

Batch effects in large-scale proteomics analysis

Introduction

Recent advances in liquid chromatography mass spectrometry (LCMS)-based proteomics analysis have enabled the efficient profiling of thousands of proteins from single LCMS runs. The ability to run untargeted, high throughput LCMS experiments has opened the door to large-scale cohort studies for biomarker and drug target discovery¹. When conducting large-scale cohort studies, technical confounding can be introduced as samples are run across different MS instruments, LC columns, dates, and geographic locations. In order to integrate these samples across datasets for joint analyses, one needs to both diagnose this batch effect and apply methods to correct for it.

Here we compare methods for characterizing batch effects in proteomics data. We have evaluated the presence of a batch effect using multiple batch effect diagnosis methods, including Principal Components Analysis-based approaches, local-neighborhood diversity measures, and machine learning classifier-based methods. Next, we benchmark batch effect correction methods for protein abundance data. These include traditional methods often used in proteomics and genomics as well as a novel deep learning-based batch correction method we developed.

LCMS Data

We collect a batch-diverse dataset using Seer's Proteograph™ Product Suite. Our dataset includes **882 LCMS runs** (DDA) across:

- Two types of control plasma samples
- Three Seer Proteograph nanoparticles (NPs)
- Three LCMS instruments
- Eight LC columns.

Methods

We applied each method on our dataset to produce batch-corrected representations, which we then evaluated with batch characterization metrics. We applied baseline methods that are traditionally used in proteomics or other omics data analysis pipelines; like PCA, MSstats², and ComBat³; methods based on nearest neighbor matching like MNN⁴ and Scanorama⁵; and Harmony⁶ which is an iterative clustering + translating algorithm; and deep learning-based approaches like scVI⁷. We then applied our developed, deep learning approach by extending Domain Adversarial Neural Networks⁸ (DANN).

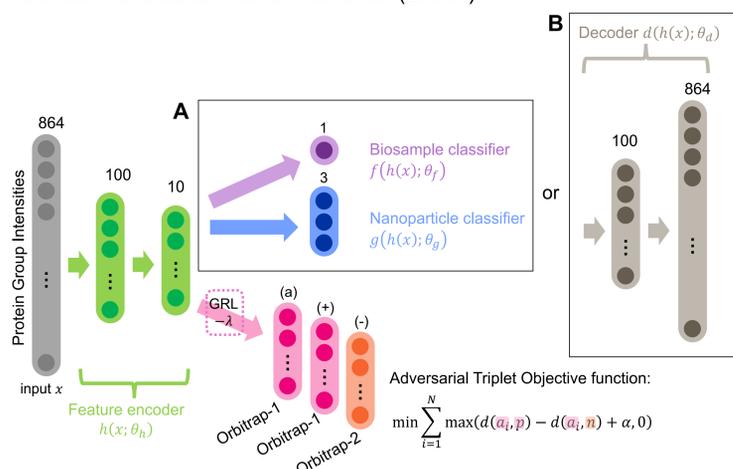


Figure 1. Adversarial neural network architectures to learn batch-invariant representations. Protein group intensity data is fed forward through a fully connected ReLU encoder stage (green), that is trained to perform poorly on a Triplet Loss which tries to discriminate technical batches. At the same time, this representation is trained to minimize either (A) two classification tasks as in DannClf, or (B) a reconstruction loss as in DannRecon.

Deep learning approaches to integrate large-scale datasets by correcting batch effects

Results

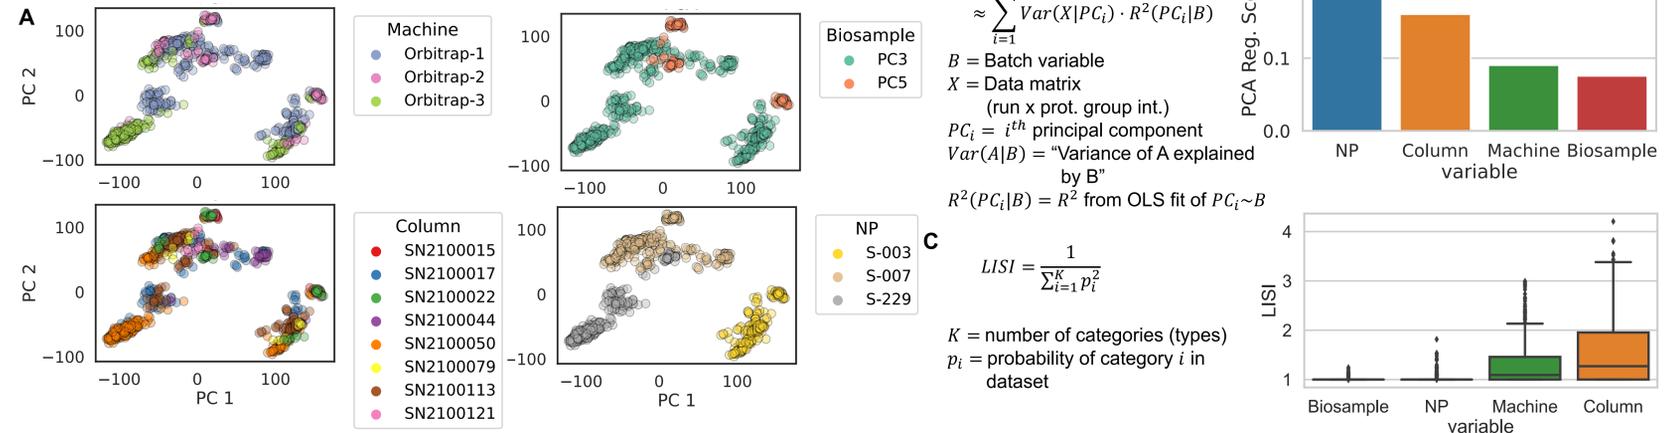


Figure 2. Characterization of batch effects in the data. A) PCA embeddings of protein group log intensities of each run, colored by four different covariates. B) Principal Components Regression, showing that batch variables (LC column and MS Instruments) are major contributors to the variance in the data, over analysis of these control plasma samples. C) Local Inverse Simpson's Index⁸ (LISI) score, measuring effective diversity of a label within small neighborhoods, which shows low levels of integration of batch variables. While (B) shows where signal resides in a data matrix, (C) the level of mixing, and is better suited for comparing batch effect correction methods.

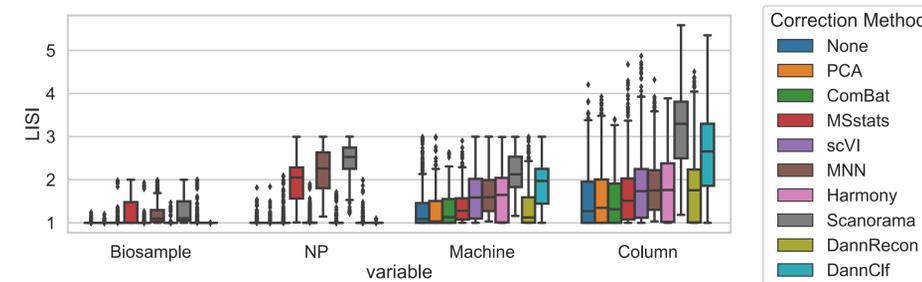


Figure 3. Dataset mixing and biological signal preservation of batch effect correction methods. LISI scores of all correction methods we applied. Though Scanorama mixed the best with respect to Machine and Column, it overmixed the biological variables. Our DannClf and DannRecon do not have this issue.

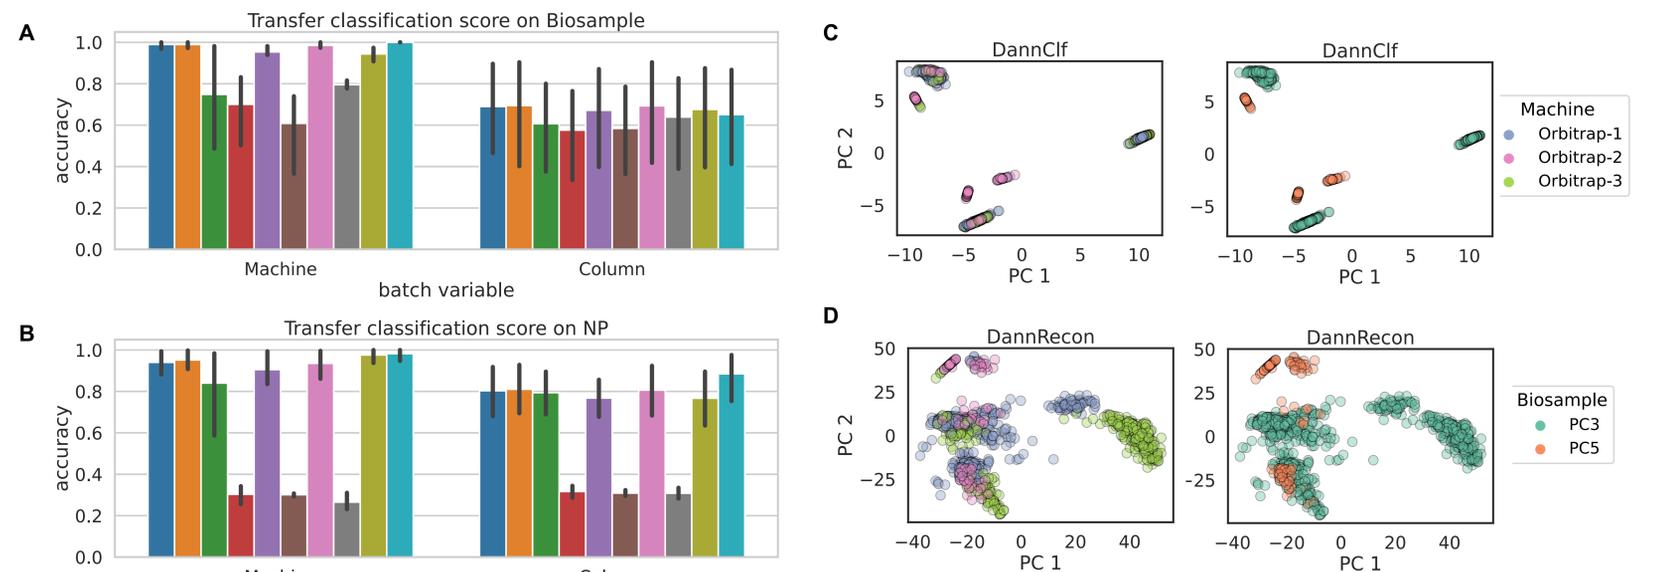


Figure 4. Quantitative and qualitative assessment of batch corrected representations for downstream tasks. A) Comparison of using batch corrected representations for classifying biological phenotypes across batches. For example: for MS Instruments (Machine), a KNN classifier is trained on Orbitrap-1 and Orbitrap-2 data to classify between the two control plasma samples (PC3 or PC5), and the test accuracy is assessed on Orbitrap-3. This is repeated for testing on Orbitrap-1 and Orbitrap-2. Note, this same procedure is repeated for Column. B) The same but the prediction task is to classify amongst the three Nanoparticles. C) PCA embeddings of the learned features from our DannClf and D) DannRecon models.

Conclusions

- Batch effects can contribute to a large amount of the noise in large-scale proteomics datasets, compared to biological variations. We have observed a large batch effect attributed to mass spectrometers and LC columns in our dataset.
- Deep learning-based approaches can learn to integrate diverse proteomics datasets well. Our extension of DANN can harmonize data across technical factors, while maintaining the fidelity of the biological signal in the data.
- While DannClf can harmonize the data well, the representations it learns are most useful for classification. Future work such as our unsupervised variant, DannRecon, may learn more general-purpose batch corrected representations.

References

1. Blume et al. Nat. Comm. (2020)
2. Choi et al. Bioinformatics (2014)
3. Johnson et al. Biostatistics. (2007)
4. Haghverdi et al. Nat. Biotechnol. (2018)
5. Hie et al. Nat. Biotechnol. (2019)
6. Korsunsky et al. Nat. Methods (2019)
7. Lopez et al. Nat. Methods (2018)
8. Ganin et al. J. Mach. Learn. Res. (2016)

