Analysis of tandem mass spectrometry data with CONGA: Combining Open and Narrow searches with Group-wise Analysis

Jack Freestone[1], William Stafford Noble[2,3], and Uri Keich[1]

[1]School of Mathematics and Statistics F07, University of Sydney
[2]Department of Genome Sciences, University of Washington
[3]Paul G. Allen School of Computer Science and Engineering, University of Washington

May 2, 2023

Abstract

Searching tandem mass spectrometry proteomics data against a database is a well-established method for assigning peptide sequences to observed spectra but cannot identify peptides harboring unexpected post-translational modifications (PTMs). Open modification searching aims to address this problem by allowing a spectrum to match a peptide even if the spectrum's precursor mass differs from the peptide mass. However, expanding the search space in this way can lead to a loss in statistical power to detect peptides. We therefore developed a method, called CONGA, that takes into account results from both types of searches—a traditional "narrow window" search and an open modification search—while carrying out rigorous false discovery rate (FDR) control. The result is an algorithm that provides the best of both worlds: the ability to detect unexpected PTMs without a concomitant loss of power to detect unmodified peptides.

1 Introduction

Tandem mass spectrometry provides an efficient way to study proteins in a high-throughput fashion but is typically limited in its ability to detect post-translationally modified (PTM) peptides due to the exponentially large search space associated with searching for PTMs. One proposed solution to this problem is to employ "open modification" searching, in which each observed spectrum is compared against peptides whose masses differ—often by hundreds of Daltons—from the observed precursor mass associated with the spectrum [2]. Open modification searching has become increasingly popular, with one report suggesting that this type of approach can achieve higher statistical power to detect peptides than a traditional "narrow window" database search [14].

We recently showed that open searches on their own often produce fewer discoveries than narrow searches applied to the same data. Only when coupled with a machine learning post-processor such as Percolator [11] or PeptideProphet [12] do open searches become typically better than narrow window searches. At the same time we showed that in practice both Percolator and PeptideProphet can apparently fail to control the FDR. Hence, that apparent improvement in power attributed to open modification searching must be taken with a grain of salt [7].

Motivated by these observations, we propose an alternative post-processing approach that allows us to consistently deliver more discoveries compared to the traditional narrow search. The key idea is to search the observed spectra twice, once using a wide (open) window and once using a narrow window, and then combine the search results in a novel way that allows us to extract

information from both searches while rigorously and empirically controlling the FDR. The method, Combining Open and Narrow searches with Group-wise Analysis (CONGA), draws on the same concept of target-decoy competition (TDC) that is commonly used to control the FDR: each real ("target") peptide is paired with a randomly shuffled or reversed decoy peptide, and each spectrum is searched against the concatenated target-decoy database. Assuming that a false match is equally likely to involve a target or a decoy peptide, the optimal matches to the decoy peptides allow us to estimate the number of false discoveries and hence to control the FDR. In CONGA's case this step essentially involves a competition scheme that produces a filtered list of peptides, which are then grouped according to the difference between the observed precursor mass and the peptide mass. This grouping is then followed by our FDR analysis, which achieves something akin to separately controlling the FDR in each group while still controlling the overall FDR.

We provide empirical evidence for CONGA's strong performance. First, we use an entrapment experiment to verify CONGA's FDR control. We then show that CONGA achieves better power to detect peptides than either a narrow or open modification search, and that its power is comparable with that of Percolator. We also demonstrate that CONGA's approach allows us to detect and better utilize chimeric spectra, because CONGA's group-wise analysis considers multiple matching peptides for each spectrum while still controlling the overall FDR.

2 Results

2.1 The CONGA algorithm

Motivated in part by the observed failure of existing post-processors to control the FDR [7], we developed CONGA. The algorithm takes as input two sets of search results—the top PSM for each spectrum from a narrow search against a concatenated target-decoy database, and the top n_t PSMs for each spectrum from an open search against the same concatenated database. CONGA also requires the target-decoy peptide pairs that make up the concatenated database. The algorithm then proceeds through four main steps, which are outlined in Figure 1 and below, with further details in Section 3.1.

CONGA's first step weeds out potentially problematic PSMs. These PSMs arise because CONGA considers more than one PSM per spectrum, allowing the algorithm to properly handle chimeric spectra generated by more than one peptide. The challenge is that some PSMs might "hitchhike," if the peptide from the hitchhiker closely resembles the true, generating peptide. To filter out these hitchhikers, CONGA combines the list of narrow- and open-search PSMs for each spectrum and removes sub-optimal PSMs with peptides whose theoretical spectra share a non-negligible number of peaks with higher scoring peptides matching the same spectrum (Figure 1A).

CONGA's second step introduces "dynamic-level competition," which is designed to help TDC overcome some inherent shortcomings of decoy construction. Specifically, TDC relies on the assumption that a false discovery—in our case a peptide that won its competition but is not in the sample—is equally likely to be a target or a decoy win. However, when variable modifications are specified the database will often contain several clusters of highly similar peptides corresponding to variable modifications of the same unmodified ("stem") form of a single peptide. These clusters can create a problem for the TDC approach. Indeed, as typically constructed, by shuffling or reversing the target peptides, the decoys cannot account for mismatches within a cluster. For example, consider a spectrum that was generated by P[10]EPTIDE but because of random noise is optimally matched to PEP[10]TIDE. In this case, the decoys cannot offer real competition to such a mismatch because, typically, the "closest" decoy to P[10]EPTIDE would still be significantly different from its cluster-neighbor PEP[10]TIDE. This problem can be further exacerbated in an

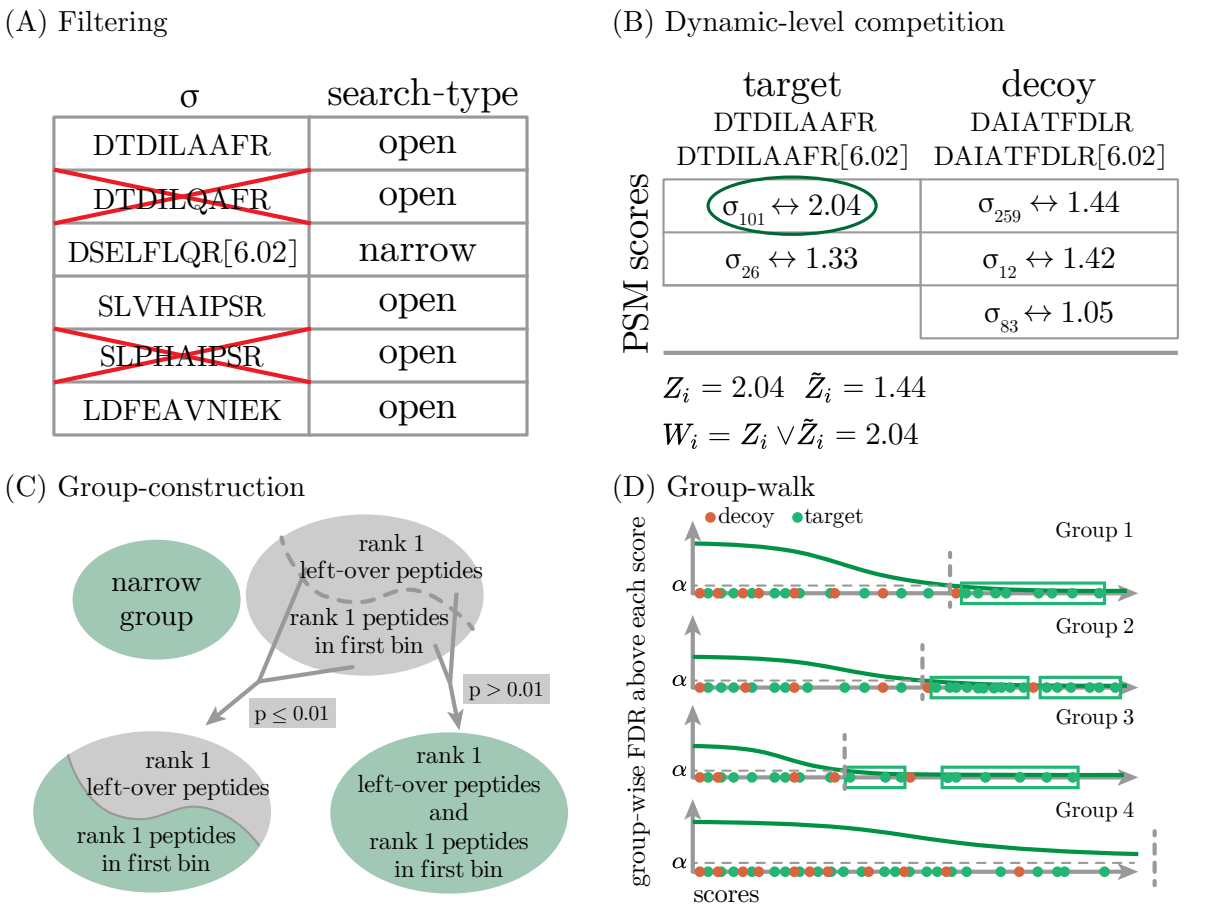

Figure 1: **Schematics of CONGA.** (A) Considering matches to the experimental spectrum σ, some sub-optimal PSMs need to be weeded out because they involve peptides whose theoretical spectra share a non trivial number of peaks with higher scoring matched peptides. (B) The best scoring peptide among all PSMs that have the same unmodified stem form, along with their decoys, is then selected. (C) Partitioning into groups: the (representative) peptides are first divided into a static "narrow group" consisting of peptides that were maximally-matched by a PSM appearing in the narrow search, and a second "open group" that initially contains all other peptides. We next adaptively split and merge sub-groups of the open group starting with comparing the winning scores of a "candidate set", consisting of the rank 1 peptides with the most commonly occurring mass-difference (labelled as the *first bin*), with the winning scores of the remaining rank 1 peptides of the open group. If the p-value of the Kolmogorov-Smirnov test is smaller than 0.01, a new group is created from the candidate set, else a new group is created from all of the rank 1 peptides of the open group (see Section 3.1.3 for further details). (D) Once the groups are defined, we employ the Group-walk algorithm which iteratively updates a vector of group-wise thresholds until the estimated FDR from the combined peptides above (to the right of) each threshold is less than α. The target peptides above each group-threshold (indicated by the rectangles) are then reported.

open search, where in our example P[10]EP[10]TIDE would also belong to the cluster.

To address this problem, CONGA dynamically adjusts the level at which the target-decoy competition takes place. When no variable modifications are specified, each target or decoy peptide is assigned a score, which is the same as the maximally scoring PSM associated with the peptide, and of each target-decoy pair of peptides only the higher scoring one is kept, together with its target/decoy label. However, when variable modifications are specified, CONGA first combines peptides into groups prior to the competition. Each group consists of all the modified forms of the same unmodified stem peptide. CONGA then replaces its peptide-level analysis by stem-level analysis, where each cluster of peptides related by variable modifications is represented by its maximally scoring representative peptide. The FDR analysis is done on this list of stem peptides, each with its own score; however, as explained below, some of the modified versions which share the same stem form as a discovered representative can still be reported separately.

CONGA's third step is the essential one of dividing the list of stem peptides into groups. This fairly involved process, detailed in Section 3.1.3, defines the groups based on (a) whether the peptide was maximally-matched by a PSM that appeared in the narrow search ("narrow group"), (b) the frequency of the difference between the precursor mass and the stem peptide mass, equivalently of the mass of the unaccounted-for PTM, and (c) the open-search rank R_p of the maximal PSM that gave the peptide its score relative to the other matches to the same spectrum (CONGA only considers peptides with $R_p \leq 2$). While the narrow group remains static, the remaining peptides are adaptively split and merged into groups based on the similarity of the winning peptides score distributions as determined by a Kolmogorov-Smirnov test.

Finally, CONGA employs the Group-walk algorithm, which is analogous to conducting a separate TDC analysis on each group while maintaining overall FDR control [8].

CONGA's dynamic level competition is designed with rigorous FDR control in mind. However, regardless of whether searching with variable or static modifications, CONGA's approach necessarily leads to some loss of information. For example, particularly when searching without any specified variable modifications, the user might still be interested to learn which unexpected mass modifications of a particular peptide have support in the data. This information is not directly available from CONGA's FDR-controlled list of discovered peptides. Therefore, CONGA augments this list by reporting for each peptide in the list any PSMs involving the same peptide that (a) is sufficiently high scoring (is a top-ranked narrow, or top-2 open PSM), and (b) is the highest scoring of all the PSMs with the same peptide that fall into the same mass modification bin. The mass bins are defined by a greedy clustering algorithm (Supplementary Algorithm 5) that takes into account the specified isolation window. This clustering is necessary to avoid overwhelming the user with PSMs that share the same peptide and have very similar precursor masses. The augmented list is further supported with a label for each PSM which indicates whether the PSM scored above its CONGA-assigned group-specific threshold, giving the user additional confidence in those discoveries. When variable modifications are specified, the process is similar, except that for each representative peptide in the FDR-controlled list CONGA reports all the different variants of the same stem in each mass modification bin that are sufficiently high scoring (as in (a) above), where the bins are defined relative to the unmodified (stem) form.

2.2 CONGA successfully controls the FDR using target-decoy competition

TDC established itself as the canonical approach to FDR control in mass spectrometry analysis even before the corresponding mathematical theory was developed independently [1, 10]. That theory is based on the following two assumptions: (a) an incorrect peptide (one that is not present in the sample) that won its target-decoy competition is equally likely to be a target or a decoy peptide,

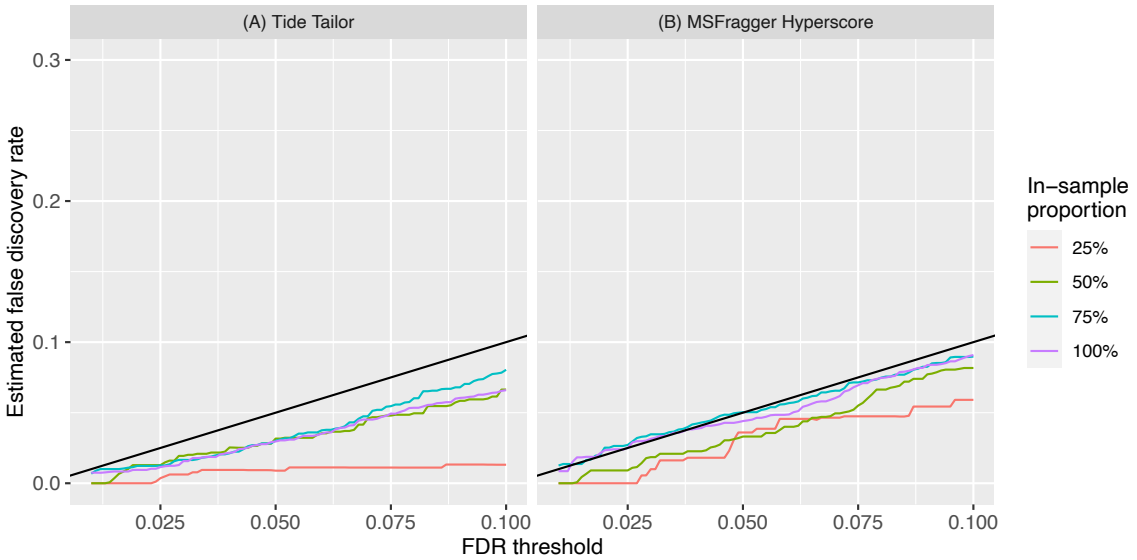

Figure 2: **CONGA's empirical FDR.** The estimated FDR when applying CONGA to our entrapment setup using (A) Tide and (B) MSFragger. The FDP is estimated at a range of FDR thresholds ($[0.01, 0.1]$), and its average over the 20 randomly generated decoys is the empirical FDR. Each solid curve represents the FDR using a different target database constructed with the specified proportion of the in-sample database. The estimated FDR is always below or just above the specified threshold (solid black line), indicating that the FDR is apparently controlled.

and (b) this happens independently of everything else (the scores of all the peptides, including the considered one, as well as all the target/decoy label of all other peptides).

CONGA relies on an extension of that theory, where the above independence is assumed to hold given the scores as well as the group structure. More precisely, we assume that, conditioned on all the winning (stem) peptide scores (step 2), on CONGA's group partition (step 3), and on all the target/decoy labels of the peptides that are truly in the sample, the peptides that won their competitions and are not in the sample are independently equally likely to be a target or a decoy win. Under these assumptions Group-walk, and hence CONGA, are guaranteed to control the FDR [8].

In practice these theoretical assumptions are naturally met with some approximation; thus, it is important that we test CONGA's ability to control the FDR using the same entrapment setup in which Percolator and PeptideProphet apparently fail to do so [7]. Specifically, we applied CONGA to PSMs generated by Tide and MSFragger, searching the same concatenated databases in both narrow and open mode. Reassuringly, the results of the CONGA analysis (Figure 2) show that the maximal estimated FDR is always below or very close to the corresponding FDR threshold, suggesting that CONGA successfully controls the FDR.

2.3 CONGA detects more peptides than narrow or open search

CONGA's primary goal is to combine the results from narrow and open searches of the same data to yield better statistical power than either of those two searches on its own. Accordingly, we compared the number of peptides detected by CONGA with the number detected by TDC on narrow and open search results. For this analysis, we employ the same PRIDE-20 datasets, and we used Tide, Comet and MSFragger.

This experiment shows a consistent improvement in power for CONGA over both narrow and open searching, using all three search engines (Figure 3A–F). Specifically, the median ratio of the

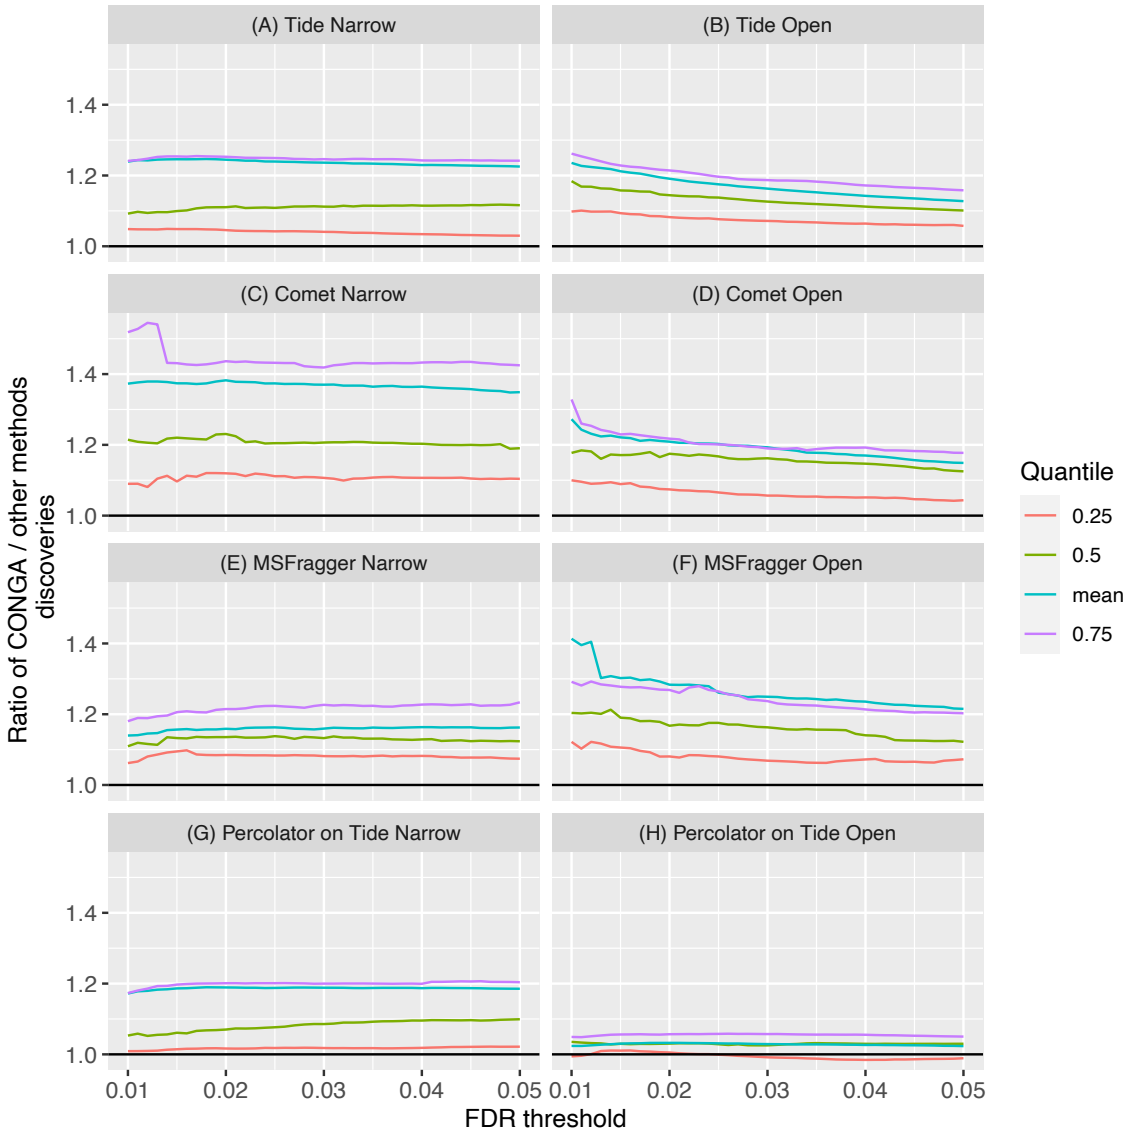

Figure 3: **CONGA detects more peptides than narrow or open search.** The quartiles and average, taken over the 20 PRIDE-20 datasets, of the mean ratio of CONGA-to-TDC discovered peptides. TDC was applied to the (A) Tide search results in narrow mode, (B) Tide search results in open mode, (C) Comet search results in narrow mode, (D) Comet search results in open mode, (E) MSFragger search results in narrow mode and (F) MSFragger search results in open mode. Panels (G) and (H) are similar, but compare CONGA with Percolator applied to the Tide narrow and open search results, respectively. For results that used Tide the mean ratio is taken with respect to each dataset's 20 decoys, while Comet's and MSFragger's searches used a single reversed set of decoys.

number of peptides discovered by CONGA relative to either the narrow or open search is always greater than 1. At 1% FDR, CONGA delivers a median increase over the narrow search of 9.23% for Tide, 21.45% for Comet, 10.95% for MSFragger, and increases over the open search of 18.40%, 17.73%, and 20.37%, respectively. Even at the lower quartiles, CONGA maintains an increase of at least 3% over TDC across all three search engines at both 1% and 5% FDR thresholds.

Perhaps the key challenge associated with open searching is that the large number of candidate peptides per spectrum leads to a loss of power relative to narrow searching. This occurs when a correctly matched peptide in the narrow search gets out-competed by a candidate from the wide window with a randomly high score. We therefore examined how many peptides discovered in the narrow search are lost in the corresponding list of CONGA discoveries. Specifically, we calculated the percentage of the peptides discovered by Tide with a narrow search that were missing from those reported by CONGA (Supplementary Figure S2). At 1% FDR, the median percentage loss is only 2.52%. While not completely negligible, this loss is clearly more than offset by CONGA's gains from the open search.

A key remaining question is how CONGA compares to a post-processor such as Percolator. Carrying out this comparison is challenging because, as pointed out in [7], Percolator's FDR control seems to be liberally biased. Thus, it is particularly interesting that, overall, our experiments suggest that CONGA marginally edges out Percolator when both are applied to the PRIDE-20 datasets (Figure 3G–H). Specifically, both post-processors use Tide open and narrow searches, where Percolator can only use one type of search at a time but is able to use significantly more features than CONGA can. For both open and narrow search, the median ratio of CONGA-to-Percolator discoveries is consistently greater than 1 across FDR thresholds from 1% to 5%.

2.4 CONGA accurately detects chimeric spectra

A chimeric spectrum is generated when the isolation following the MS1 scan fails to separate co-eluting but distinct peptides. During the subsequent fragmentation step, ions are generated from multiple parent peptides, and these ions are simultaneously recorded in the following MS2 scan. Those chimeric spectra often pose a challenge for traditional "one-spectrum-one-peptide" tools that do not account for them [4].

CONGA's design allows it to inherently process chimeras, identifying their multiple generating peptides, while still rigorously controlling the FDR in the overall reported list of detected peptides. To understand how this works, recall that CONGA considers multiple matches for each spectrum: one from the narrow search and up to $n_t = 5$ from the open search. The subsequent filtering process (Section 2.1) guarantees that all remaining PSMs are made of peptides that account for different peaks of the experimental spectrum. This requirement allows CONGA to distinguish between a truly co-generating peptide and a neighbor or "hitchhiking" peptide, where the latter offers a good match simply because its theoretical spectrum happens to share many peaks with the generating peptide. The multiple PSMs are then assigned to different groups, and CONGA's group-wise FDR control ensures the overall FDR level is below the selected threshold. In particular, for non-chimeric spectra the secondary matches that get through the filtering will typically be low-scoring and therefore below the specific group threshold, whereas for some of the chimeric spectra multiple peptides can score high enough to be above the group threshold. We refer to peptides that CONGA identifies using a single spectrum as "co-peptides." During group construction CONGA only considers the narrow and the top 2 open PSMs; thus, in principle, a chimeric spectrum can contribute to the identification of up to three peptides in three distinct groups, although in practice using a 1% FDR threshold this seems to be an extremely rare event. Note that a PSM that comes up in the open search can still have an essentially 0 mass-modification, so CONGA can detect

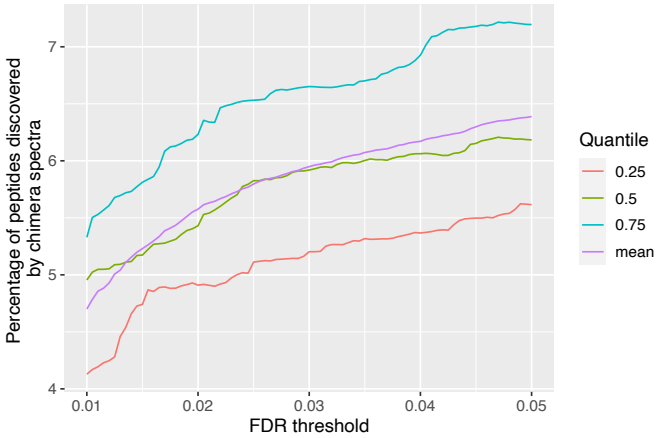

Figure 4: **The percentage of peptides CONGA discovers through chimeric spectra.** The quartiles and mean of the percentage of peptides detected by CONGA for which the top-scoring PSM is from a presumed chimeric spectra, plotted as a function of the FDR level. The quartiles and mean are taken over the 24 HEK293 spectrum files, and the percentages are taken with respect to the number of peptides discovered at the same FDR level by CONGA.

chimeric spectra regardless of whether each of the generating co-peptides includes a PTM.

To better understand CONGA's ability to detect chimeric spectra, we set out to quantify the proportion of accepted peptides that are discovered via chimeric spectra. For this analysis, we consider the 24 spectrum files of the HEK293 dataset [2], and we focus on the MSFragger search results. We observe that at a 1% FDR threshold a median of 4.96% of the total peptides discovered by CONGA were identified thanks to chimeric spectra, and at 5% the median increases to 6.18% (Figure 4).

Although CONGA provides theoretical guarantees related to FDR control, we wanted to further validate the co-peptides in comparison to peptides detected using standard TDC applied to a narrow window search. To this end, we used two types of predictions from the machine learning tool Prosit [9]: predicted peak intensities and predicted retention times. We hypothesized that the co-peptides should be indistinguishable from the peptides detected via TDC, in terms of the accuracy of both types of Prosit predictions. We tested this hypothesis by randomly selecting one of the 24 HEK293 spectrum files (Section 3.2.3), collecting all of the co-peptides from CONGA and the peptides from MSFragger at a 1% FDR threshold, and using Prosit to predict peak intensities and retention times for both sets of peptides. Because Prosit generally does not allow for peptides with mass modifications as inputs, we only considered co-peptides where both have a mass that is no more than 1 Da away from that of the corresponding spectrum's precursor.

We find that the accuracies of the Prosit predictions, for both peak intensities and retention times, are highly similar between the CONGA co-peptides and the MSFragger peptides (Figure 5). For the peak intensity prediction, visually there is very little that distinguishes between the two sets of discovered peptides (Figure 5A–B), and moreover, the overall correlation between the experimental and theoretical spectra in the two groups is essentially the same. For the retention time prediction, it is evident from Figure 5C–D that the relationship between the Prosit-predicted and experimental retention time is non-linear. Accordingly, we fitted a LOESS model, and we also used the Spearman correlation rather than the usual Pearson correlation. Again, we find the fit between the predicted and actual retention times seems to be invariant to whether the peptide was one identified by a presumed chimeric spectrum. The two corresponding Spearman correlation coefficients are also very close: 0.978 for the co-peptides and 0.972 for the TDC-discovered peptides.

Finally, in Supplementary Figures S3 and S4 we used PDV [15] to manually annotate some of those presumed chimeric spectra identified by CONGA.

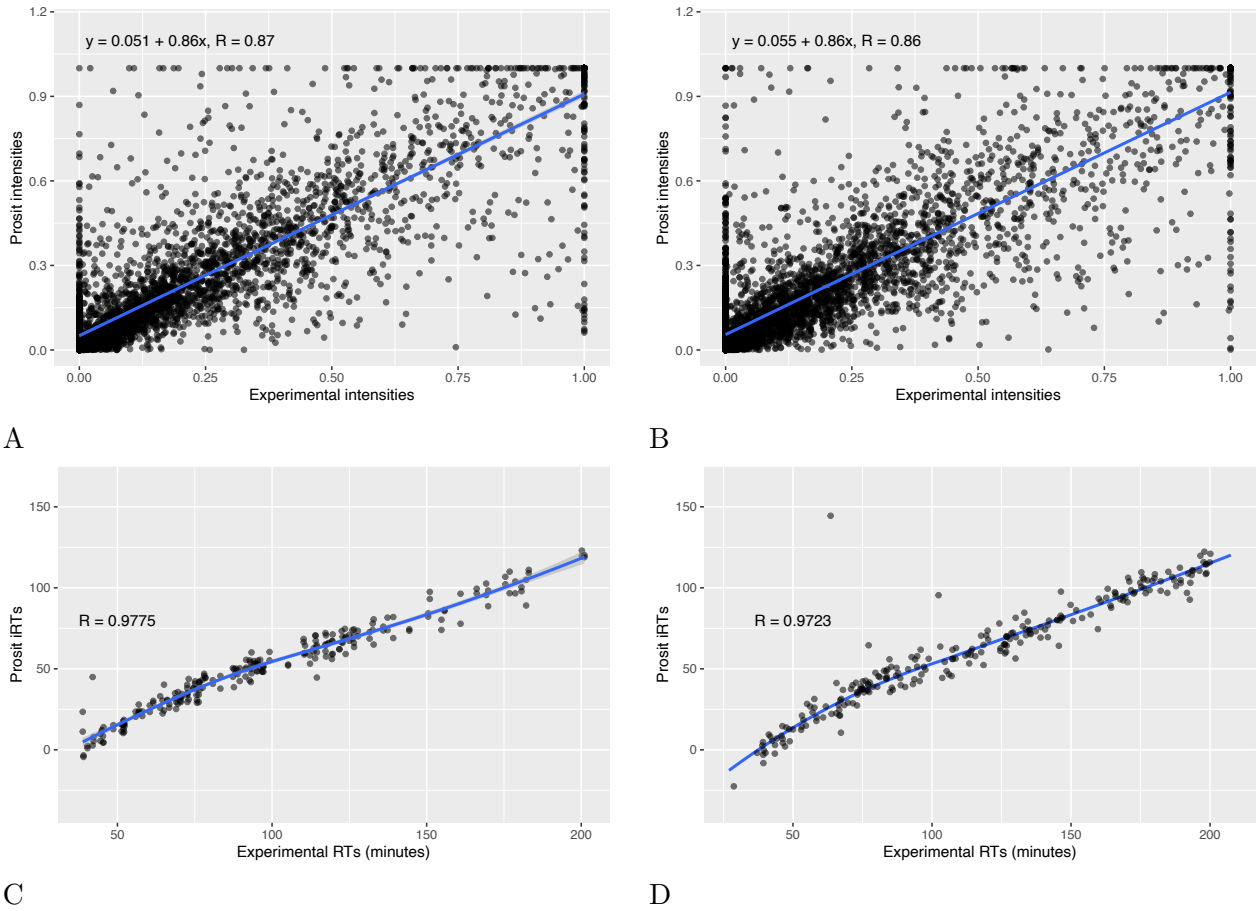

Figure 5: **Comparisons of Prosit-predicted and observed peak intensities and retention times.** (A) Scatter plot of experimental (x-axis) versus Prosit-predicted (y-axis) peak intensities from co-peptides discovered by CONGA at 1% FDR. (B) Similar scatter plot as panel A, but for peptides discovered by MSFragger in a narrow-window search at 1% FDR. For visualization, we only plot a randomly selected subset of the peak pairs, but the fitted line and correlation are computed over the entire set of peaks (see Section 3.2.4 for further details). (C) Scatter plot of observed versus Prosit-predicted retention times for co-peptides discovered by CONGA at 1% FDR. Note that technically we used Prosit's normalized retention time indices (which can be negative), but for simplicity we refer to them as predicted retention times [6]. (D) Similar plot as panel C, but for peptides discovered by MSFragger in a narrow-window search at 1% FDR. Again, panel D only includes a randomly selected subset of points, but the LOESS-fitted curve and the Spearman correlation are computed over the entire set of peptides in both C and D.

3 Methods

3.1 The CONGA algorithm

3.1.1 Filtering neighbors

CONGA's first step focuses on removing neighbors: a peptide is considered a *neighbor* of another peptide if the two theoretical spectra share a non-trivial proportion of peaks. Neighbors pose a problem, particularly if, as in CONGA's case, we consider more than the top match to a spectrum σ: if a peptide offers a good match to σ then it is likely that so do all its neighbors.

Supplementary Algorithm 2 implements CONGA's filtering process. Its input is the list of peptides corresponding to the top PSMs associated with each spectrum. Specifically, associated with each spectrum σ are $k^\sigma \leq n_t+1$ unique PSMs, of which $\leq n_t$ are from the open search and ≤ 1 from the narrow search (we remove an open PSM if it happens to coincide with the narrow PSM). Sorting those k^σ peptides by decreasing PSM scores, CONGA sequentially filters out peptides that share more than a fraction of τ peaks with a higher ranking peptide (we used $\tau = 0.05$, a choice we discuss in Section 4). Following [16], we define the fraction of shared peaks between two peptides π_1 and π_2, t_{12}, as the proportion of shared b- and y-ions:

$$t_{12} := \frac{2B_{12}}{B_1 + B_2}, \tag{1}$$

where B_i is the number of charged b- and y-ions associated with π_i between 200 m/z and 3000 m/z, and B_{12} is the number of such shared ions. If the presumed charge of π_i is $+1$ or $+2$, then B_i considers only singly charged b- and y-ions, and if the charge is $+3$ or more, then B_i considers both singly and doubly charged ions.

Comet's E-value [5] and Tide's Tailor score [19] normalization of the XCorr score will generally differ between the open and narrow searches. Therefore, to ensure a fair comparison when sorting the combined list of k^σ peptides we use the PSMs' XCorr score. The PSMs that are not filtered out are then reverted back to their original E-value or Tailor score.

3.1.2 Dynamic Competition

CONGA uses both open and narrow search PSMs generated by searching the concatenated dataset. After filtering, the score assigned to each peptide is that of the maximal PSM associated with the peptide (or the smallest possible score if no such PSM exists). This selection of the top-scoring target or decoy PSM per spectrum is the first round of competition.

The second level of competition differs depending on whether variable modifications are specified. When no variable modifications are specified, the second target-decoy competition takes place at the peptide level: each target peptide is compared with its corresponding decoy, and only the higher scoring of the two is kept. On the other hand, as mentioned in Section 2.1, when variable modifications are specified, CONGA clusters together all peptides that share a common stem form. Each cluster is represented by its maximally scoring peptide. In this case, the second target-decoy competition takes place at the representative rather than the peptide level: each target representative is compared with its corresponding decoy. Note that, regardless of whether variable modifications are considered, even when the FDR analysis is done at the representative level, CONGA's augmented list includes all sufficiently high scoring PSMs with distinct unaccounted-for mass modification as described in Supplementary Algorithm 6.

After the second-level competition takes place, each winning peptide is assigned four values: its winning score (W_i), the target/decoy label ($L_i = 1$ if the winner is a target object), a narrow-search indicator ($N_i = 1$ if the maximal PSM of the winner came from a narrow search), and the

spectrum-specific rank R_i of the winning peptide among the top PSMs reported for that spectrum in the original search file. Because we only keep the top PSM per spectrum in the narrow search, $R_i = 1$ if $N_i = 1$, but at any rate CONGA only further considers peptides/representatives for which $R_i \leq 2$. Finally, δ_i, is the mass difference between the precursor mass and the peptide.

3.1.3 Group construction

Group-walk, and by extension CONGA, gains a power boost when the data inherently segregates into distinct groups. We expect the narrow peptides ($N_i = 1$) to be distinct from the open ones; hence, those peptides form their own "narrow group". The other groups, made up of the open peptides ($N_i = 0$), are essentially defined based on the mass differences δ_i. As outlined next, and detailed in Supplementary Algorithm 3, CONGA accomplishes this grouping in two phases, in which it dynamically splits and merges the open peptides according to δ_i, in an attempt to define distinct groups in terms of their winning scores distribution. Note that throughout this section we refer only to "peptides," with the understanding that we mean representatives when using variable modifications.

With the narrow group defined, CONGA first partitions the open peptides with rank $R_i = 1$ into a set of disjoint bins according to their associated mass differences, δ_i. Specifically, each bin is centered at $k\lambda$, an integer multiple of $\lambda := 1.0005079/4$, and it contains all peptides with $\delta_i \in (k\lambda - \lambda/2, k\lambda + \lambda/2]$. We then rank the bins in terms of their occupancy rate from largest to smallest, and we denote by b_i the rank of the bin containing the ith considered peptide.

Relying on Group-walk, CONGA cannot make efficient use of a group which is too small: it is better to merge such a group with another one. CONGA's minimal group size is set to $2K$, where K is the window size used in Group-walk (Section 3.1.4: $K = 40$ by default). By merging any bin that is smaller than this threshold with the subsequently ranked bin, we can assume without loss of generality that every bin, except possibly the lowest ranked one, clears the minimal group size threshold.

Starting with the top ranked bin (i.e., the most common mass difference), CONGA sequentially defines a candidate set C as the set of peptides populating the highest ranked bin that is yet to be processed. Defining the current left-over set, L, as the set of peptides in all lower ranked bins, CONGA applies the Kolmogorov-Smirnov test to gauge the similarity between the winning scores of (i) C and L, and (ii) C and each group of open peptides previously defined in this sequential process. For example, when considering the 3rd ranked bin this would typically generate $1 + 2 = 3$ p-values, one for each application of the test.

If all the p-values are less than or equal to a pre-specified threshold $\beta := 0.01$, then the winning scores in C are deemed distinct enough to merit defining C as a new group, and the process moves on to the next candidate set. Otherwise, if the the p-value of the test comparing C and L is $\leq \beta$ but the p-values of the similarity tests between C and some of the previously defined groups are $> \beta$, then C is deemed to be insufficiently distinct from any one of those groups. In that case, one of those groups is selected at random, C is merged with it, and the process moves on to the next candidate set. Otherwise, the p-value of the similarity test between C and L is $> \beta$ so C and L are joined into a new group, and this phase of the group construction ends.

If the process did not terminate before C coincides with the lowest ranked bin, then at that point C is added to the most recently defined group. Note that in practice the winning PSM scores often have ties which implies that the Kolmogorov Smirnov p-values are only approximate. However, this is not overly concerning because we are only using the tests here to merge groups rather than for hypothesis testing per se.

Next, CONGA groups together all open peptides with rank $R_i = 2$ with an associated δ_i that

would place it in one of the top ranking bins that were processed in the previous phase (excluding the terminating step). If this group is larger than the $2K$ threshold then it is defined as the last group; otherwise, it is merged with the most recently constructed group (which is usually the final candidate set C and the left-over set L).

3.1.4 Group-walk

After the groups are constructed, we can proceed with applying the Group-walk procedure of [8], added here for convenient reference as Supplementary Algorithm 4. Briefly, Group-walk first orders the peptides (representatives) within each group $g \in \{1, \ldots, n_g\}$ in increasing winning scores order, so that with n_g being the number of groups, W_i^g the score of the ith peptide in group g, and m_g the number of peptides in group g, $W_1^g \leq \cdots \leq W_{m_g}^g$.

It then initializes a *frontier* vector $\tilde{k} = \left(k_1, \ldots, k_{n_g}\right) \coloneqq (1, \ldots, 1)$, which delineates the "active set" of peptides that the front has not yet passed over: all peptides whose (group g) index i satisfy $i \geq k_g$ (so $W_i^g \geq W_{k_g}^g$). The peptides that are not in the active set cannot be discovered (the front passed over them), but as explained below, they are still useful in determining how to advance the front.

Group-walk advances the front one coordinate at a time, and following each such advancement it estimates the FDR among the target peptides in the active set as the number of decoys in this set (plus 1) divided by the number of target peptides in this set. If this estimated FDR is $\leq \alpha$ Group-walk stops and reports all target peptides in the active set. Otherwise, Group-walk selects a group g for which it will advance the front by one peptide: $k_g \coloneqq k_g + 1$.

Initially the group that is advanced is selected sequentially: Group-walk rotates through the groups. However, once the frontier reaches $\tilde{k} = (K, \ldots, K)$, where K is the "peek-back" window size ($K = 40$ here), the decision of which group to advance is made based on the last K observed peptides in each group. Specifically, Group-walk advances the group for which the number of decoy wins among the last K peptides is maximized.

3.2 Datasets

3.2.1 Entrapment run searches

Entrapment runs for CONGA were performed using data taken from the standard protein mix database, ISB18 [13]. We used the nine `.ms2` spectrum files taken from [16] which were originally sourced from the Mix 7 spectrum files downloaded from `https://regis-web.systemsbiology.net/PublicDatasets/18_Mix/Mix_7/ORBITRAP/` (this excluded the tenth smallest spectrum file). These `.ms2` spectrum files were subsequently used for Tide. For compatibility with MSFragger, we directly downloaded the same nine spectrum files in ThermoRaw format and converted them to a single combined, as well as separate, `.mzML` files using MSConvert 3.0.22314 with the vendor peak-picking filter using the default settings.

The in-sample database comprised of the extended protein database containing the 18 peptides and 30 additional hitchhiker peptides downloaded from `https://regis-web.systemsbiology.net/PublicDatasets/database`. We used the castor plant proteome as the entrapment database as in [16]. Tide-index was used to digest the in-sample and entrapment database using trypsin, for Tide, and strict-tryspin (specified with `trypsin/p`), for MSFragger. For each digest rule, four random subsets of the in-sample peptides of varying size, plus the entrapment peptides were used to create a combined target peptide database for which 20 randomly shuffled decoy indices were created, again using Tide-index. Param-medic [18] was used to determine any variable modifications, though none were detected. Any entrapment peptide that was identical (up to any leucine/isoleucine

Trypsin for Tide			
In-sample proportion	In-sample peptides	Entrapment peptides	In-sample:entrapment
100%	898	571298	1:636
75%	674	571298	1:848
50%	449	571298	1:1272
25%	225	571298	1:2539
Strict-trypsin for MSFragger			
In-sample proportion	In-sample peptides	Entrapment peptides	In-sample:entrapment
100%	919	584350	1:636
75%	690	584350	1:847
50%	460	584350	1:1270
25%	280	584350	1:2087

Table 1: **The number of in-sample target peptides and entrapment-peptides for each combined target peptide database used in the entrapment runs.** The table reports the number of in-sample target peptides and entrapment target peptides for each of the four combined target databases, for each digest rule. The first column reports the proportion of the total in-sample target peptides used in the combined target database, the second is the number of these in-sample peptides used, the third column is the number of entrapment peptides, and the last column is the ratio of in-sample peptides to entrapment peptides.

substitution) to an ISB18 peptide was removed from the entrapment database. Table 1 lists the resulting number of in-sample and entrapment peptides in each combined target peptide database, along with their in-sample-to-entrapment ratio.

Tide search was used to search the nine `.ms2` spectrum files against each of the combined target and decoy databases prepared using trypsin with the following settings: for narrow searches we used `--auto-precursor-window warn --auto-mz-bin-width warn --use-tailor-calibration T --concat T`, and for open searches we used `--auto-mz-bin-width warn --precursor-window-type mass --precursor-window 100 --top-match 5 --concat T --use -tailor-calibration T`. MSFragger v3.3 was used to search the combined `.mzML` and the separate `.mzML` spectrum files against each of the combined target and decoy databases prepared using strict-trypsin. Because we used tide-index to prepare these databases, we need to use the following settings to ensure MSFragger has the same digestion parameters. For narrow searches we used `allowed_missed_cleavage = 0, digest_min_length = 6, digest_max_length = 50, digest_mass_range = 200.0 7200.0, allowed_missed_cleavage = 0` with all other options set to default, and for open searches we used `localize_delta_mass = 0, allowed_ missed_cleavage = 0, digest_min_length = 6, digest_max_length = 50, digest_mass _range = 200.0 7200.0, allowed_missed_cleavage = 0, output_report_topN = 5`, with all other option set to default. No variable modifications were set, so CONGA's analysis was done at the peptide level. Note we used MSFragger v3.3 at the time, as the subsequent version at the time (v3.4) did not allow for `output_report_topN > 1` for DDA data. We used Tide in Crux v4.1.decd99ff.

3.2.2 PRIDE-20 searches

We downloaded 20 high-resolution spectrum files from the Proteomics Identifications Database, PRIDE [17]. Seven of these spectrum files were taken from [8] and were originally obtained by randomly selecting a spectrum file from randomly selected PRIDE projects (submitted no earlier than 2018). The remaining 13 spectrum files were similarly obtained by randomly selecting a single spectrum file from a randomly selected PRIDE project (submitted no earlier than 2019). The sampling was constrained to generate a collection of 10 spectrum files that had modifications and

Project ID	Spectrum file	Var
PXD008920	QEP1_ZADAM_deg1_2_1_170615.mzid_QEP1_ZADAM_deg1_2_1_-170615.MGF	F
PXD008996	A431_01uM_DON_3_2.mgf	F
PXD010504	QX05437.mgf	F
PXD014277	Q06965_MS18-017_37_J9(55).mgf	F
PXD016274	190322-SI-0149-F5-01.mgf	F
PXD019354	Progenesis-Trophoplast-top5-140416.mgf	F
PXD024284	p2830_NaClvsHank_mascot.mgf	F
PXD025130	Q26431_MS20-025_Virus-purif_1.mgf	F
PXD029319	MP_16072020_LvN_S_layer_gps_5_DDA01.mgf	F
PXD030118	Q26756_MS20-013_C15.mgf	F
PXD002470	Mudpit_LC-iTRAQ1.RAW_(F004785).mzid_Mudpit_LC-iTRAQ1.RAW_-(F004785).MGF	T
PXD006856	export_Mzidentml_F134282.mgf	T
PXD010413	F373129.mgf	T
PXD012528	04222014_111_AF_160_22.mgf	T
PXD012611	Temp-3-Speclib-1.mzid_Temp-3-Speclib-1.MGF	T
PXD013274	PDL-2-11.mzid_PDL-2-11.MGF	T
PXD019186	QE1_Chrysi_Exp1_Sample12b.raw.-1.mgf	T
PXD022257	Sitaset1F6.mgf	T
PXD023571	Experiment_new_run_P6-GFP-PBANKA_1028300-Schizont-8hrs_from_OT_-200625_105_MS20-039_01.mzid_Experiment_new_run_P6-GFP__PBANKA_-1028300__Schizont__8hrs_from_OT_200625_105_MS20-039_01.MGF	T
PXD026895	191004_MK_ChRS19_100Z10C_pink_A.mzid_191004_MK_ChRS19 _100Z10C_-pink_A.MGF	T

Table 2: **The PRIDE-20 data set** The list of 20 spectrum files used in the PRIDE-20 data set and their associated project IDs. The "Var" column indicates whether the data includes variable modifications.

10 spectrum files for which no modifications were detected, as determined in both cases by Param-medic [18]. The protein FASTA database files were also downloaded from the associated PRIDE projects or in the case of human data, the UniProt database UP000005640 was used (downloaded 9/11/2021). Table 2 reports the list of the 20 spectrum files and PRIDE projects used.

For each of the PRIDE-20 data sets, we used Tide-index 20 times to create 20 different target-decoy peptide databases using `--auto-modifications T` to call Param-medic and apply the variable modifications. For narrow searches using Tide, we used the following options `--auto-precursor-window warn --auto-mz-bin-width warn --use-tailor-calibration T --concat T` and for open searches using Tide we used `--auto-mz-bin-width warn --precursor-window-type mass --precursor-window 100 --use-tailor-calibration T --concat T`.

For searches using Comet, we used the builtin version within Crux. For narrow searches we used the options `--decoy_search 1 --auto_peptide_mass_tolerance warn --allowed_ missed_cleavage 0 --auto_fragment_bin_tol warn --auto_modifications T --peptide _mass_units 2` and for open searches we used `--decoy_search 1 --peptide_mass_tolerance 100 --allowed_missed_-cleavage 0 --auto_fragment_bin_tol warn --auto_modifications T --peptide_mass_units 0`. Comet digests and produces decoy peptides by reversing each target sequence. Hence, only a single target-decoy database was used for each of the PRIDE-20 data sets. PSMs matched to peptides that were identified as both a target and decoy were deleted. In the case of E-values, lower scores

indicate a better PSM and so we take the negative of these scores when applying TDC or CONGA.

Lastly for searches using MSFragger we converted each `.mgf` file into `.mzmL` files using MSConvert with default settings. We constructed a decoy protein database by taking each target protein and reversing the position of the amino acids that lie between K, R or the N- and C- terminal, while keeping the position of K and R fixed. Importantly, the digestion of the target and decoy proteins allow us to pair each resulting target peptide with a corresponding decoy peptide. This is not achievable if we reverse the entire target protein sequence, which is the behaviour of many softwares. For narrow searches, we used the following options `allowed_missed_cleavage = 0, use_topN_peaks = 100` and for open searches we used `output_report_topN = 5, allowed_missed_cleavage = 0, use_topN_peaks = 100`. Variable modifications were set according to those detected by Parammedic using Tide-index. All other options were set to the default settings. Similar to Comet, PSMs matched to peptides that were identified as both a target and decoy were deleted.

3.2.3 HEK293 searches

We downloaded the 24 HEK293 spectrum files from `https://ftp.pride.ebi.ac.uk/pride/data/archive/2015/06/PXD001468` which were converted to `.mzML` format using MSConvert with the vendor peak-picking filter using the default settings. We used the UniProt database UP000005640 (downloaded 18/05/2022). As discussed in Section 3.2.2, we prepared a combined target-decoy protein database by taking each target protein and reversing the position of the amino acids that lie between K, R or the N- and C- terminal, while keeping the position of K and R fixed. We ran narrow MSFragger searches for each of the 24 spectrum files using the settings `precursor_mass_lower = -100, precursor_mass_upper = 100, allowed_missed_cleavage = 1, use_topN_peaks = 100` and open searches using the settings `precursor_mass_lower = -500, precursor_mass_upper = 500, precursor_mass_units = 0, localize_delta_mass = 0, allowed_missed_cleavage = 1, output_report_topN = 5`. No variable modifications were considered, and all other options were set to default. These settings were used to follow as best we reasonably could the setup used by [14].

3.2.4 Chimera spectra analysis

We used Prosit to validate our co-peptide discoveries. We first randomly selected one of the HEK293 datasets (`b1924_293T_proteinID_04A_QE3_122212.mzML`) and converted it to `mgf` format with MSConvert for easy importing in R. We determined the list of co-peptide discoveries using CONGA at the 1% FDR level, and filtered this list for co-peptides that were both less than 1 Da away from the precursor mass of the corresponding experimental spectrum. This is because Prosit does not generally support peptides with mass-modifications. In addition to these co-peptides, we determined their precursor charge and set the collision energy for each peptide as 35%, as described in [2]. The peptides, charges, and collision energy were then submitted to Prosit. The resulting output returned the predicted fragment m/z locations, the relative intensities and the normalized retention times (iRT). For each co-peptide, we used a 0.01 Th fragment tolerance to pair the Prosit fragment and experimental fragment locations, setting the experimental intensity to zero if no experimental fragment exists. We then transformed the experimental intensities into relative intensities, by dividing the experimental peak intensities by the maximal experimental peak intensity that was paired with a Prosit peak.

To benchmark our co-peptide discoveries, we determined the list of narrow-discoveries using TDC at the 1% FDR level, and similarly used Prosit along with the subsequent analysis as described above for the co-peptide discoveries using CONGA.

3.2.5 Target-decoy competition

For the TDC procedure, we employ the double competition protocol, "PSM-and-peptide" as originally described in [16] and summarized for convenience below as well in pseudocode in Supplementary Algorithm 1a. For the PSM-level competition, we determine the best peptide-spectrum match to the target database and decoy database separately, and then record the best scoring peptide-spectrum match out of the two. This is achieved implicitly by taking the top 1 PSM for each spectrum when searching against the concatenated target-decoy database. For the subsequent peptide-level competition, we define a score for each peptide as the maximal PSM score associated to that peptide. Then for each target-decoy pair, we take the peptide with highest peptide score. The resulting peptides are then ordered according to their peptide score from largest to smallest. Denoting the resulting scores as $W_1 \geq \cdots \geq W_n$, we determine the largest index k that the estimated FDR is $\leq$ the threshold α, and report all target peptides up to and including that k. More specifically, denoting $L_i = 1$ as a target peptide and $L_i = -1$ as a decoy peptide, then all target peptides with rank $i \leq K_{\mathrm{TDC}}$ are reported, where

$$K_{\mathrm{TDC}} = K(\alpha) := \max\{k : \frac{\#\{i \leq k : L_i = -1\} + 1}{\#\{i \leq k : L_i = 1\} \vee 1} \leq \alpha\}, \tag{2}$$

and $x \vee y := \max\{x, y\}$. For convenience we provide a pseudocode

As described in Section 3.1.3 of CONGA's method, if variable modifications are specified, the second, peptide-level competition is replaced with the representative-level one (Supplementary Algorithm 1b).

3.2.6 Percolator settings

For each of the Tide-searches where Percolator was subsequently used, we constructed `.pin` files by either specifying `--pin-output T` during Tide-search or by using the `make-pin` utility command in Crux. When creating `.pin` files from an open search we used the default number of top 5 PSMs which allows for generating the feature `deltLCn` (the difference between the XCorr score of the top 1 PSM and the last ranked PSM). We then kept only the top 1 PSM for each spectrum to overcome the problem of neighbors. For narrow searches, we directly used the top 1 PSM. Lastly, to each PSM in the `.pin` file, we set the `enzInt` feature (the number of internal enzymatic sites) to zero as explained in [3].

For each of the .pin files, we used the following options in the crux percolator command `--only-psms T, --tdc F`. This reports only the PSMs and does not apply the TDC procedure within Percolator since it does not use the double competition protocol, "PSM-and-peptide". Instead, we apply our TDC procedure as described in the previous Section 3.2.5 to the resulting list of PSMs, using the learned percolator score as the PSM score.

4 Discussion

Open search is an increasingly popular approach to detect peptides that harbor unexpected post-translational modifications. However, applying both open and narrow search to the same data, followed by TDC, often yields fewer discovered peptides for the open search than the narrow one. Only when applying a post-processor such as Percolator or PeptideProphet do we see a fairly consistent improvement using an open search. On the other hand, we demonstrated recently that both Percolator and PeptideProphet apparently fail to control the FDR in our entrapment experiments.

CONGA is designed as an alternative: it combines the results from the open and narrow searches in an approach that is built on a recently developed extension of the canonical TDC. CONGA's control of the FDR draws on rigorous theoretical results; however, in practice, some of the necessary theoretical assumptions are only approximately satisfied. For example, the target database can often contain neighbors to in-sample peptides, and such neighbors can compromise the assumption that an incorrect peptide is equally likely to be a target or a decoy. With its PSM-filtering as well as dynamic-level competition CONGA takes steps to mitigate the impact of neighbors.

Notably, in practice we see that CONGA apparently controls the FDR in the same entrapment experiments in which Percolator and PeptideProphet failed. Interestingly, in spite of its tighter FDR control, CONGA's power is comparable to, if not marginally better, than Percolator (we did not compare CONGA's power with PeptideProphet's). This is particularly surprising given that Percolator uses significantly more information than CONGA does, and at the same it suggests a promising avenue for future gains by incorporating into CONGA the additional features that Percolator uses.

When introducing Group-walk we showed that we can use the mass difference between the precursor and the peptide to increase the power of the open search. However, the resulting open analysis occasionally fell significantly behind the results of narrow-TDC (Supplementary Figure S1). The current approach addressed that problem by co-opting the narrow search results into the analysis, allowing for multiple PSMs per spectrum, and applying a significantly more sophisticated, and automated, group construction procedure.

Yu et al. [20] recently introduced a localization-aware open search procedure, which matches the peaks of fragment ions containing up to one unknown PTM in the experimental spectrum to fragment ions in the theoretical peptide. This approach has the potential to significantly improve the open search analysis, and therefore it is fair to ask whether it makes CONGA redundant. However, we demonstrate in Supplementary Figure S6 that even this localization-aware open search procedure benefits from using CONGA over TDC.

CONGA requires as input a pairing between targets and decoys, which unfortunately is ill-defined if the decoy peptides were generated by first reversing or shuffling the entire protein. For this reason, CONGA only supports Tide, Comet and MSFragger, and the latter only if the target-decoy database is prepared beforehand, as explained in the last paragraph of Section 3.2.2).

We finish with a couple of tangential comments. The first is regarding the fraction of shared peaks threshold we use for filtering PSMs ($\tau := 0.05$). This is probably lower than what is needed to get rid of neighbors, but setting it to a low value helps to filter out many secondary, incorrect PSMs in the same way that the added spectrum-level competition helps the PSM-and-peptide method improve on peptide-only competition [16]. Although this low threshold incurs the cost of potentially throwing away some correctly matched co-peptides due to chimeric spectra, it overall delivered more discoveries than the higher threshold of $\tau = 0.25$ (data not shown).

As explained in Section 2.1, CONGA's dynamic-level competition is motivated by theoretical FDR-control considerations. However, Supplementary Figure S5A shows that in the 10 PRIDE-20 datasets whose analysis invoked variable modifications, there was little difference between the representative-level and peptide-level analyses. Specifically, we examined the power loss due to collapsing-to-representative, defined as one minus the ratio of representative-to-peptide CONGA discoveries (averaged over the 20 decoys). At the 1% FDR level, and expressed as percentages, we see median losses of 0.134% (Tailor) and 0.208% (XCorr), while at the 5% FDR the corresponding median percentage losses are 0.450% and 0.543%. Panel B of the same figure shows the effect is even milder when used in conjunction with narrow TDC. This is particularly reassuring given that all power analyses in this paper applied the representative-level analysis to *all compared methods*

when the dataset was deemed to have variable modifications.

References

[1] R. F. Barber and Emmanuel J. Candès. Controlling the false discovery rate via knockoffs. *The Annals of Statistics*, 43(5):2055–2085, 2015.

[2] J. M. Chick, D. Kolippakkam, D. P. Nusinow, B. Zhai, R. Rad, E. L. Huttlin, and S. P. Gygi. A mass-tolerant database search identifies a large proportion of unassigned spectra in shotgun proteomics as modified peptides. *Nature Biotechnology*, 2015. Epub ahead of print.

[3] Yulia Danilova, Anastasia Voronkova, Pavel Sulimov, and Attila Kertész-Farkas. Bias in false discovery rate estimation in mass-spectrometry-based peptide identification. *Journal of proteome research*, 18(5):2354–2358, 2019.

[4] V. Dorfer, P. Pichler, T. Stranzl, J. Stadlmann, T. Taus, S. Winkler, and K. Mechtler. MSAmanda, a universal identification algorithm optimized for high accuracy tandem mass spectra. *Journal of Proteome Research*, 13(8):3679–3684, 2014.

[5] J. K. Eng, T. A. Jahan, and M. R. Hoopmann. Comet: an open source tandem mass spectrometry sequence database search tool. *Proteomics*, 13(1):22–24, 2012.

[6] Claudia Escher, Lukas Reiter, Brendan MacLean, Reto Ossola, Franz Herzog, John Chilton, Michael J MacCoss, and Oliver Rinner. Using iRT, a normalized retention time for more targeted measurement of peptides. *Proteomics*, 12(8):1111–1121, 2012.

[7] J. Freestone, W. S. Noble, and U. Keich. Re-investigating the correctness of decoy-based false discovery rate control in proteomics tandem mass spectrometry. In preparation, 2023.

[8] J. Freestone, T. Short, W. S. Noble, and U. Keich. Group-walk: a rigorous approach to group-wise false discovery rate analysis by target-decoy competition. *Bioinformatics*, 38(Supplement 2):ii82–ii88, 09 2022.

[9] S. Gessulat, T. Schmidt, D. P. Zolg, P. Samaras, K. Schnatbaum, J. Zerweck, T. Knaute, J. Rechenberger, B. Delanghed, A. Huhmer, U. Reimer, H. Ehrlich, S. Aiche, B. Kuster, and M. Wilhelm. Prosit: proteome-wide prediction of peptide tandem mass spectra by deep learning. *Nature Methods*, 16(6):509, 2019.

[10] K. He, Y. Fu, W.-F. Zeng, L. Luo, H. Chi, C. Liu, L.-Y. Qing, R.-X. Sun, and S.-M. He. A theoretical foundation of the target-decoy search strategy for false discovery rate control in proteomics. *arXiv*, 2015. https://arxiv.org/abs/1501.00537.

[11] L. Käll, J. Canterbury, J. Weston, W. S. Noble, and M. J. MacCoss. A semi-supervised machine learning technique for peptide identification from shotgun proteomics datasets. *Nature Methods*, 4:923–25, 2007.

[12] A. Keller, A. I. Nesvizhskii, E. Kolker, and R. Aebersold. Empirical statistical model to estimate the accuracy of peptide identification made by MS/MS and database search. *Analytical Chemistry*, 74:5383–5392, 2002.

[13] J. Klimek, J. S. Eddes, L. Hohmann, J. Jackson, A. Peterson, S. Letarte, P. R. Gafken, J. E. Katz, P. Mallick, H. Lee, A. Schmidt, R. Ossola, J. K. Eng, R. Aebersold, and D. B. Martin. The standard protein mix database: a diverse data set to assist in the production of improved peptide and protein identification software tools. *Journal of Proteome Research*, 7(1):96–1003, 2008.

[14] A. T. Kong, F. V. Leprevost, D. M. Avtonomov, D. Mellacheruvu, and A. I. Nesvizhskii. MSFragger: ultrafast and comprehensive peptide identification in mass spectrometry-based proteomics. *Nature Methods*, 14(5):513–520, 2017.

[15] K. Li, M. Vaudel, B. Zhang, Y. Ren, and B. Wen. PDV: an integrative proteomics data viewer. *Bioinformatics*, 35:1249–1251, April 2019.

[16] Andy Lin, Temana Short, William Stafford Noble, and Uri Keich. Improving peptide-level mass spectrometry analysis via double competition. *Journal of Proteome Research*, 21(10):2412–2420, 2022.

[17] L. Martens, H. Hermjakob, P. Jones, M. Adamsk, C. Taylor, D. States, K. Gevaert, J. Vandekerckhove, and R. Apweiler. PRIDE: The proteomics identifications database. *Proteomics*, 5(13):3537–3545, 2005.

[18] D. H. May, K. Tamura, and W. S. Noble. Detecting modifications in proteomics experiments with Param-Medic. *Journal of Proteome Research*, 18(4):1902–1906, 2019.

[19] P. Sulimov and A. Kertész-Farkas. Tailor: A nonparametric and rapid score calibration method for database search-based peptide identification in shotgun proteomics. *Journal of Proteome Research*, 19(4):1481–1490, 2020.

[20] F. Yu, G. C. Teo, A. T. Kong, S. E. Haynes, D. M. Avtonomov, D. J. Geiszler, and A. I. Nesvizhskii. Identification of modified peptides using localization-aware open search. *Nature Communications*, 11(1):4065, August 2020.