HHS Public Access

Author manuscript

J Proteome Res. Author manuscript; available in PMC 2020 April 05.

Published in final edited form as:

J Proteome Res. 2019 April 05; 18(4): 1902–1906. doi:10.1021/acs.jproteome.8b00954.

Detecting modifications in proteomics experiments with Param-Medic

Damon H. May¹, Kaipo Tamura¹, and William S. Noble^{1,2}

¹Department of Genome Sciences, University of Washington

²Paul G. Allen School of Computer Science and Engineering, University of Washington

Abstract

Searching tandem mass spectra against a peptide database requires accurate knowledge of various experimental parameters, including machine settings and details of the sample preparation protocol. In some cases, such as in re-analysis of public data sets, this experimental metadata may be missing or inaccurate. We describe a method for automatically inferring the presence of various types of modifications, including stable-isotope and isobaric labeling and tandem mass tags, as well as enrichment of phosphorylated peptides, directly from a given set of mass spectra. We demonstrate the sensitivity and specificity of the proposed approach, and we provide open source Python and C++ implementations in a new version of the software tool Param-Medic.

Graphical Abstract

Keywords

proteomics; mass spectrometry; tandem mass tags; isobaric labeling; phosphorylation; database search

1 Introduction

Properly searching shotgun proteomics data requires accurate metadata. For example, searching with suboptimal tolerances for precursor and fragment mass error can lead to a dramatic loss of sensitivity at a given false discovery rate.¹ Search parameters must also account for any appropriate modifications, such as isotopic labeling, or for enrichment of phosphopeptides. If such modifications are not included in the search, then modified peptides cannot be identified. On the other hand, if variable modifications are included in the search parameters but are not present in the sample, then the search space is unnecessarily expanded, and search sensitivity at a given false discovery rate threshold suffers.

In most cases, the appropriate parameters for a particular mass spectrometry run are known to the person performing the search. When a lab is searching data that it produced, presumably information is easily obtained about which runs were enriched for phosphorylated peptides or included labels for quantitation and which runs did not. However, in other situations this information may not be readily available. For example, when searching data downloaded from a public repository, the appropriate metadata may not be available, may not be trustworthy, or may not apply to all of the runs in an experiment. Repositories performing re-searches of submitted data may also lack this information. Also, a researcher may want to perform a quick assessment of their own data in order to determine whether an expected attribute (e.g., phospho-enrichment) is present to the extent that it will lead to detectable peptides containing the attribute.

We previously introduced Param-Medic,¹ a software tool that enables automatic inference of precursor and fragment mass accuracy parameters. In this work, we expand Param-Medic's functionality to automatically detect the presence of the following sample attributes: enrichment for phosphorylated peptides, SILAC labeling,² and two different isobaric labeling schemes (iTRAQ^{3,4} and TMT^{5,6}). We demonstrate Param-Medic's sensitivity and specificity on each of these tasks by using publicly available data. We make this functionality available in both open source implementations of Param-Medic: the standalone Python tool (<http://noble.gs.washington.edu/proj/param-medic>) and the implementation integrated into the Crux proteomics toolkit (<http://crux.ms>).

2 Methods

For each of three general types of attributes—phospho-enrichment, isobaric labels, and SILAC labels—we developed a statistic that aims to distinguish between runs that contain that attribute and runs that do not. Param-Medic computes each of these statistics and applies an empirically determined threshold to each one to determine the presence or absence of each attribute. We developed and evaluated each classifier using public data from

the MassIVE repository.⁷Below, we describe the statistics, the data sets, and the methods for setting decision thresholds and evaluating performance.

2.1 Detecting Phospho-enrichment

Studying phosphorylation in the proteome often entails enrichment of phosphorylated proteins.⁸In phospho-enriched runs, spectra generated by phosphorylated peptides often contain a fragment peak representing a neutral loss of H_3PO_4 ,⁹occurring at an m/z value corresponding to a singly-charged mass ~98 Da less than the precursor mass.

Accordingly, Param-Medic assesses whether a given run is phospho-enriched by quantifying the observed signal corresponding to a loss of H_3PO_4 . Specifically, Param-Medic assesses the proportion of total fragment signal accounted for by fragments falling in a nominal mass bin (with width 1.000507 Da, corresponding to the distance between the centers of two adjacent peptide mass clusters¹⁰) representing a 98-Da loss. This proportion is converted to a Z-score by comparing with the proportions of signal in eight nearby control bins (nominal m/z differences from precursor: -20, -15, -12, -10, 10, 12, 15, 20). If the Z-score exceeds an empirically determined threshold, then the run is determined to be phospho-enriched, and Param-Medic suggests that searches performed on the run use an appropriate variable modification of 79.966331 Da on serine, threonine and tyrosine residues.

2.2 Detecting Isobaric Labels

Isobaric labeling systems have become very popular for relative and absolute quantitation of peptides. Param-Medic now supports automatic detection of two labeling schemes: iTRAQ and TMT. Both iTRAQ and TMT use isobaric labels, with the intensity of reporter ions of fixed m/z ratio providing a measure of abundance in each sample.

To determine whether isobaric labels of a particular type (whether iTRAQ or TMT) are present, Param-Medic assesses the proportion of total MS/MS fragment signal that is accounted for by a list of nominal reporter ion m/z values for each type of label, as well as for a set of control m/z values (111, 112, 120, 122, 123, 124, 125, 133, and 134 m/z , chosen because they are nearby in m/z to the reporter ions of interest but do not appear to represent known modifications). Then, separately for iTRAQ and for TMT, Param-Medic calculates a t -statistic describing the extent to which the signal from all reporter ions for the label type is higher than the signal from the controls in order to determine whether the label responsible for the reporter ions is present. Finally, Param-Medic compares the t -statistic to an empirically determined threshold of 1.5 for iTRAQ labeling and 2.0 for TMT labeling.

Each labeling method comes in multiple varieties. For example, iTRAQ has both 4-plex and 8-plex versions, where the 4-plex reporter ions are a subset of the 8-plex reporter ions. Accordingly, if the 8-plex ions are determined to be present then the run is classified as 8-plex labeled rather than 4-plex labeled. Similarly, TMT comes in 2-plex, 6-plex and 10-plex varieties,¹¹and each variety's reporter ions are a subset of the next variety. 2-plex TMT is rarely used, and when 6-plex or 10-plex TMT is used, reporter ions representing as few as three of the available channels may be present. Therefore, if any 6-plex ions that are not also 2-plex ions are determined to be present, then the experiment is labeled "6-plex or 10-plex," rather than "2-plex."

Although the 2-plex and 6-plex reporter ions have different nominal masses, four of the 6-plex ion nominal masses are duplicated in the 10-plex reporter ions and are only distinguishable with a high-accuracy measurement. Therefore, if 6-plex ions are detected, then a separate test is run to determine whether 10-plex ions are also present. This test consists of counting the peaks within a small m/z range (0.003 Da) of each reporter ion and calculating the proportion of that count over a count of peaks for each nominal mass. If, for two or more reporter ion nominal masses, the small range around the 10-plex reporter ions makes up more than 10% of the number of peaks in the window, then the run is classified as 10-plex labeled.

A further complication of TMT 10-plex labeling is that, due to the small mass difference between label masses, reporter ions are sometimes measured in MS3 spectra (e.g., in the “SPS-MS3” analysis workflow). In this case, they may not be detectable in the associated MS/MS spectra if the detector used for MS/MS is not capable of measuring the low mass range. Therefore, if MS3 spectra are present in a run, then they are evaluated for the presence of TMT 10-plex reporter ions in the same way described above for MS/MS spectra, even if no TMT reporter ions are observed in the run’s MS/MS spectra.

If iTRAQ or TMT reporter ions are detected in a run, then Param-Medic suggests searching with a variable n-terminal modification appropriate to the specific label detected.

2.3 Detecting SILAC Labeling

SILAC labeling is a popular method for comparing two or three samples using labels incorporated into lysine and/or arginine residues. Accordingly, in a SILAC-labeled mixed sample, the same peptide occurring in two differently-labeled samples will have a different precursor mass in each sample, potentially generating a pair of MS/MS spectra nearby in retention time, with precursor masses having a nominal mass difference of 4 Da, 6 Da, 8 Da or 10 Da, depending on the labels used.

For each SILAC mass separation, Param-Medic counts the number of such pairs separated by no more than 50 scans and calculates a Z-score comparing the observed number of pairs to the counts of pairs found at six different control mass separations (11, 14, 15, 21, 23, and 27 Da). If the Z-score exceeds an empirically-determined threshold, then the run is determined to contain the corresponding SILAC label, and Param-Medic suggests that the run be searched with an appropriate variable modification on both lysine and arginine residues.

2.4 Data sets

To select decision thresholds for each statistic, and to evaluate the performance of the resulting classifiers, we downloaded experiments containing each attribute from the MassIVE proteomics repository (Table 1). We used individual filenames and information from the associated manuscripts to determine which specific runs in each experiment contained the modification and which did not. For each attribute, we assembled a list of runs not containing the attribute from all of the appropriate datasets for the other attributes and the unmodified runs, if any, from the experiment containing the attribute in question. In this way, we assembled a training set with a total of 196 runs and a test set with a total of 174

runs, ensuring that all runs from a given experiment are assigned either to the training or test set. Some experiments contained hundreds of runs; to ensure that no experiment dominated the test set, in those cases we chose a subset of runs that came first in lexicographic order, in an attempt to use data from similar types of samples.

2.5 Selecting decision thresholds

To select a decision threshold for each of the statistics described above, we ran Param-Medic on all of the training set runs. In each case, a higher value of the statistic indicates a stronger signal representing the presence of the attribute. For each statistic, we selected a decision threshold that maximized classification accuracy on the training set runs.

2.6 Evaluation

Finally, we ran Param-Medic on the test set runs. For each classifier, we considered the runs containing the associated attribute to be positive examples and the remaining runs to be negative examples. We created a receiver operator characteristic (ROC) curve describing the performance of the classifier, and we computed sensitivity and specificity at the selected decision boundary.

3 Results

Overall, evaluation of Param-Medic on the test set suggests that the software performs very well across the 10 binary classification tasks that we considered (Table 2 and Figure 1).

TMT labeling can be detected accurately, except that Param-Medic has a difficult time distinguishing between 6-plex and 10-plex TMT, even on the training data. Since the two varieties of TMT require the same search parameter settings (but different quantitative analysis), detecting 6-plex TMT in run with 10-plex labeling is still valuable. Therefore, if 10-plex TMT is detected, then TMT 6-plex is also indicated as present.

In the case of iTRAQ, both settings (8-plex and 4-plex) can be identified very accurately.

For SILAC, the performance for 4 Da, 8 Da and 10 Da labeling is quite good, but classification of 6 Da SILAC labeling is not as successful, with only 57% sensitivity at 99% specificity. Looking more closely at these results, we observe a stark difference between the classifier's performance on the 25 runs from experiment MSV000079583 (100% sensitivity), which had only 6 Da and 10 Da SILAC separations, and the 23 runs from experiment MSV000079097 (0% sensitivity), which had all four SILAC separations. The relative proportion of pairs of ions separated by 6 Da in experiment MSV000079097, although somewhat higher than background, is far weaker than the relative proportion of pairs of ions with the other three SILAC separations, and also far smaller than the relative proportion of pairs of ions separated by 6 Da in experiment MSV000079583. It may be that the other three SILAC separations interfere with the approach, or that experiment MSV000079097 simply doesn't contain as many 6 Da-separated SILAC pairs.

Finally, we find that detection of enrichment for phosphorylated peptides is somewhat difficult, possibly due to the presence of a small proportion of phosphorylated peptides in un-enriched samples. Our classifier achieves 94% sensitivity at 98% specificity.

In the case of phosphorylation, we found ourselves using Param-Medic in precisely the way we anticipate it being used by other researchers. One of our datasets (with MassIVE accession MSV000080777) contained both TMT tags and phospho-enrichment. Initially, our phosphorylation classifier appeared to have very low sensitivity (57%) on the 28 runs from this experiment. However, upon investigation, all 11 of the files the classifier declared to be non-phosphorylated contained “_FT_” in their filenames. The associated manuscript¹⁶ described an analysis that may have been performed on the flow-through, or phospho-depleted portions of the samples, but the manuscript and MassIVE submission did not indicate which files represented phosphorylated samples. Subsequently, through private communication with the authors, we confirmed that the “_FT_” files were indeed flow-through samples. We therefore changed the labels of those runs in our analysis.

4 Discussion

We have described enhancements to Param-Medic that detect the presence of SILAC labeling, TMT and iTRAQ labeling, and enrichment of phosphorylated peptides with high sensitivity and specificity. These tools, combined with the existing Param-Medic estimation of precursor and fragment mass error, provide many of the parameter values needed to successfully search mass spectrometry data files whose metadata are not known *a priori*.

In fact, the classifiers described here were useful in the development of the classifiers themselves. The apparent lack of signal representing phospho-enrichment in 11 of the runs from one of our phosphorylated experiments, all from files with “_FT_” in the filename, led us to investigate and determine that those 11 files represented flow-through fractions, as confirmed by the authors. Param-Medic can be helpful in unraveling the puzzle of which specific runs from a publicly available experiment represent which analyses. In this way, Param-Medic could save a great deal of time for researchers engaged in large-scale meta-analyses of public data.

Param-Medic has been implemented as a standalone Python 2.7 tool which may be downloaded (including source code) at <https://github.com/dhmay/param-medic> or simply added to a Python installation with the “pip” tool. It has also been incorporated into version 3.2 of the Crux Toolkit, available at <http://crux.ms>.

Acknowledgments

We are grateful to Nuno Bandeira for helpful discussions and for pointing us to relevant data sets in MassIVE. Research reported in this publication was supported by the National Defense Science and Engineering Graduate Fellowship (NDSEG) Program, National Science Foundation award 1633939, and National Institutes of Health award R01 GM121818.

References

- [1]. May Damon H., Tamura Kaipo, and Noble William S.. Param-Medic: A Tool for Improving MS/MS Database Search Yield by Optimizing Parameter Settings. *Journal of Proteome Research*, 16(4):1817–1824, 2017. [PubMed: 28263070]
- [2]. Ong Shao-En, Blagoev Blagoy, Kratchmarova Irina, Kristensen Dan Bach, Steen Hanno, Pandey Akhilesh, and Mann Matthias. Stable isotope labeling by amino acids in cell culture, SILAC, as a simple and accurate approach to expression proteomics. *Molecular & cellular proteomics : MCP*, 1(5):376–86, 2002. [PubMed: 12118079]
- [3]. Wiese Sebastian, Reidegeld Kai A., Meyer Helmut E., and Warscheid Bettina. Protein labeling by iTRAQ: A new tool for quantitative mass spectrometry in proteome research. *Proteomics*, 7(3): 340–350, 2007. [PubMed: 17177251]
- [4]. Pottiez Gwenael, Wiederin Jayme, Fox Howard S., and Ciborowski Pawel. Comparison of 4-plex to 8-plex iTRAQ quantitative measurements of proteins in human plasma samples. *Journal of Proteome Research*, 11(7):3774–3781, 2012. [PubMed: 22594965]
- [5]. Thompson Andrew, Schäfer Jürgen, Kuhn Karsten, Kienle Stefan, Schwarz Josef, Schmidt Günter, Neumann Thomas, and Hamon Christian. Tandem mass tags: A novel quantification strategy for comparative analysis of complex protein mixtures by MS/MS. *Analytical Chemistry*, 75(8):1895–1904, 2003. [PubMed: 12713048]
- [6]. Dayon Loïc, Hainard Alexandre, Licker Virginie, Turck Natacha, Kuhn Karsten, Hochstrasser Denis F., Burkhard Pierre R., and Sanchez Jean Charles. Relative quantification of proteins in human cerebrospinal fluids by MS/MS using 6-plex isobaric tags. *Analytical Chemistry*, 80(8): 2921–2931, 2008. [PubMed: 18312001]
- [7]. Mingxun Wang, Wang Jian, Carver Jeremy, Pullman Benjamin S., Cha Seong Won, and Bandeira Nuno. Assembling the Community-Scale Discoverable Human Proteome. *Cell Systems*, 7(4): 412–421, 2018. [PubMed: 30172843]
- [8]. Filla Jan and Honys David. Enrichment techniques employed in phosphoproteomics. *Amino Acids*, 43(3):1025–1047, 2012. [PubMed: 22002794]
- [9]. Lehmann Wolf D., Krüger Ralf, Salek Mogjiborahman Hung Chien Wen, Wolschin Florian and Weckwerth Wolfram. Neutral loss-based phosphopeptide recognition: A collection of caveats. *Journal of Proteome Research*, 6(7):2866–73, 2007. [PubMed: 17569551]
- [10]. Wolski Witold E, Farrow Malcolm, Emde Anne-Katrin, Lehrach Hans, Lalowski Maciej, and Reinert Knut. Analytical model of peptide mass cluster centres with applications. *Proteome science*, 4:18, 2006. [PubMed: 16995952]
- [11]. Werner Thilo, Sweetman Gavain, Savitski Maria Fälth, Mathieson Toby, Bantscheff Marcus, and Savitski Mikhail M.. Ion coalescence of neutron encoded TMT 10-plex reporter ions. *Analytical Chemistry*, 86(7):3594–3601, 2014. [PubMed: 24579773]
- [12]. Udeshi Namrata D., Svinkina Tanya, Mertins Philipp, Kuhn Eric, Mani DR, Qiao Jana W., and Carr Steven A.. Refined Preparation and Use of Anti-diglycine Remnant (K-ε-GG) Antibody Enables Routine Quantification of 10,000s of Ubiquitination Sites in Single Proteomics Experiments. *Molecular & Cellular Proteomics*, 12(3):825–831, 2013. [PubMed: 23266961]
- [13]. Jiang Yanxialei, Lee Jeeyoung, Lee Jung Hoon, Lee Joon Won, Kim Ji Hyeon, Choi Won Hoon, Yoo Young Dong, Cha-Molstad Hyunjo, Kim Bo Yeon, Kwon Yong Tae, Noh Sue Ah, Kim Kwang Pyo, and Lee Min Jae. The arginylation branch of the N-end rule pathway positively regulates cellular autophagic flux and clearance of proteotoxic proteins. *Autophagy*, 12(11): 2197–2212, 2016. [PubMed: 27560450]
- [14]. Shiromizu Takashi, Adachi Jun, Watanabe Shio, Murakami Tatsuo, Kuga Takahisa, Muraoka Satoshi, and Tomonaga Takeshi. Identification of missing proteins in the neXtProt database and unregistered phosphopeptides in the phosphositeplus database as part of the chromosome-centric human proteome project. *Journal of Proteome Research*, 12(6):2414–21, 2013. [PubMed: 23312004]
- [15]. Clark David J., Fondrie William E., Liao Zhongping, Hanson Phyllis I., Fulton Amy, Mao Li, and Yang Austin J.. Redefining the Breast Cancer Exosome Proteome by Tandem Mass Tag Quantitative Proteomics and Multivariate Cluster Analysis. *Analytical Chemistry*, 87(20):10462–10469, 2015. [PubMed: 26378940]

- [16]. Hoffman Nolan J., Parker Benjamin L., Chaudhuri Rima, Fisher-Wellman Kelsey H., Kleinert Maximilian, Humphrey Sean J., Yang Pengyi, Holliday Mira, Trefely Sophie, Fazakerley Daniel J., Stöckli Jacqueline, Burchfield James G., Jensen Thomas E., Jothi Raja, Kiens Bente, Wojtaszewski Jørgen F.P., Richter Erik A., and James David E.. Global Phosphoproteomic Analysis of Human Skeletal Muscle Reveals a Network of Exercise-Regulated Kinases and AMPK Substrates. *Cell metabolism*, 22(5):922–35, 2015. [PubMed: 26437602]
- [17]. Xu Benhong, Gao Yanpan, Zhan Shaohua, Xiong Feng, Qiu Wenying, Qian Xiaojing, Wang Tao, Wang Naili, Zhang Di, Yang Qian, Wang Renzhi, Bao Xinjie, Dou Wanchen, Tian Rui, Meng Shu, Gai Wei Ping, Huang Yue, Yan Xiao Xin, Ge Wei, and Ma Chao. Quantitative protein profiling of hippocampus during human aging. *Neurobiology of Aging*, 39:46–56, 2016. [PubMed: 26923401]
- [18]. Sahasrabuddhe Nandini A., Barbhuiya Mustafa A., Bhunia Shushruta, Subbannayya Tejaswini, Gowda Harsha, Advani Jayshree, Shrivastav Braj R., Navani Sanjay, Leal Pamela, Roa Juan Carlos, Chaerkady Raghothama, Gupta Sanjeev, Chatterjee Aditi, Pandey Akhilesh, and Tiwari Pramod K.. Identification of prosaposin and transgelin as potential biomarkers for gallbladder cancer using quantitative proteomics. *Biochemical and Biophysical Research Communications*, 446(4):863–9, 2014. [PubMed: 24657443]
- [19]. Guo Jiao, Ren Yan, Hou Guixue, Wen Bo, Xian Feng, Chen Zhen, Cui Ping, Xie Yingying, Zi Jin, Lin Liang, Wu Song, Li Zesong, Wu Lin, Lou Xiaomin, and Liu Siqu. A Comprehensive Investigation toward the Indicative Proteins of Bladder Cancer in Urine: From Surveying Cell Secretomes to Verifying Urine Proteins. *Journal of Proteome Research*, 15(7):2164–77, 2016. [PubMed: 27265680]
- [20]. Preil Simone A.R., Kristensen Lars P., Beck Hans C., Jensen Pia S., Nielsen Patricia S., Steiniche Torben, Bjørling-Poulsen Marina, Larsen Martin R., Hansen Maria L., and Rasmussen Lars M.. Quantitative Proteome Analysis Reveals Increased Content of Basement Membrane Proteins in Arteries from Patients with Type 2 Diabetes Mellitus and Lower Levels among Metformin Users. *Circulation: Cardiovascular Genetics*, 8(5):727–35, 2015. [PubMed: 26371159]

Figure 1: Attribute detection performance.

ROC curves describing the performance of each of the classifiers. The dot on each line corresponds to the selected decision threshold. All ROC curves are for test data sets, except that TMT10 is for training data.

Table 1:
Experiments used in the construction and assessment modification-detection classifiers.

All accession numbers are MassIVE accessions except PXD002053, which is a PRIDE accession. The score threshold determined for the 6 Da SILAC classifier, on 6 Da SILAC data, was used for all four SILAC classifiers; hence, no training data are listed for the other three. Experiment MSV000079583 had SILAC-labeled runs, iTRAQ-labeled runs and phosphorylated runs but was only used for training the iTRAQ classifier. There were no test data for 10-plex TMT, and results below are from evaluation of training data.

Modification	Training data set(s)	Test data set(s)
SILAC (4 Da)	None	MSV000079097 ¹²
SILAC (6 Da)	MSV000080791 ¹³	MSV000079097, MSV000079583 ¹⁴
SILAC (8 Da)	None	MSV000079097
SILAC (10 Da)	None	MSV000079097, MSV000079583
TMT (6-plex)	PXD002053 ¹⁵	MSV000080777, ¹⁶ MSV000080710 ¹⁷
TMT (10-plex)	MSV000079068 ¹⁷	None
iTRAQ (4-plex)	MSV000079583	MSV000082884 ¹⁸
iTRAQ (8-plex)	MSV000080696 ¹⁹	MSV000080848 ²⁰
Phosphorylation	MSV000080773	MSV000079583, MSV000080777

Table 2:
Performance of the classifier on each attribute.

Classifiers were evaluated on 228 total runs. Some runs contained multiple attributes. “AUROC” is the area under the ROC curve.

Modification	Runs with Modification	Sensitivity	Specificity	AUROC
TMT (6-plex)	53	0.98	0.98	0.98
TMT (10-plex)	5	0.80	0.86	0.75
iTRAQ (4-plex)	58	0.98	0.99	0.99
iTRAQ (8-plex)	30	0.97	1.00	1.00
SILAC (4 Da)	12	0.92	1.00	1.00
SILAC (6 Da)	37	0.57	0.99	0.96
SILAC (8 Da)	12	1.00	0.99	1.00
SILAC (10 Da)	37	0.89	0.99	0.98
Phosphorylation	17	0.94	0.98	0.95