Published in final edited form as:

Nat Methods. 2015 July ; 12(7): 605–608. doi:10.1038/nmeth.3450.

Mass spectrometrists should only search for peptides they care about

William Stafford Noble*

Department of Genome Sciences, Department of Computer Science and Engineering, University of Washington

Abstract

Analysis pipelines that assign peptides to shotgun proteomics spectra often discard identified spectra deemed irrelevant to the scientific hypothesis being tested. To improve statistical power, irrelevant peptides should be removed from the database prior to searching, rather than assigning them to spectra and then discarding the matches.

Over the past decade, perhaps the most significant trend in the analysis of shotgun proteomics data has been the increasing statistical rigor of most analysis pipelines. No longer are reviewers or editors satisfied with spectrum identifications defined with respect to arbitrary score thresholds or by using “rules of thumb” based on multiple thresholds. Most journals now require that statistical confidence estimates be reported. A variety of methods have been devised for assigning confidence estimates to individual matches, using parametric¹ or exact procedures.^{2,3} Perhaps most importantly, target-decoy analysis^{4,5} provides a straightforward method for estimating the *false discovery rate* (FDR), defined as the percentage of incorrect identifications, associated with nearly any data set and scoring procedure.

However, despite these advances in statistical confidence estimation, one problematic protocol remains in common use. The protocol is quite general and involves testing more hypotheses than we are actually interested in. To illustrate the idea, consider the analysis of mass spectra derived from the erythrocytic cycle of the malaria parasite *Plasmodium falciparum*. Because erythrocytic *Plasmodium* parasites inhabit human red blood cells, any *Plasmodium* proteomics experiment will inevitably generate spectra from a mixture of human and *Plasmodium* peptides. It is therefore common to search the observed spectra against a combined database of human and *Plasmodium* peptides and then to discard the spectra that match to human peptides.

From a statistical perspective, this protocol is suboptimal, in the sense that it needlessly sacrifices statistical power. In particular, at a fixed FDR threshold, we can obtain a larger set of identifications by searching the spectra against only the *Plasmodium* peptides. To understand why this is the case, we need only consider a single spectrum *s* searched against

*3720 15th Ave NE, Seattle, WA 98195-5065; Tel: 206 543-8930; Fax: 206 685-7301; william-noble@uw.edu.

Statement of Competing Financial Interests

The author declares no competing financial interests.

either the *Plasmodium* database or the *Plasmodium*+human database. If the spectrum has an associated precursor mass of, say, 2000.0 Da, and if we search the databases using a tolerance of 10 parts-per-million, then the *Plasmodium* database yields 17 candidate peptides and the *Plasmodium*+human database yields 29 candidate peptides. Let us assume that we are using a method such as MS-GF+² that assigns a *P*-value to each possible match. After selecting the best-scoring candidate peptide (i.e., the one with the smallest *P*-value), a Bonferroni adjustment requires that, to achieve a statistical significance of $P < 0.01$, we must observe a *P*-value of $0.01/17 = 5.9 \times 10^{-4}$ from the *Plasmodium* database but a *P*-value of $0.01/29 = 3.4 \times 10^{-4}$ from the *Plasmodium*+human database. Thus, whenever our best-scoring match against the *Plasmodium* database receives a *P*-value in between these two values, we will fail to identify that match if we choose also to search the human database. This is ironic, because the only potential gain we achieve by searching the human database is an identification that we have no interest in and will ultimately discard.

Many mass spectrometrists, when presented with this idea, express concern that the spectrum identifications produced by this simpler protocol will be contaminated with human spectra. The reasoning goes like this: even if spectrum *s* obtained a score that exceeds a given threshold when we searched it against the *Plasmodium* database, that spectrum might have received an *even larger score* if we had searched it against the combined *Plasmodium* and human databases. While this statement is certainly true, it misses the point of our statistical confidence estimation procedure, which is to accurately estimate the FDR associated with a given collection of identified spectra. Well calibrated statistical confidence estimates allow us to skip the testing of these extraneous hypotheses.

One potential source of confusion in the assignment of confidence estimates to identified spectra is that the most commonly used assignment method—assigning FDRs using target-decoy competition—does not explicitly make use of *P*-values. It is not obvious, therefore, that notions such as Bonferroni adjustment and multiple hypothesis testing are applicable in this domain. The following two case studies, which were carried out using target-decoy FDR estimation, illustrate that these concepts do indeed apply. Any valid confidence estimation procedure, even one that does not make explicit use of *P*-values, must take into account the large number of peptides in the database and the large number of spectra in a given experiment.

Two case studies

To illustrate the loss of statistical power associated with searching a large database, we analyzed two recently published collections of mass spectra. The first contains 12,478 high-resolution *Plasmodium* spectra.⁶ We searched the spectra against a *Plasmodium* database, with and without the human database appended. We used the MS-GF+ search engine², and we estimated FDRs using target-decoy competition.⁵ Searching the combined database always yielded fewer *Plasmodium* identifications across FDR thresholds up to 10% (Fig. 1a). In particular, at an FDR threshold of 1%, the combined search assigned *Plasmodium* peptides to 2339 spectra, whereas the *Plasmodium*-only search assigned 2530 such spectra, an increase of 8.1%. Furthermore, only two of the 2530 spectra identified in the *Plasmodium*-only search received a better score in the corresponding *Plasmodium*+human

search, indicating that the *Plasmodium*-only results are not highly contaminated with spurious matches due to the presence of human-derived spectra.

A potential complication arises when peptides are shared between *Plasmodium* and human. In a tryptic digestion with no missed cleavages and no variable modifications, the *Plasmodium* and human databases contain 221,567 and 432,840 peptides, respectively, with an overlap of 916 peptides. One risk associated with searching only the *Plasmodium* database is that human peptides from this overlap set might be misidentified as *Plasmodium* peptides. The solution to this problem is to check whether any of the identified *Plasmodium* peptides occur in human, and to handle these identifications accordingly. Indeed, one might wish to eliminate these overlapping peptides from the *Plasmodium* database *a priori*, thereby further reducing the multiple testing burden.

We also analyzed a collection of 1503 fragmentation spectra generated during an investigation of coding variants and isoforms of pulmonary surfactant protein A (SP-A).⁷ The protocol involved purifying SPA proteins from bronchial lavage fluid and subjecting the purified proteins to tandem mass spectrometry analysis. Because the purification procedure is necessarily imperfect, the resulting spectra were searched against the entire human proteome, augmented with a collection of six known SP-A variants. We used Tide³ to search the spectra against this combined database, as well as against a small database containing only variants of SP-A. Similar to the *Plasmodium* analysis, this experiment showed that searching only against the proteins of interest yields much better statistical power (Fig. 1b). At a 1% FDR threshold, the combined search assigned 345 spectra to peptides from SP-A, whereas the search against the SP-A database identified 448 spectra, an increase of 29.9%.

This practice of sacrificing statistical power by considering irrelevant hypotheses is common. For example, any proteomics experiment that targets a particular pathway or set of pathways would benefit from using a protein database consisting only of the proteins of interest. Similarly, any study that aims only to identify phosphorylation sites could gain statistical power by not searching for unphosphorylated peptides. When searching for cross-linked peptides, uninteresting species such as non-cross-linked, self-loop and dead-end peptides should be left out of the database. A good rule of thumb is that if you are going to do a post-filtering step to eliminate some of the matches from your analysis, and if those matches are not relevant to the scientific hypothesis you are testing, then you should consider eliminating those peptides beforehand to avoid having to correct for these extra statistical tests.

A practical follow-up question to both the *Plasmodium* and SP-A examples is how the magnitude of the loss in statistical power varies as a function of the number of irrelevant proteins included in the search. To address this question empirically, we carried out searches using a series of randomly downsampled human databases. For the *Plasmodium* data set, we observe a rapid loss of power even when only 10,000 human proteins are added to the database: at a 1% FDR threshold, the number of spectra assigned to *Plasmodium* peptides drops by 3.1%, from 2729 to 2643 (Fig. 2a). Power continues to drop relatively smoothly as the database grows larger. In the SP-A case, the initial drop in power is more dramatic. Even adding 1000 human peptides reduces power by 13.4% (434 to 376 peptides) at 1% FDR

(Fig. 2b). The reason for this behavior is that, in this case, less than half (709 of 1503) of the observed spectra have associated precursor m/z values that lie within ± 3 m/z of the mass of at least one peptide in the SP-A database. The remaining 794 spectra are therefore not matched at all in the SP-A search. When we add human proteins to the database, most of those 794 spectra match at least one candidate peptide. Thus, in this case, adding human peptides to the database increases the multiple testing burden in two ways: increasing the number of candidate peptides per spectrum, and increasing the total number of spectra under consideration.

Beyond maximizing statistical power

In practice, simply maximizing statistical power may not be the only or even the primary concern. For example, focusing an experiment on a small, targeted collection of peptides necessarily eliminates the possibility of serendipitous, unexpected discoveries. Indeed, one of the great powers of tandem mass spectrometry is its depth of coverage, with the potential to query a large proportion of the proteome in a single experiment.

A second important consideration, before filtering out irrelevant peptides, is that the analysis of those peptides can provide a useful sanity check. For example, common contaminants such as human keratin should probably always be included in any search. Overall, the total number of such contaminants is generally quite small, and their presence in the database will likely have a negligible impact on statistical power.

Irrelevant peptides can also be helpful to identify problems with the statistical analysis. For example, Bromenshenk et al. reported an analysis of honeybee-derived spectra that yielded sequences from two previously unreported RNA viruses—*Varroa destructor*-1 and Kukugo virus—in North American honeybees as well as an invertebrate iridovirus (IIV) associated with collapsed colonies.⁸ The analysis involved searching the spectra against a database containing only bacterial and viral sequences, without including the honeybee proteins that are expected to match the majority of the spectra. A total of 3000 peptides were confidently identified from over 900 different microbial species. Three critiques of this study were subsequently published by two different research groups, pointing out various problems with the reported results and with the analysis procedures.^{9;10;11} Knudsen and Chalkley hypothesized that the tool PeptideProphet¹² was provided with an input data set containing no correct identifications, and as a result the tool gave invalid results. Both critical reassessments claimed that, had Bromenshenk et al. included the honeybee proteins in their initial database search, then mistakes could have been avoided.

To investigate this claim, we re-analyzed the full set of 262,572 spectra from the original Bromenshenk et al. study. Similar to the *Plasmodium* and SP-A case studies, the spectra were searched against a viral database, a honeybee database, and a concatenation of the two. At an FDR threshold of 1%, the Tide search engine identified 36,371 spectra in the honeybee search. The search of the viral database, on the other hand, identified only 152 spectra. These viral identifications correspond to only nine distinct peptides (Table 1). In agreement with previous re-analyses of this data set,^{9;11} none of the nine peptides come from *Varroa destructor*-1, Kukugo or invertebrate iridoviruses. Two of the nine

peptides occur in both the honeybee and viral databases, and one occurs in both databases modulo a substitution of isobaric amino acids ($K \rightarrow Q$). Together, these three shared peptides account for 115 of the 152 spectra. Among the remaining six viral peptides, four appear to be matches to actual viral sequences from the Kashmir bee virus, the Acute bee paralysis virus, and the Sacbrood virus. Only two peptides identified as viral in the initial search, and corresponding to 10 distinct spectra, receive a better match to a homologous honeybee peptide in the combined search.

This analysis suggests two conclusions. First, in rare cases, leaving irrelevant peptides out of the database may lead to incorrectly identified spectra due to homology between the proteins in the database and the left-out proteins: in this case, 10 identifications out of 152. Thus, the risk of such false positive identifications must be weighed against the potential gain in statistical power (loss of false negatives) when deciding upon an analysis strategy. Second, in the particular case of the honeybee data set, the choice of protein database does not explain the incorrect results obtained by Bromenshenk et al. Instead, the problem apparently lay in how the analysis was performed.

Challenges and future directions

One challenge associated with our proposed rule of thumb is that if we use a decoy-based estimation procedure, and if our database consists of only a handful of proteins, then the resulting confidence estimates will likely be inaccurate. Indeed, one published standard explicitly warns about the difficulty of achieving accurate confidence estimates when the database contains fewer than 1000 proteins.¹³ Three points are worth considering with respect to this problem. First, any empirical confidence estimation protocol such as target-decoy competition yields intrinsic variance due to the stochastic generation of decoys. This variance is not removed by using a deterministic decoy generation scheme, such as reversing peptides. Furthermore, as the confidence estimates get small, the relative magnitude of the variance increases relative to the magnitude of the estimate (Fig. 2). Second, if we use a shuffling procedure to generate decoys, then it is possible to reduce the variance in our confidence estimates significantly by repeating the target-decoy competition many times.¹⁴ Such an approach is particularly applicable when the protein database is small, because the cost of repeated searches is relatively low. Third, the use of analytic methods to compute exact P -values with respect to a particular null model avoids the variance associated with empirical confidence estimation schemes.^{2,3}

Note that an alternative protocol to the one proposed here involves first searching the spectra against a database containing only the irrelevant peptides and eliminating from the data set all spectra that match this “garbage” database with high confidence. The remaining spectra could then be searched against the database of interesting peptides. For stringent FDR thresholds, this approach is likely to yield slightly better power than the simple strategy proposed here, at the expense of being somewhat more complicated to implement.

Thus far, our discussion has focused on assigning confidence estimates to peptide-spectrum matches. Confidence estimation at the level of peptides—taking into account multiple spectra matching to the same peptide—or at the level of proteins—taking into account the

many-to-many mapping between peptides and proteins—is considerably more challenging. Nonetheless, existing methods for assigning peptide- and protein-level confidence estimates typically do so by aggregating spectrum-level evidence^{15;16;17}; hence, any gains in statistical power at the spectrum level achieved by leaving out irrelevant peptides should in principle yield concomitant gains in power at the peptide and protein levels.

The relatively simple idea proposed here, of limiting the hypothesis space to only hypotheses of interest, is an expression of a more general statistical goal, which aims to explore the hypothesis space in a fashion that maximizes statistical power. A clear avenue for future work at the interface of statistics and proteomics lies in the application of existing protocols for stratified multiple testing control.¹⁸

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was funded by National Institutes of Health awards R01 GM096306 and P41 GM103533.

References

1. Klammer AA, Park CY, Noble WS. Statistical calibration of the sequest XCorr function. *Journal of Proteome Research*. 2009; 8:2106–2113. [PubMed: 19275164]
2. Kim S, Gupta N, Pevzner PA. Spectral probabilities and generating functions of tandem mass spectra: a strike against decoy databases. *Journal of Proteome Research*. 2008; 7:3354–3363. [PubMed: 18597511]
3. Howbert JJ, Noble WS. Computing exact p-values for a cross-correlation shotgun proteomics score function. *Molecular and Cellular Proteomics*. 2014; 13:2467–2479. [PubMed: 24895379]
4. Moore RE, Young MK, Lee TD. Qscore: An algorithm for evaluating sequest database search results. *Journal of the American Society for Mass Spectrometry*. 2002; 13:378–386. [PubMed: 11951976]
5. Elias JE, Gygi SP. Target-decoy search strategy for increased confidence in large-scale protein identifications by mass spectrometry. *Nature Methods*. 2007; 4:207–214. [PubMed: 17327847]
6. Pease BN, et al. Global analysis of protein expression and phosphorylation of three stages of *Plasmodium falciparum* intraerythrocytic development. *Journal of Proteome Research*. 2013; 12:4028–4045. [PubMed: 23914800]
7. Foster MW, et al. Identification and quantitation of coding variants and isoforms of pulmonary surfactant protein A. *Journal of Proteome Research*. 2014; 13:3722–3732. [PubMed: 25025725]
8. Bromenshenk JJ, et al. Iridovirus and Microsporidian linked to honey bee colony decline. *PLOS One*. 2010; 5:e13181. [PubMed: 20949138]
9. Knudsen GM, Chalkley RJ. The effect of using an inappropriate protein database for proteomic data analysis. *PLOS One*. 2011; 6:e20873. [PubMed: 21695130]
10. Foster LJ. Interpretation of data underlying the link between colony collapse disorder (CCD) and an invertebrate iridescent virus. *Molecular and Cellular Proteomics*. 2011; 10:M110.006387.
11. Foster LJ. Letter to editor. *Molecular and Cellular Proteomics*. 2012; 11:A110.006387–1. [PubMed: 22238354]
12. Keller A, Nesvizhskii AI, Kolker E, Aebersold R. Empirical statistical model to estimate the accuracy of peptide identification made by MS/MS and database search. *Analytical Chemistry*. 2002; 74:5383–5392. [PubMed: 12403597]

13. [Accessed 6 April 2015] Revised publication guidelines for documenting the identification and quantification of peptides, proteins, and posttranslational modifications by mass spectrometry. http://www.mcponline.org/site/misc/peptide_and_protein_identification_guidelines.pdf
14. Keich U, Noble WS. On the importance of well calibrated scores for identifying shotgun proteomics spectra. *Journal of Proteome Research*. 2015; 14:1147–1160. [PubMed: 25482958]
15. Nesvizhskii AI, Keller A, Kolker E, Aebersold R. A statistical model for identifying proteins by tandem mass spectrometry. *Analytical Chemistry*. 2003; 75:4646–4658. [PubMed: 14632076]
16. Serang O, MacCoss MJ, Noble WS. Efficient marginalization to compute protein posterior probabilities from shotgun mass spectrometry data. *Journal of Proteome Research*. 2010; 9:5346–5357. [PubMed: 20712337]
17. Cox J, Mann M. MaxQuant enables high peptide identification rates, individualized p.p.b.-range mass accuracies and proteome-wide protein quantification. *Nature Biotechnology*. 2008; 26:1367–1372.
18. Dudoit, S.; van der Laan, MJ. *Multiple Testing Procedures with Applications to Genomics*. Springer; New York, NY: 2008.

Figure 1. Boosting statistical power by eliminating irrelevant hypotheses

(a) The number of accepted peptide-spectrum matches (PSMs) as a function of FDR threshold for the *Plasmodium* data set. Searching a combined database of *Plasmodium* and human sequences yields more identified spectra overall (blue) but fewer *Plasmodium* identifications (violet) than searching just the *Plasmodium* database (green). (b) Similar to (a), but analyzing isoforms of SP-A alone or in conjunction with the entire human proteome.

Figure 2. Loss of statistical power as a function of database size

The *Plasmodium* and SP-A data sets were searched, using Tide, against a series of databases containing various numbers of human proteins. The figure plots the number of (a) *Plasmodium* and (b) SP-A proteins identified at various FDR thresholds as a function of the number of human proteins included in the database. Results are averaged over ten different decoy databases, and errors correspond to standard deviation.

Table 1

Viral peptides identified in the honeybee data set.

Spectra	Viral	Honeybee	Combined	Species
52	IWHHTFYNELR	same	both	honeybee
46	HKGVMVGMGQK	HKGVMVGMGQK	both	honeybee
17	LAVNMVPFPR	same	both	honeybee
7	IIAQVVSSITASLR	LIGQIVSSITASLR	honeybee	honeybee
3	SYELPDGQVIKIGSER	SYELPDGQVITIGNER	honeybee	honeybee
16	IGPISEVASGVK	various (6 decoys)	viral	Kashmir bee virus
6	DYMSYLSYLYR	various (5 decoys)	viral	Acute bee paralysis virus
4	IDTPMAQDTSSAR	various (4 decoys)	viral	Acute bee paralysis/Kashmir bee virus
1	VNNLHEYTK	NVNITFPQ GK	viral	Sacbrood virus