samples by preparation set, which indicated the batches to which samples were randomized for purification, digestion and other processing prior to the M/S run, revealed pronounced patterns (Fig. 1d). In a dataset without batch effects we would expect to see no discernible structure in this plot.

LIMBR was developed to model and remove bias trends in this type of data (Fig. 2). After missing data were imputed, batch effects were modeled using an SVA-based method. By modeling replicate and time series correlations, we produced a matrix of residuals containing the unknown batch effects. We modeled these batch effects with singular value decomposition (SVD) to produce linearly independent bias trend. By permuting the residual matrix, we estimated the significance for each bias trend and removed those passing a user-specified significance threshold. The LIMBR software implements these methods alongside imputation methods to provide a toolkit for the analysis of large proteomic experiments.

3 Materials and methods

3.1 Imputation of missing values

In the proof-of-concept dataset only 32 561 peptides were measured in all samples; however, 137 340 peptides were measured in at least one sample. We sought to increase the coverage of samples to avoid the situation where increased numbers of replicates led to fewer analyzable peptides. For this, we implemented KNN imputation, which has been found to be fast and reliable at the imputation levels and dataset sizes relevant to such proteomics experiments (Batista and Monard, 2001; Mandel *et al.*, 2015; Troyanskaya *et al.*, 2001; Wasito and Mirkin, 2005). The KNN algorithm imputes missing data by finding the K nearest data points with complete data for a given data point and imputes the missing value as the average of the nearby points' values. Where n is the number of observed peptides, m is the number of MS experiments, $\mathcal{D}^{n \times m}$ is our dataset and d is a

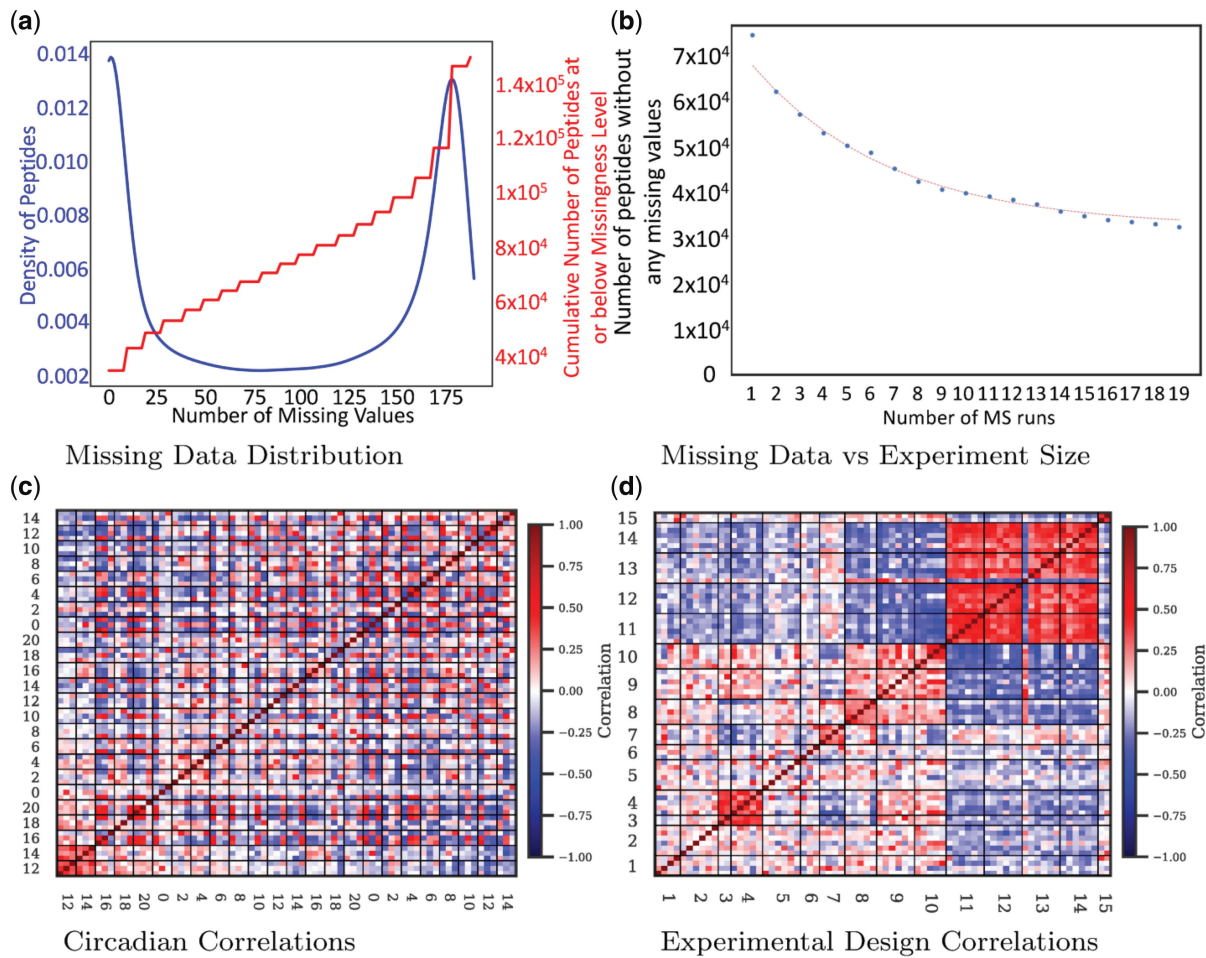


Fig. 1. Illustration of dataset issues addressed by LIMBR. (a) Distribution of missing data is shown in blue with the number of missing values on the x axis and the kernel density estimate of peptides with that number of missing values on the left axis. The number of peptides observed with fewer than the given number of missing values on the x axis is shown in red with corresponding values on the right axis. (b) Number of peptides with no missing values (complete cases) as the number of MS runs is increased. For a complete-case analysis, the number of peptides considered decreases exponentially as the size of the experiment is increased (exponential fit shown in red). (c) Correlation matrix of WT experimental data. The matrix is sorted by the time the sample was taken and the triplicate structure is indicated by the 3×3 black boxes. Numbers indicate the circadian time of the observation. Minimal correlation is observed between replicates and between samples with matched circadian times. (d) The same correlation data as in (c) is here shown sorted by the sample preparation set. Strong correlations are seen within sets, indicating that sample preparation was a major contributor to the observed signal

datum (all observations on a peptide), for $d \in \mathcal{D}$ we found the KNN of d by Euclidean distance. We then constructed the matrix $\mathcal{N}^{K \times m}$ with rows being nearest neighbors of d . Finally we imputed the missing values in d as the mean of the corresponding columns of \mathcal{N} (second panel Fig. 2). We report results using 10 nearest neighbors, which have been shown to be appropriate for datasets of roughly this size (Wang et al., 2017b). In light of the triplicate experimental design employed for this data and the fact that KNN has been shown to function effectively on biological datasets without replicates up to 35% missingness, we imputed peptides missing in fewer than 30% of samples (Mandel et al., 2015). Both parameters may be set at runtime in LIMBR by the user. An exploration of the impact of parameter selection on the analysis of simulated data is provided in Supplementary Figure S1. Values for k in the range of 5–15, along with missingness thresholds in the 30–40% range are recommended.

Even with this conservative imputation threshold, applying KNN increased the size of the proof-of-concept dataset by $\sim 70\%$ (32 561–56 353 peptides), substantially improving both the number of proteins detected and eliminating the decrease in analyzable peptides with increasing numbers of replicates.

Algorithm 1 KNN Imputation

procedure KNN(\mathcal{D} , K)

$\mathcal{D}^{n \times m} \leftarrow$ Our dataset

$n \leftarrow$ number of observed peptides

$m \leftarrow$ number of MS experiments (timepoints \times replicates)

$d \leftarrow$ a datum (all observations on a peptide)

for $d \in \mathcal{D}$

 find KNN of d by Euclidean distance

 construct the matrix $\mathcal{N}^{K \times m}$ with rows being nearest neighbors of d

 impute missing values in d as mean of corresponding columns of \mathcal{N}

end

return complete \mathcal{D}

end procedure

3.2 Modeling of bias trends

MS experiments are subject to complex batch effects which can contribute a sizeable fraction of the total variability recorded in an experiment. Because circadian time courses require many time points and replicates and the biological trends are relatively subtle, these bias trends are often stronger than the biological signal in such datasets.

Algorithm 2 SVD Bias Modeling

procedure SVD_BIAS(\mathcal{D}, α)

$\mathcal{D}^{n \times m} \leftarrow$ Our dataset

$n \leftarrow$ number of observed peptides

$m \leftarrow$ number of MS experiments (timepoints \times replicates)

$\alpha \leftarrow$ bias trend significance threshold

$S \leftarrow$ Sample timepoints

$c_t \leftarrow$ correlation threshold

calculate the correlation to the primary variable of interest for each peptide c

form a reduced data matrix \mathcal{D}_r of peptides for which $c < c_t$

fit the Lowess model $\mathcal{D}_r = \beta S^T + E$

Calculate the residual matrix as $\hat{E} = \mathcal{D}_r - \beta S^T$

Calculate the SVD of the residual matrix $\hat{E} = UDV^T$

With d_l as the l^{th} singular value, for right singular value $k = 1, \dots, n$ calculate the observed test statistic as:

$$T_k = \frac{d_k^2}{\sum_{l=1}^n d_l^2}$$

Permute each row of the matrix \hat{E} independently to form a matrix \hat{E}^*

Calculate the SVD of the matrix $\hat{E}^* = U_0 D_0 V_0^T$

For right singular value k calculate the null statistic:

$$T_k^0 = \frac{d_{0k}^2}{\sum_{l=1}^n d_{0l}^2}$$

Repeat calculation of the null statistic B times

Calculate the P -value for right singular vector k as:

$$p_k = \frac{\#\{T_k^{0b} > T_k; b=1, \dots, B\}}{B}$$

Estimate number of significant trends sb as:

$$sb = \sum_{k=1}^n \mathcal{I}(p_k \leq \alpha) \text{ While } p_k \leq \alpha$$

For each significant trend v_k , regress v_k on each row of \mathcal{D}_r calculating a P -value for their association

Estimate the number of truly associated features as $\hat{m}_1 = [(1 - \hat{\pi}_0) \times m]$ and form a subset of features with the \hat{m}_1 smallest P -values

Calculate the right singular vectors of the reduced subsetted matrix as v_j^r for $j = 1, \dots, n$

Estimate the surrogate variable j^* as the eigengene of the reduced subset matrix most correlated with the corresponding residual eigengene

$$j^* = \operatorname{argmax}_{1 < j < n} \operatorname{cor}(v_k, v_j^r)$$

Set $\hat{G}_k = v_{j^*}^r$

Find least squares solution for M where $\mathcal{D} = M \times \hat{G}$

return $\mathcal{D} - M \times \hat{G}$

end procedure

We first recognize that some peptides will be more closely correlated with the process of interest to investigators and that our modeling of bias trends can be improved by excluding such peptides (Leek, 2007). In the case of circadian experiments we use the ratio of autocorrelations at 12 and 24 h as a measure of correlation (c) to our variable of interest (circadian expression). The correlation threshold at which peptides are not analyzed for batch effects (c_t), can be specified by the user and should be set relatively low (we use 25%) to screen out only those peptides where the signal of interest greatly outweighs any batch effects. In this way, we formed the reduced data matrix, \mathcal{D}_r of peptides for which $c < c_t$.

Next, we proceeded to separating batch effects from our dataset. For a time series experiment, we expect some simple correlations to describe the data. Namely, we expect that replicates will agree with one another and that adjacent timepoints will be more similar than distant ones. Given a series of timepoints for each observation S , this structure is captured mathematically by a LOWESS model which we fitted to our data:

$$\mathcal{D}_r = \beta S^T + \hat{E}$$

We then analyzed the residuals, \hat{E} from this model calculated as:

$$\hat{E} = \mathcal{D}_r - \beta S^T$$

These residuals represent variability in the dataset which did not agree with the experimental design and should therefore contain any batch effects. Since completely eliminating all of the model residuals would have drastically underestimated our uncertainty, we chose to model batch effects within the residual matrix (Jaffe et al., 2015; Leek, 2007). We did this by breaking the residual matrix \hat{E} , down into a set of linearly independent bias trends with SVD $\hat{E} = UDV^T$. SVD has been used to model bias trends in micro-array, RNAseq and M/S data (Karpievitch et al., 2009; Parsana et al., 2017; Tsang et al., 2014). Once we had broken down the residual matrix into bias trends, we needed to determine which were likely to have been contributed by batch effects and should therefore be removed from the dataset. We did this by estimating which of these trends contributes more variability than we would expect at random. By repeatedly permuting the residual matrix and performing SVD, we estimated the null distribution of variance explained by each singular vector and thereby assessed the 'significance' of individual bias trends, removing only those which explain more variance than we would expect by chance at a user-specified threshold P . We designated bias trends with $P < 0.05$ as significant. With d_l as the l^{th} singular value, for right singular value $k = 1, \dots, n$ we calculated the observed test statistic T_k as:

$$T_k = \frac{d_k^2}{\sum_{l=1}^n d_l^2}$$

We then permuted each row of the matrix \hat{E} independently to form a matrix \hat{E}^* . We calculated the singular values and the null statistic T_k^0 in similar fashion ($\hat{E}^* = U_0 D_0 V_0^T$ and $T_k^0 = \frac{d_{0k}^2}{\sum_{l=1}^n d_{0l}^2}$). These calculations were repeated B times and the P -value for right singular vector k was calculated as:

$$p_k = \frac{\#\{T_k^{0b} > T_k; b = 1, \dots, B\}}{B}$$

Having calculated the significance of each bias trend, we could have simply removed those passing a given threshold. Instead we chose to further remove our modeling choices from the bias trends

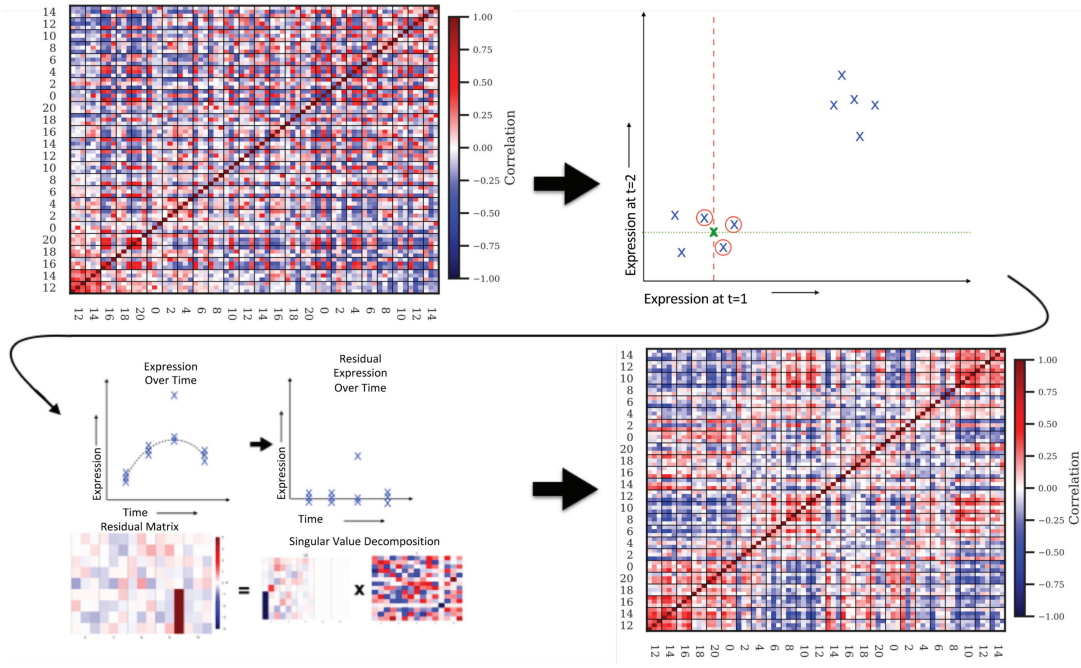


Fig. 2. Algorithm outline. The outline of processes performed by the LIMBR algorithm is diagrammed, beginning with raw data heavily influenced by bias trends, proceeding to imputation of missing data, modeling of bias trends and the final removal of those bias trends producing a denoised dataset

we calculated (Leek, 2007). We did this by moving through the trends in order adding one for each trend where $p_k < \alpha$ and stopping when this condition is no longer met, thereby estimating the number of significant trends sb as:

$$sb = \sum_{k=1}^n \mathcal{I}(p_k \leq \alpha) \text{ While } p_k \leq \alpha$$

We then regressed each significant trend \mathbf{v}_k against the rows of \mathcal{D}_r calculating a P -value for their association. Given an estimate of the background uniform distribution of P -values $\hat{\pi}_0$, the number of truly associated peptides for each trend, \hat{m}_1 , was estimated as:

$$\hat{m}_1 = [(1 - \hat{\pi}_0) \times m]$$

A subset of the peptides with the \hat{m}_1 smallest P -values for that trend was then formed. The batch effect j^* was then modeled as the right singular vector of the reduced subset matrix most correlated with the bias trend being modeled:

$$j^* = \operatorname{argmax}_{1 \leq j \leq n} \operatorname{cor}(\mathbf{v}_k, \mathbf{v}_j^r)$$

where the right singular vectors of the reduced subsetted matrix are:

$$\mathbf{v}_j^r j = 1, \text{ for } \dots, n$$

The matrix of batch effects $\hat{\mathbf{G}}_k$ was then taken to be $\hat{\mathbf{G}}_k = \mathbf{v}_j^r$ and these batch effects were removed by returning $\mathcal{D} - \mathbf{M} \times \hat{\mathbf{G}}$ where \mathbf{M} is found by a least squares solution of $\mathcal{D} = \mathbf{M} \times \hat{\mathbf{G}}$. By first reducing the data considered to only those peptides not correlated with our variable of interest, we avoid learning circadian expression patterns as a bias trend. We then quantify the variability contributed by the linearly independent bias trends produced by SVD of the residual matrix. By repeatedly randomizing the residual matrix and reperforming those calculations we estimate the likelihood of each bias trend contributing the variability observed by chance and can label some bias trends as significant based on this estimate. We then

reduced the impact of the already limited modeling assumptions we had employed by regressing the significant trends against the original data and calculating our bias trends from a subset of that data most correlated with the original trends. In this way, we were able to infer and remove batch effects without knowledge of the sample handling steps that produced them.

3.3 Simulation studies

To verify the effectiveness of the batch effect modeling and removal procedure, we conducted simulation studies. First, we generated synthetic data with 50% circadian peptides taken from sine waves equally distributed between opposite phases and added Gaussian random noise. The other 50% of peptides were non-circadian and generated solely from Gaussian random noise. We added randomly generated batch effects to these data and processed the resulting datasets with LIMBR. We analyzed the datasets before noise (Fig. 3a panel 1), after the addition of noise (Fig. 3a panel 2) and after processing with LIMBR (Fig. 3a panel 3) with eJTK-cycle. We compared the resulting circadian classification accuracies using ROC curves.

In order to provide a more stringent evaluation, we conducted an additional set of simulations in which three randomly generated batch effects were applied to the simulated data. Since the batch effects were applied at random, this resulted in many fewer unaffected peptides ($\sim 12.5\%$ versus $\sim 50\%$) and many more combinations of bias trends (7 versus 1) making this data much more challenging to analyze. Both the single batch effect and triple batch effect simulations were repeated a total of 20 times and the results for the more challenging simulations are visualized in Figure 3b and c. As an additional point of comparison, we processed the triple batch effect simulations with eigenMS the current state of the art SVA-based software for removing batch effects in MS data (Karpievitch et al., 2009). The addition of three bias trends completely obliterated circadian classification accuracy. Processing with eigenMS, which

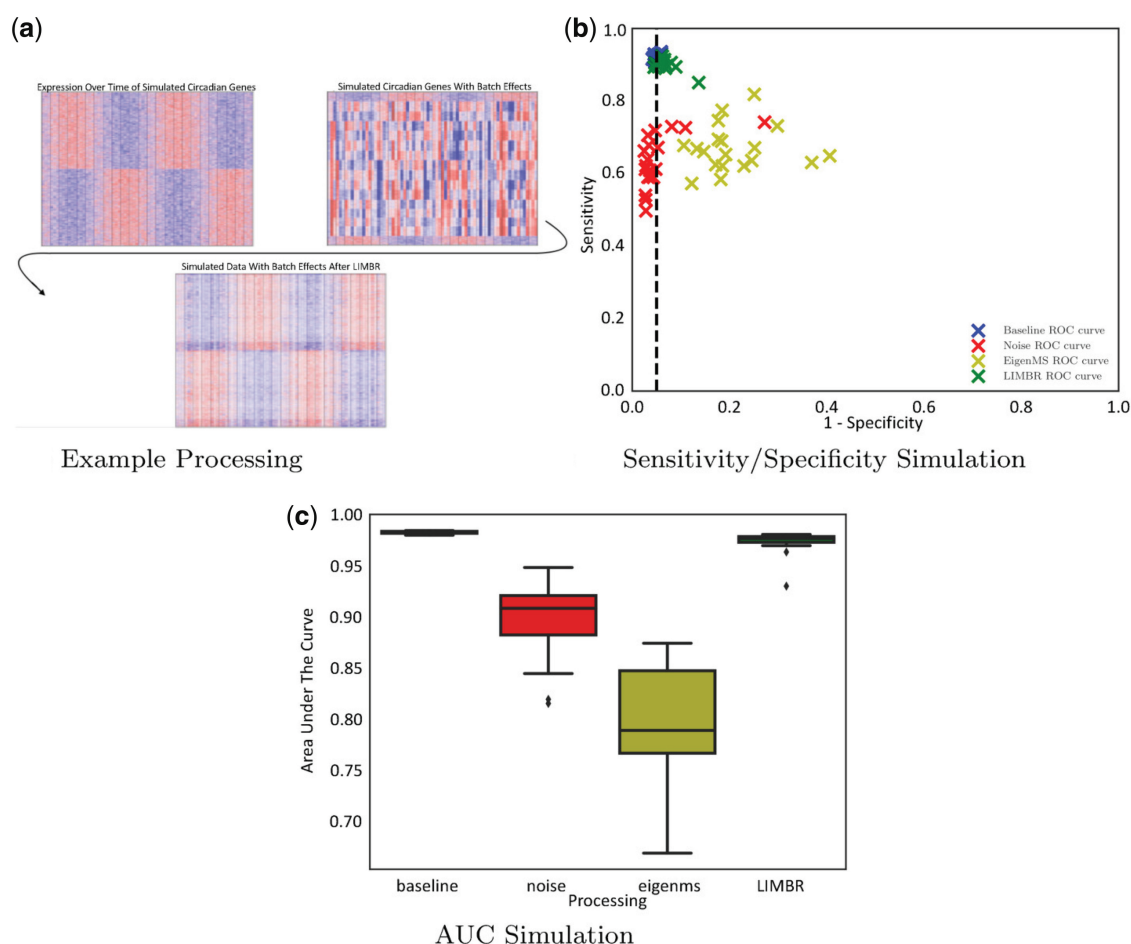


Fig. 3. Simulation results. (a) Example results showing initial circadian signal, data with the addition of three random bias trends and data after the application of LIMBR. (b) Representative sensitivity and specificity data showing the effectiveness of circadian classification for 20 rounds of simulation using a $P < 0.05$ threshold in eJTK for each of the three stages of data processing shown in (a) with the alternative eigenMS method shown for comparison. (c) Area under the receiver operating characteristic curves for the same 20 simulations shown in (b)

cannot exploit the time series structure of the dataset was largely ineffective. Processing with LIMBR; however, almost completely recovered classification accuracy. These simulations indicated that LIMBR was modeling and removing unknown bias trends accurately and reproducibly.

3.4 Application to biological data

Since the two genotypes were randomized and run on the M/S together, the batch effects observed in each dataset should be the same. Additionally, while we do expect some differences in expression between the two, globally expression should be similar. When analyzing the data however, we applied LIMBR to the time courses from each of the two genotypes independently. Since information was not shared with the algorithm between these runs, comparing the results of these analyses allows us to examine the reproducibility of our techniques.

First, we compared the bias trends identified by LIMBR with the sample handling steps performed prior to M/S. We found that for both genotype time courses, the bias trends correlated well with sample preparation steps (Fig. 4a and b). The primary sources of bias in sample preparation were strongly correlated with sample preparation and Tandem Mass Tag (TMT) sets indicating that these steps made major contributions to the observed bias. This fits well with previous reports that peptide digestion and charge state are

major sources of bias in M/S quantification (Piehowski *et al.*, 2013; Rudnick *et al.*, 2014). We then compared the bias trends found in each genotype, finding that LIMBR had independently reproduced analogous bias trends for each dataset (Fig. 4c).

After these final evaluations of the software, we analyzed the results of circadian classification for our LIMBR processed data and found 324 circadian proteins out of 4754 analyzed proteins (7%) with a predominantly bimodal distribution of phases (Fig. 5a). Analysis of the sample correlation matrix for the wild type (WT) dataset revealed a much stronger agreement between replicates along with the diurnal checkerboard pattern (Fig. 5b). When we focused the analysis on only circadian proteins, these patterns were even more pronounced (Fig. 5c).

We compared the predicted phases of proteins identified as circadian in both genotypes and found a tight correspondence between the two (Fig. 5d). A detailed discussion of the biological implications of these analyses is underway (Hurley *et al.*, in preparation).

4 Discussion

The combination of KNN based imputation and SVD based error modeling improve the results of time-course proteomics analysis markedly. These techniques, which we implement into the LIMBR software package, allow for the recovery of circadian signals in

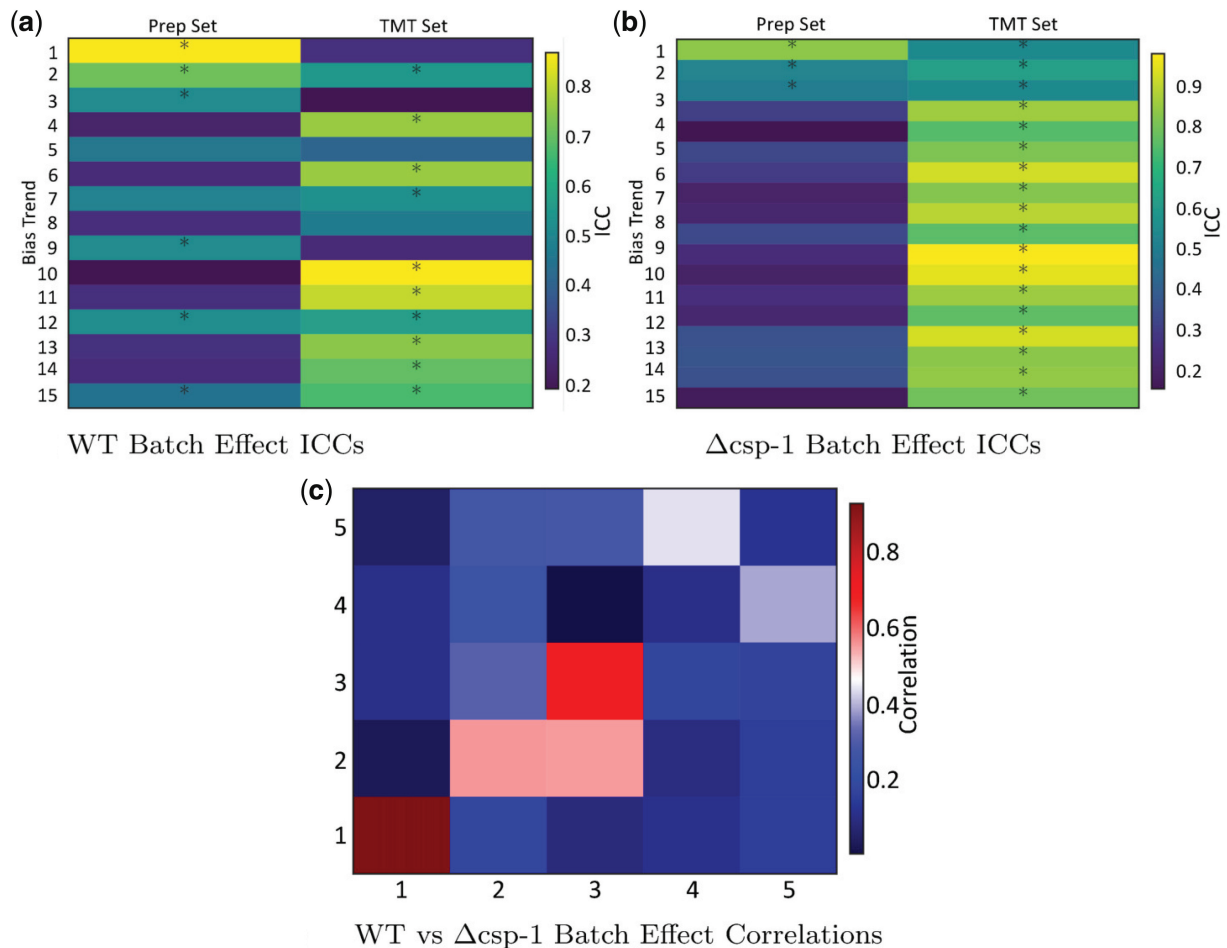


Fig. 4. Analysis of batch effects identified in biological data. Interclass correlation coefficients between batch effects identified as significant and the prep set and TMT set of each sample in WT (a) and Δ csp-1 (b) are shown with * indicating significance at the $P < 0.001$ level for a one way F-test. Almost all bias trends are significantly associated with sample preparation or TMT labeling sets. In some cases, due to incomplete randomization, bias trends are associated with both. (c) Correlation matrix of the top 5 bias trends as measured by explained variance in the WT and delta csp-1 experiments. The top 3 bias trends independently identified in each experiment are quite similar, and trends 4 and 5 seem to be flipped in their rankings but still correspond closely

simulated and biological datasets for which no such signal is initially observable.

Simulation studies reveal that LIMBR has high sensitivity and specificity along with effective recovery of phase. LIMBR also produces superior results to other SVA-based methods for the removal of batch effects when applied to time series data. This increase in performance can be attributed to the reduction and subsetting steps, but especially to the LOWESS based calculation of residuals. The reduction step serves to prevent learning signal of interest as a batch effect, while the subsetting step helps to draw batch effects more directly from the data with less influence from our modeling decisions. The LOWESS-based calculation of residuals allows for an improved estimate of batch effects by incorporating experimental structure which is lost when processing with eigenMS. While the performance improvements of LIMBR over eigenMS for the simulated data in this study are drastic, this difference can be attributed to the much greater amount of information LIMBR has access to for time series data. LIMBR should therefore be expected to offer larger improvements over other methods for time series data than for block experimental designs where LIMBR's improvements in residual calculations do not come into play. Even in the case of a smaller block design experiment with relatively fewer and smaller batch

effects however, LIMBR still offers the added benefit of integrated imputation of missing values.

In the case of biological data, we also show that LIMBR produces effective and consistent results. When comparing between proteomics datasets for two genotypes, we see that the inferred batch effects match well with experimental parameters. This indicates that the batch effects being removed are well grounded in reality. Additionally, batch effects inferred from independent processing of the two genotypes match well with one another, indicating the consistency of LIMBR. Finally, the phases of circadian proteins found in these independently processed datasets closely agree. The consistency of batch effects with experimental parameters and between LIMBR runs along with the agreement of phases of circadian genes between LIMBR runs indicates the consistency of our methods, while the recovery of circadian signal from a dataset otherwise dominated by batch effects speaks to the efficacy of these techniques.

While it is possible for the removal of bias trends such as batch effects to inflate false positive rates in studies of differential expression, these effects should only be a concern when there is a highly uneven distribution of samples to batches (Nygaard et al., 2016). Additionally, the application of bias modeling techniques has been

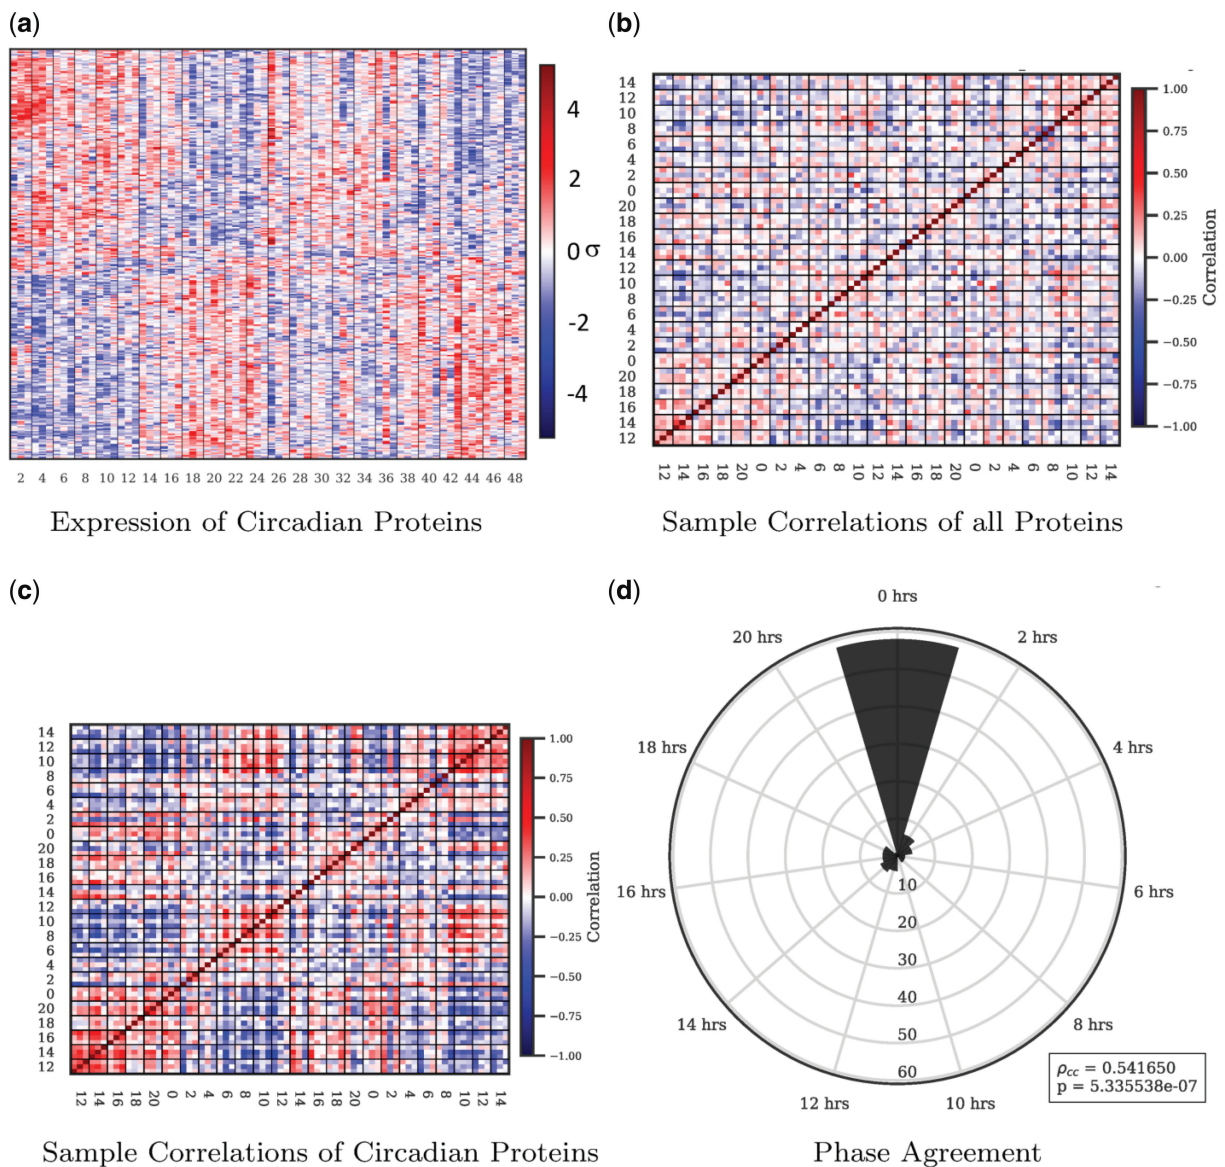


Fig. 5. Results of application of LIMBR to biological data. (a) Expression profiles of genes identified as circadian in the WT dataset after processing. (b) Correlation matrix of WT dataset after processing, showing increased similarity between replicates and checker boarding pattern of similarity between repeated circadian times (compare to Fig. 1b). (c) Correlation matrix of WT dataset generated only from proteins identified as circadian showing pronounced checker boarding similarity between repeated circadian times. (d) For proteins identified as circadian in both genotypes the distribution of differences in predicted phases between the genotypes is shown. A strong peak at 0 indicates that the predicted phases of these genes match well between the independent processing runs

shown to increase the sensitivity and specificity of differential expression studies by removing bias trends and to be a net positive when applied to RNAseq data (Li *et al.*, 2014). One limitation of any strategy for imputation of MS data is the generation of novel batch effects. It is possible for batch effects to bleed through from peptides with complete observations to imputed peptides. Because only a subset of sample abundances will be imputed for these peptides, it is possible for only a subset of the original batch effect to be transferred, thereby generating a novel batch effect. While this is an unavoidable limitation of imputation of MS data, any batch effects generated in this manner would be removed by LIMBR with the same efficiency as endogenous batch effects. The algorithm employed here is based on earlier work but is notable both for being a modern implementation in python and for implementing a more advanced reduced subset method originally proposed by Leek in

tandem with LOWESS-based calculation of residuals (Leek *et al.*, 2012; Leek and Storey, 2007b).

5 Conclusion

The application of LIMBR to this dataset demonstrates the potential these techniques hold for the analysis of large-scale circadian proteomics time courses. In three separate prior studies of circadian proteomics (Mauvoisin *et al.*, 2014; Robles *et al.*, 2014, 2017), the issues of lost data resulted in circadian data for <200 proteins, even after accepting time courses in which many observations were missing and where rhythmicity was called with an FDR of 0.25–0.3. With LIMBR, we can apply a much lower FDR of 0.05, reject datasets in which more than 30% of the points are missing, and still

identify many more circadianly regulated proteins. While this highlights the strengths of the technique, it is also worth noting the limitations, specifically that a high number of time points and replicates are critical to effectively model batch effects. Additionally, randomization of samples to MS runs and processing steps must be performed independently to ensure the separability of biological signal and batch effects. Finally, we recommend that normalization with pooled controls, which is optionally incorporated into the LIMBR pipeline, be used for the MS runs where possible. While such large experiments require the investment of additional resources, their importance to this type of study cannot be overemphasized. Although there is still room for improvement in the design of circadian proteomics experiments both *in vitro* and *in silico*, this work makes clear the need for analysis techniques and software for such datasets.

Funding

Funding for A.M.C. was provided by National Institutes of Health grants [R01GM083336] and [1R35GM118022] to J.J.L., [1R35GM118021, 1U01EB022546] to J.C.D. and through an Albert J. Ryan Fellowship. Funding for C.S.G. was provided by a grant from the Gordon and Betty Moore Foundation [GBMF 4552].

Conflict of Interest: none declared.

References

- Batista, G.E. and Monard, M.C. (2001) A study of K- nearest neighbour as a model-based method to treat missing data. In *Proceedings of the Argentine Symposium on Artificial Intelligence*, Vol. 30, pp. 1–9.
- Benjamin, J.S. et al. (2017) A ketogenic diet rescues hippocampal memory defects in a mouse model of Kabuki syndrome. *Proc. Natl. Acad. Sci. USA*, **114**, 125–130.
- Chakraborty, S. et al. (2013) svapls: an R package to correct for hidden factors of variability in gene expression studies. *BMC Bioinformatics*, **14**, 236–242.
- Chick, J.M. et al. (2016) Defining the consequences of genetic variation on a proteome-wide scale. *Nature*, **534**, 500–505.
- Gelman, A. and Hill, J. (2007) Missing-data imputation. In: *Data Analysis Using Regression and Multilevel/Hierarchical Models*. 1st edn. Cambridge University Press, Cambridge, pp. 529–543.
- Hughes, M.E. et al. (2010) JTK_CYCLE: an efficient nonparametric algorithm for detecting rhythmic components in genome-scale data sets. *J. Biol. Rhythms*, **25**, 372–380.
- Hultin-Rosenberg, L. et al. (2013) Defining, comparing, and improving iTRAQ quantification in mass spectrometry proteomics data. *Mol. Cell. Proteomics*, **12**, 2021–2031.
- Hutchison, A.L. et al. (2015) Improved Statistical Methods Enable Greater Sensitivity in Rhythm Detection for Genome-Wide Data. *PLoS Comput. Bio.*, **11**, 1–29.
- Jaffe, A.E. et al. (2015) Practical impacts of genomic data “cleaning” on biological discovery using surrogate variable analysis. *BMC Bioinformatics*, **16**, 372–381.
- Karpievitch, Y.V. et al. (2009) Normalization of peak intensities in bottom-up MS-based proteomics using singular value decomposition. *Bioinformatics*, **25**, 2573–2580.
- Karpievitch, Y.V. et al. (2012) Normalization and missing value imputation for label-free LC-MS analysis. *BMC Bioinformatics*, **13**, S5.
- Leek, J.T. (2007) Surrogate variable analysis. PhD Thesis, University of Washington, Seattle, WA, USA.
- Leek, J.T. and Storey, J.D. (2007) Capturing heterogeneity in gene expression studies by surrogate variable analysis. *PLoS Genet.*, **3**, 1724–1735.
- Leek, J.T. et al. (2012) The sva package for removing batch effects and other unwanted variation in high-throughput experiments. *Bioinformatics*, **28**, 882–883.
- Li, S. et al. (2014) Detecting and correcting systematic variation in large-scale RNA sequencing data. *Nature*, **32**, 888–895.
- Lopez, J.P. et al. (2014) miR-1202 is a primate-specific and brain-enriched microRNA involved in major depression and antidepressant treatment. *Nat. Med.*, **20**, 764–768.
- Mallick, P. et al. (2007) Computational prediction of proteotypic peptides for quantitative proteomics. *Nat. Biotechnol.*, **25**, 125–131.
- Mandel, J. et al. (2015) A Comparison of Six Methods for Missing Data Imputation. *J. Biom. Biostat.*, **6**, 1–6.
- Mauvoisin, D. et al. (2014) Circadian clock-dependent and -independent rhythmic proteomes implement distinct diurnal functions in mouse liver. *Proc. Natl. Acad. Sci. USA*, **111**, 167–172.
- Nygaard, V. et al. (2016) Methods that remove batch effects while retaining group differences may lead to exaggerated confidence in downstream analyses. *Biostatistics*, **17**, 29–39.
- Parker, H.S. et al. (2014) Preserving biological heterogeneity with a permuted surrogate variable analysis for genomics batch correction. *Bioinformatics*, **30**, 2757–2763.
- Parsana, P. et al. (2017) Identifying global expression patterns and key regulators in epithelial to mesenchymal transition through multi-study integration. *BMC Cancer*, **17**, 447–460.
- Piehowski, P.D. et al. (2013) Sources of Technical Variability in Quantitative LC-MS Proteomics: human Brain Tissue Sample Analysis. *J. Proteome Res.*, **12**, 2128–2137.
- Robles, M.S. et al. (2014) In-vivo quantitative proteomics reveals a key contribution of post-transcriptional mechanisms to the circadian regulation of liver metabolism. *PLoS Genet.*, **10**, 15.
- Robles, M.S. et al. (2017) Phosphorylation Is a Central Mechanism for Circadian Control of Metabolism and Physiology. *Cell Metab.*, **25**, 118–127.
- Rudnick, P.A. et al. (2014) Improved normalization of systematic biases affecting ion current measurements in label-free proteomics data. *Mol. Cell. Proteomics*, **13**, 1341–1351.
- Tabb, D.L. et al. (2003) Statistical characterization of ion trap tandem mass spectra from doubly charged tryptic peptides. *Anal. Chem.*, **75**, 1155–1163.
- Troyanskaya, O. et al. (2001) Missing value estimation methods for DNA microarrays. *Bioinformatics*, **17**, 520–525.
- Tsang, J.S. et al. (2014) Global analyses of human immune variation reveal baseline predictors of postvaccination responses. *Cell*, **157**, 499–513.
- Wang, Z. et al. (2016) Extraction and analysis of signatures from the Gene Expression Omnibus by the crowd. *Nat. Commun.*, **7**, 1–11.
- Wang, J. et al. (2017a) Nuclear Proteomics Uncovers Diurnal Regulatory Landscapes in Mouse Liver. *Cell Metab.*, **25**, 102–117.
- Wang, J. et al. (2017b) In-depth method assessments of differentially expressed protein detection for shotgun proteomics data with missing values. *Sci. Rep.*, **7**, 273–278.
- Wasito, I. and Mirkin, B. (2005) Nearest neighbour approach in the least-squares data imputation algorithms. *Inf. Sci.*, **169**, 1–25.
- Weekes, M.P. et al. (2014) Quantitative Temporal Viromics: an Approach to Investigate Host-Pathogen Interaction. *Cell*, **157**, 1460–1472.