

MS²Rescore: Data-Driven Rescoring Dramatically Boosts Immunopeptide Identification Rates

Authors

Arthur Declercq, Robbin Bouwmeester, Aurélie Hirschler, Christine Carapito, Sven Degroeve, Lennart Martens, and Ralf Gabriels


Correspondence

lennart.martens@vib-ugent.be

Graphical Abstract

In Brief

The integration of newly trained immunopeptide MS²PIP models, DeepLC, and Percolator into one software package called MS²Rescore allows for a significant boost in immunopeptide identification rate as well as a substantial increase in specificity. MS²Rescore is search engine-agnostic and unbiased toward HLA types. MS²Rescore, therefore, shows great promise to extend the current neo- and xeno-epitope landscape in existing and future immunopeptidomics experiments.



Highlights

- MS²Rescore significantly boosts immunopeptide identification rates
- Data-driven post-processing allows for a ten-fold increase in specificity
- MS²PIP and DeepLC predictors are integrated with Percolator post-processing
- MS²Rescore accepts identification results from MaxQuant, PEAKS, MS-GF+ and X!Tandem
- MS²Rescore shows great promise to extend current neo- and xeno-epitope landscapes

MS²Rescore: Data-Driven Rescoring Dramatically Boosts Immunopeptide Identification Rates

Arthur Declercq^{1,2}, Robbin Bouwmeester^{1,2}, Aurélie Hirschler³,
Christine Carapito³, Sven Degroeve^{1,2}, Lennart Martens^{1,2,*}, and Ralf Gabriels^{1,2}

Immunopeptidomics aims to identify major histocompatibility complex (MHC)-presented peptides on almost all cells that can be used in anti-cancer vaccine development. However, existing immunopeptidomics data analysis pipelines suffer from the nontryptic nature of immunopeptides, complicating their identification. Previously, peak intensity predictions by MS²PIP and retention time predictions by DeepLC have been shown to improve tryptic peptide identifications when rescoring peptide-spectrum matches with Percolator. However, as MS²PIP was tailored toward tryptic peptides, we have here retrained MS²PIP to include nontryptic peptides. Interestingly, the new models not only greatly improve predictions for immunopeptides but also yield further improvements for tryptic peptides. We show that the integration of new MS²PIP models, DeepLC, and Percolator in one software package, MS²Rescore, increases spectrum identification rate and unique identified peptides with 46% and 36% compared to standard Percolator rescoring at 1% FDR. Moreover, MS²Rescore also outperforms the current state-of-the-art in immunopeptide-specific identification approaches. Altogether, MS²Rescore thus allows substantially improved identification of novel epitopes from existing immunopeptidomics workflows.

The immune system is a complex, yet remarkable system that protects us from both invaders from outside the body, that is, pathogens, as well as from inside the body, that is, malignancies (1). Increased understanding of the immune system allowed for great medical achievements such as vaccination, which is currently available for over 29 diseases, enabled the eradication of smallpox, and prevents over 3 million deaths each year (https://www.cdc.gov/vaccines/vpd/vaccines-diseases.html?CDC_AA_refVal=https%3A%2F%2Fwww.cdc.gov%2Fvaccines%2Fvpd-vac%2Fdefault.htm). However, many diseases such as *Mycobacterium tuberculosis* or malignancies lack effective vaccines due to improper T-cell activation. A key issue in developing effective vaccines for these diseases is the lack of accurately identified major

histocompatibility complex (MHC)-presented epitopes or immunopeptides. These epitopes are presented on the cell surface and enable T-cells to discern healthy cells from infected or malignant cells. While much effort has recently been invested in accurate prediction of these epitopes *in silico* (2), these are mostly limited to viruses as these contain fewer potential protein antigens (3). Moreover, these tools are not yet sufficiently precise to confidently identify epitopes (4, 5). Therefore, experimental immunopeptidomics workflows, such as epitope detection through LC-MS, are still the best way to accurately identify these immunopeptides (6).

While immunopeptidomics workflows have been readily developed and applied (7), acquisition of immunopeptides through LC-MS suffers from some major problems. First, the acquisition of immunopeptide spectra is hampered due to the low abundance of immunopeptides and even more so of neo-epitopes. Infrequently occurring epitopes are still very challenging to identify through LC-MS, despite enrichment efforts and sample preprocessing (8). Second, in contrast to standard proteomics experiments where proteins are usually digested with trypsin before LC-MS, immunopeptides are captured through immune purification with antibodies followed by acidic elution resulting in mostly nontryptic peptides. The nontryptic nature of immunopeptides results in one less positive charge due to the missing arginine or lysine at the peptide's C-terminus, causing many immunopeptides to be singly charged during MS acquisition. These singly charged peptides are much harder to analyze because, during fragmentation of the peptide, the charge resides on one of the fragments, leaving the other uncharged and thus lost (9). Moreover, most contaminants are singly charged as well, making identifications of immunopeptides much harder (10). The nontryptic nature of immunopeptides hampers not only the acquisition but also the identification of immunopeptide spectra. To match each acquired spectrum with the peptide from which the spectrum originated, proteomics database search engines such as SEQUEST (11), X! Tandem (12), Andromeda (13), or PEAKS DB (14) generate *in*

From the ¹VIB-UGent Center for Medical Biotechnology, VIB, Zwijnaarde, Belgium; ²Department of Biomolecular Medicine, Ghent University, Ghent, Belgium; ³Laboratoire de Spectrométrie de Masse BioOrganique (LSMBO), Université de Strasbourg, CNRS, Strasbourg, France

* For correspondence: Lennart Martens, lennart.martens@vib-ugent.be.

silico spectra for all potentially matching peptides. The complete list of peptides that could be in the sample is called the search space. It is important to note that spectra from peptides that are not included in this search space cannot be identified, even though they were acquired. In standard shotgun proteomics experiments, *in silico* tryptic digestion of relevant proteins yields a broad yet representable search space. In immunopeptidomics, however, the search space tends to be two orders of magnitude larger due to (i) seemingly random cleavage from the protein of origin, (ii) the variable length of MHC class I bound peptides, 8 to 11 amino acids, and MHC-II peptides, 6 to 24 amino acids (15) and (iii) the potential occurrence of conformational cis- and trans-spliced immunopeptides, which are nonlinear peptides that originate from the same or different proteins, respectively (16). Additionally, sequence variants and noncanonical protein sequences are often considered as well, even further increasing the search space. Such a search space expansion leads to considerably more ambiguity between candidate peptide-spectrum matches (PSMs) (17), lower PSM scores, drastically elevated false discovery rate (FDR) score thresholds, and ultimately in fewer identified immunopeptides (18). Furthermore, because tryptic peptides have been the longtime standard in proteomics, search engines as well as bioinformatics tools that aid in identifying LC-MS spectra are tailored toward tryptic peptides, making them less accurate or not applicable at all for immunopeptidomics.

The high need for neo- and xeno-epitope discovery led to the development of many bioinformatics tools to improve or validate identifications in immunopeptidomics. On the one hand, motif deconvolution tools have been developed that leverage binding motifs of immunopeptides to validate immunopeptide identifications. On the other hand, full pipelines have been developed to improve immunopeptide identification. For example, MHCquant (19), which is a recent computational workflow designed specifically for neo-epitope identification, and PEAKS DB (14). Even though PEAKS DB is not specifically designed for immunopeptides, it is highly interesting due to its *de novo*-assisted database searches, which tend to work well for large search spaces. Even though these tools can help with immunopeptide identification, they do not use all available information, such as retention time and fragment ion intensity patterns. Previously, it has been proven that integrating retention time predictions in standard proteomics workflows can improve identification rates (20). Similarly, adding peak intensity predictions to postprocessing tools such as Percolator can also improve identification rates drastically (21), which has already been proven to work for immunopeptides as well by efforts such as Prosit (22, 23). Similarly, tools such as DeepLC (24) and MS²PIP (25–27) can provide accurate retention time predictions and peak intensity predictions, respectively, to aid in postprocessing. Indeed, when combined with Percolator, identification rates at a fixed FDR have been proven to substantially increase (21). However, currently, DeepLC and MS²PIP are solely trained on

tryptic peptides. This absence of lysine and arginine at the C-terminus is less of a problem for DeepLC as the effect on retention time is small and is accounted for through feature encoding (28). However, this is not the case for MS²PIP, as alterations in peptide composition as well as fragmentation patterns and labeling methods heavily alter peak intensity patterns (25). Therefore, we here present greatly improved MS²PIP models that include immunopeptides and nontryptic peptides in general. Moreover, we have integrated MS²PIP and DeepLC with Percolator into the free and open-source MS²Rescore software tool, which enables improved rescoring of peptide identifications from various proteomics search engines. Altogether, we show that well-adapted fragmentation spectrum and retention time predictions integrated into MS²Rescore drastically increase immunopeptide identification rates and outperform existing postprocessing methods.

EXPERIMENTAL PROCEDURES

Training and Evaluation of New MS²PIP Spectrum Prediction Models for Immunopeptides

To train and test new MS²PIP models, five publicly available immunopeptide data sets and one publicly available chymotrypsin-digestion data set were downloaded from PRIDE Archive (29, 30). Similarly, for evaluation on representative unseen data, four distinct data sets were downloaded: (i) a data set containing HLA-I immunopeptides, (ii) a data set containing HLA-II immunopeptides, (iii) the data set with tryptic peptides that was previously used to evaluate the existing MS²PIP higher-energy collision-induced dissociation (HCD) models, and (iv) a data set containing chymotrypsin-digested peptide data. The corresponding ProteomeXchange project identifiers as well as the number of unique peptides and HLA patterns for each data set are listed in supplemental Table S1.

As tandem mass spectrometry (MS2) fragmentation patterns are highly dependent on the instrument, instrument settings, fragmentation method, and any applied labeling methods (25), all MS²PIP train, test, and evaluation data must originate from experiments with the same experimental parameters. Unlabeled HCD data from Quadrupole-Orbitrap instruments was used, as this makes the newly trained models widely applicable and plenty of training data is readily available on public repositories. For each PRIDE Archive project, the raw mass spectrometry files were converted to Mascot Generic Format (MGF) files using ThermoRawFileParser (v1.3.4) (31). The corresponding identification files were converted to MS²PIP input files using custom Python scripts and were further filtered to retain unique combinations of peptide sequence, modifications, and precursor charge at 1% FDR. Next, all spectra were combined into one MGF file. The universal spectrum identifier (32) was used as unique identifier for each PSM to ensure reproducibility and a one-on-one mapping between peptide identifications and spectra. Data from each PRIDE Archive project was either used as train/test data or as evaluation data to ensure fully independent data sets, except for the chymotrypsin data, where the same project was used to provide both training/testing data (70%) and evaluation data (30%). This split was made after selecting unique peptide-modification-charge combinations, to ensure no overlap in samples between both splits.

Similarly to the 2019 MS²PIP models, new models were trained with the XGBoost machine learning algorithm (33). The Hyperopt (34) package (v0.2.5) was used in combination with four-fold cross-validation for hyperparameter optimization, allowing 400 boosting rounds

and early stopping fixed at 10 rounds. Hyperparameters were optimized for each training data set separately, as well as for b- and y-ion models. All selected hyperparameters are shown in [supplemental Table S2](#). To evaluate each model, the Pearson correlation coefficient (PCC) was calculated between observed and predicted b- and y-ion peak intensities for each spectrum. The model performances were further analyzed by peptide length and precursor charge. Ultimately, three models were trained: the *immunopeptide* model solely trained on immunopeptides, the *immunopeptide-chymotrypsin* model trained on immunopeptides supplemented with chymotrypsin-digested peptides, and the *nontryptic immunopeptide* model solely trained on nontryptic immunopeptides. Ultimately, two models were integrated into MS²PIP: (i) the *immunopeptide* model and (ii) the *immuno-chymotrypsin* model. The former can be used for immunopeptide peak intensity predictions and the latter for tryptic and more general nontryptic peptide predictions. For further analysis into rescoring immunopeptide PSMs, only the *immunopeptide* model was used, as it showed the best performance for both HLA-I and HLA-II immunopeptides.

To compare MS²PIP predictions with ProSIT predictions, the same evaluation data sets were used as mentioned above. ProSIT (v1.1.2) was downloaded from GitHub (<https://github.com/kusterlab/prosit>) and the *hcd_hla* and *irt_prediction* models were downloaded from Figshare (<https://figshare.com/projects/prosit/35582>). MS²PIP predictions were acquired for the general proteomics and chymotrypsin evaluation data with the immuno-chymotrypsin model and for the HLA-I and HLA-II evaluation data with the immunopeptide model. Peptides that were not included in the ProSIT output were filtered out of the MS²PIP predictions. The performance was measured in both PCC and spectral angle to ensure a thorough comparison. Only correlations for singly charged fragment ions were taken into account, as the newly trained MS²PIP models only predict intensities for these ions.

Evaluation of MS²Rescore on HLA Class I Peptides and Comparison with ProSIT Rescoring

To validate the capacity of the new MS²PIP models to improve immunopeptide identification rates, the new models were implemented with DeepLC (v0.1.36) and Percolator (v3.5) into MS²Rescore. MS²Rescore calculates various meaningful features based on (i) the search engine output, (ii) the DeepLC-predicted and the observed retention times, and (iii) the MS²PIP-predicted and the observed MS² peak intensities. These features are then passed to Percolator for PSM rescoring. Search engine features were selected based on the previous publication by Granholm *et al.* (35) and replicated for use with MaxQuant search results (36). MS²PIP features were used as first described by Silva *et al.* (21). All features generated by MS²Rescore are listed in [supplemental Table S3](#).

MS²Rescore was validated on a large-scale HLA class I data set (37), which was also used to validate the recently published ProSIT-rescoring effort for immunopeptides (PXD021398) (23). This allows both an evaluation of the improved identification rates due to the new MS²PIP models and a straight-forward comparison with ProSIT rescoring. First, the *msms.txt* identification files for the projects' two MaxQuant searches (alkylated and nonalkylated samples), the ProSIT-rescored Percolator output files, and the raw mass spectrometry files were downloaded from PRIDE Archive. The mass spectrometry files were then further processed with ThermoRawFileParser (v1.3.4) (31) and the PSMs for each of the two MaxQuant searches were rescored separately. Two rescoring methods were evaluated: (i) using only search engine features, replicating a normal Percolator run, and (ii) using the full MS²Rescore feature set, including search engine-, MS²PIP-, and DeepLC-features. Additionally, these rescoring methods were compared with the original MaxQuant results and with the downloaded ProSIT-rescoring results.

Each rescoring method was evaluated at varying FDR thresholds in terms of identification rate and number of unique identified peptides. The contribution of the different feature sets in MS²Rescore was visualized using Percolator's model weights, and the distributions of retention time difference and MS²PIP prediction correlations were compared between decoy PSMs, accepted target PSMs, and rejected target PSMs.

Additionally, as reported by Wilhelm *et al.* (23), sequence motif patterns for HLA pattern C*12:03 were further analyzed with GibbsCluster (v2.0) (38), for the gained and lost peptides compared to rescoring with only search engine features.

Evaluation of MS²Rescore Across Collision Energy Settings and Peptide Abundances

To further analyze MS²Rescore performance for various experimental collision energy settings, replicate LC-MS/MS runs were performed on HL60 cells at collision energy values of 25, 27, 30, 32, and 35 NCE ([supplemental Methods](#)). The resulting spectra were searched with the Andromeda search engine (MaxQuant v1.6.14.0) against the human UniProtKB-SwissProt (14-09-2020; 20, 388 sequences, Taxonomy ID 9606) database without any enzyme specificity. A minimal peptide length of seven amino acids was required. Oxidation (M) was set as variable modification with a maximum of three modifications per peptide. Mass tolerances were set at 5 ppm and 20 ppm for MS1 and MS2 spectra, respectively. FDR was kept at 100% with the use of a decoy strategy for downstream rescoring with (i) only search engine features and (ii) the full MS²Rescore feature set. Furthermore, for all LC-MS/MS runs at all collision energy settings, precursor intensities were obtained from the MaxQuant *msms.txt* file to assess any differences in the performance of MS²Rescore between low and high abundant peptides.

Evaluation of MS²Rescore on HLA Class II Peptides

To validate MS²Rescore for HLA class II peptides, another set of raw mass spectrometry files were downloaded from PRIDE Archive (PXD015408). As the uploaded search engine results were already filtered at 5% FDR, the spectra were reanalyzed with PEAKS DB (v10.5) (14) with the same search parameters that were used in the original publication, that is, no enzyme specificity, precursor error tolerance of 10 ppm, fragment ion tolerance of 0.01 Da, with oxidation (M), deamidation (NQ), and trioxidation (C) as variable modifications, searched against the UniProtKB-SwissProt database (01-2021, 22, 235 sequences, Taxonomy ID 10090). The *mzIdentML* identification file as well as the corresponding MGF files were exported from PEAKS DB and were rescored with MS²Rescore with (i) only search engine features and (ii) the full MS²Rescore feature set, as described above for the evaluation on HLA class I peptides.

General Data Processing and Data Visualization

All plots, unless specified, were generated in Jupyter notebooks (v6.4.0) using Python (v3.8.3) with the Matplotlib (v3.4.2) (39), Seaborn (v0.11.0) (40), UpSetPlot (v0.6.0), and spectrum_utils (v0.3.5) (41) libraries.

RESULTS

Newly Trained MS²PIP Models Accurately Predict Immunopeptide Spectrum Peak Intensities

In order to improve the identification rate of immunopeptides by leveraging peak intensity predictions, new models for MS²PIP were trained specifically for immunopeptides. Despite using different training set compositions, all newly trained

models drastically improve predictions for both HLA-I and HLA-II data in comparison with the tryptic 2019 HCD model (Fig. 1A). Surprisingly, even for standard tryptic shotgun proteomics data, the predictions from the new models are slightly better, largely due to the portion of tryptic peptides within the immunopeptide training data. Indeed, when these peptides are left out of the training data, accuracy drops in comparison with the 2019 HCD model. While both immunopeptide models are well suited to predict peak intensities for tryptic and immunopeptides, the performance on chymotrypsin-digested peptides is not as high (supplemental Fig. S1). Thus, even though immuno- and chymotrypsin-digested peptides are both considered nontryptic, they are still very different for MS²PIP peak intensity predictions. Overall, immunopeptide peak intensity predictions are drastically improved by all the newly trained models, with the immunopeptide model showing the highest accuracy (median PCC of 0.94). The exact median PCCs are listed in supplemental Table S4. Examples of a prediction with median PCC values with the immunopeptide model and the corresponding, less accurate, 2019 model prediction are shown in Figure 1, B and C.

MS²Rescore Drastically Improves Immunopeptide Identification Rates

Ultimately, the goal of these newly trained models is to improve immunopeptide identification rates by providing more accurate peak intensity predictions. Therefore, (i) identification results without rescoring, (ii) rescoring with solely search

engine features (replicating a normal Percolator run), (iii) rescoring with MS²Rescore (including DeepLC and the new MS²PIP models), and (iv) rescoring with the recently published Prosit models were compared in terms of the total amount of identifications as well as the number of unique identifications based on sequence. Overall, rescoring with both MS²Rescore and Prosit substantially improved the spectrum identification rate in comparison with rescoring with search engine features alone or not rescoring and this at both 1% and 0.1% FDR. Indeed, MS²Rescore achieves an identification rate of 11.1%, out of 18 million spectra, compared to 7.6% for traditional rescoring (an increase of 46%), and only 1.9% for the Max-Quant search results, all at 1% FDR (Fig. 2, A and B). Moreover, 83% of the identified spectra at 1% FDR are retained when restricting the threshold to 0.1% FDR. Thus, providing peak intensities and retention time predictions to Percolator substantially increases the number of identified immunopeptides. This is clearly illustrated by analyzing the Percolator weights for each separate feature as well as the combined absolute weights for search engine features, MS²PIP features, and DeepLC features (supplemental Fig. S2). Similarly, the number of unique identified immunopeptides increases by 36% when adding MS²PIP and DeepLC Features for the 1% FDR and even more so for 0.1% FDR where the number of unique identified peptides reaches nearly 300% of the number of traditional Percolator identification results (Fig. 2, C–F). These gains are consistent across all 95 HLA class I alleles included in the data (supplemental Figs. S3–S4), showing that




FIG. 1. Performance evaluation of the new nontryptic MS²PIP models. Boxplots comparing Pearson correlation distributions between predicted and observed peak intensity values on spectrum level for the 2019 HCD model, immunopeptide model, nontryptic immunopeptide model, and immuno-chymotrypsin model. Performance is evaluated on general tryptic proteomics data, HLA-I immunopeptide data, HLA-II immunopeptide data, and chymotrypsin data (A). Predictions for immunopeptide "KQHGVNVSV" by new immunopeptide MS²PIP model (top) compared to observed MS2 spectrum (mzspec:PXD005231:20160513_TIL1_R1:scan:10909) (bottom) (B). This peptide was selected for visualization as its PCC lies close to the median for the immunopeptide model on all HLA-I evaluation data. Predictions by the 2019 HCD model for the same immunopeptide (C) as visualized in B. HCD, higher-energy collision-induced dissociation; MS2, tandem mass spectrometry; PCC, Pearson correlation coefficient.




FIG. 2. Percentage of identified spectra and unique identified peptides using different rescoring methods (PXD021398). Bar charts showing the spectrum identification rate out of 18,375,659 spectra (A and B), showing the total number of unique identified peptides in terms of sequence (C and D), and showing the shared (blue), gained (green), and lost (red) number of unique (by sequence) identified immunopeptides in relation to rescoring with only search engine features (E and F). All results are shown for the 1% FDR (A, C, and E) and 0.1% FDR (B, D, and F) thresholds. FDR, false discovery rate.

the newly trained MS²PIP model, and therefore MS²Rescore, is generalizable across different HLA types. In fact, MS²Rescore allows for a substantial increase in identification rate for HLA types with initially fewer identifications (e.g., A0101, A0204, B4402...), indicating that MS²Rescore especially improves the peptide identification coverage for harder-to-identify HLA alleles.

The power of providing these predictions to Percolator is further illustrated when visualizing the distributions for decoy PSMs, rejected target PSMs, and accepted target PSMs. Indeed, the distributions for decoy and rejected target PSMs are highly similar for both the retention time error as well as the PCC, while the accepted target PSMs accumulate around low retention time errors and high PCCs (Fig. 3, A and B). The accepted target PSMs are clearly separable from the decoy and rejected target PSMs using only the PCC and retention time error distributions (Fig. 3, C and D). Furthermore, while both metrics correlate with the search engine score, a large amount of decoy and rejected target PSMs can only be separated from the target PSMs by also including PCC or retention time error information (Fig. 3, E and F). This clearly illustrates how Percolator achieves its much-improved separation between true and false target PSMs when provided with peak intensity and retention time prediction features. Furthermore, PSMs that previously would have been incorrectly accepted below a 1% FDR because of a high search engine score alone are now rejected due to a low PCC, a high retention time error, or both. This most likely accounts for the small percentage of identified peptides that are lost after rescoring.

MS²Rescore Outperforms the Current State-of-the-Art

The integration of MS²PIP-, DeepLC-, and search engine-based features in MS²Rescore has proven to substantially

increase the identification rate of immunopeptides and furthermore, outperforms the recently published Prosit-rescoring method (23). In comparison with Prosit, MS²Rescore gains 5% and 35% more identifications at 1% FDR and 0.1% FDR, respectively. This trend continues for the number of unique identified peptides with a respective increase of 8% and 57% (Fig. 2). Indeed, over 32,000 unique peptides were identified below the 0.1% FDR threshold by MS²Rescore, while Prosit rescoring only identified these peptides at the 1% FDR threshold, highlighting the gain in confidence in terms of unique identified immunopeptides (supplemental Fig. S5). MS²Rescore thus substantially increases the identification rate, especially for more stringent FDR thresholds.

Peptides for which Prosit cannot predict MS2 peak intensities, that is, unmodified (C), cysteinylated (C), and acetylation (N-terminus), were left out of the Prosit-rescoring output. That is why MS²Rescore includes more PSMs in the unfiltered data set at 100% FDR, especially for the noIAA sample (supplemental Fig. S6, A and B). For a second, thorough comparison, these PSMs were left out of the MS²Rescore output as well. However, this seemed to have a rather negligible impact on the number of identifications at 1% and 0.1% FDR (supplemental Fig. S6, C and D), and thus the difference in identification rates cannot be attributed to these filtered peptides.

Furthermore, a comparison of the peptide spectrum prediction accuracy of the newly trained MS²PIP models and Prosit indicates that higher performance of MS²Rescore cannot be attributed to improved peak intensity predictions. Indeed, depending on the correlation metric that is used, either MS²PIP or Prosit performs slightly better than the other on the evaluation data (supplemental Fig. S7). It is important to note, however, that for Prosit, the correct collision energy (CE) value should be selected for optimal performance. To




FIG. 3. Distributions of MS²PIP-, DeepLC, and search engine-based features. Density plots showing the distribution of the smallest retention time error between observed and predicted retention time (A) and showing the Pearson correlation between observed and predicted peak intensities (B) for each PSM split in decoys (red), rejected targets, q value >0.01 (blue), and accepted targets, q value <0.01 (green) (note the rejected targets distribution coincides with the decoy distribution). Scatterplots showing the relation between observed and predicted retention time (C), between Pearson correlation and retention time error (D), between retention time error and search engine score (E), and between Pearson correlation and search engine score (F) for decoys (red), rejected targets (blue), and accepted targets (green) for all PSMs from “GN20170531_SK_HLA_C0102_R1_01”. Note that all retention time errors are the smallest retention time error for each precursor (see *rt_diff_best* supplemental Table S3), all Pearson correlation coefficients are calculated for each full log2 transformed spectrum, and all zero values are caused by the fact that either the observed or predicted intensities for a given ion type are all zero. PSMs, peptide-spectrum matches.

determine if the difference in rescoring between MS²Rescore and Prosit is driven by a difference in features, all MS²Rescore features that do not have a close counterpart in Prosit rescoring were removed for a separate run. This includes all DeepLC retention time features and various search engine features (supplemental Table S3). As MS²Rescore and Prosit both include various similar peak intensity prediction features, these were retained. With this reduced feature set, the performance of MS²Rescore is slightly lower than for Prosit rescoring (supplemental Fig. S6, A–D), which confirms that the retention time features and additional search engine features result in improved MS²Rescore performance over Prosit rescoring.

Similarly to Prosit, the identified sequence motif for HLA type C*12:03 was highly similar to the motif reported in the original publication of the data set (23, 37), while the peptides

that were removed by MS²Rescore compared to search engine rescoring showed quite different, less conserved sequence motifs (supplemental Fig. S8).

MS²Rescore Generalizes Well Across Collision Energy Settings and Peptide Abundances

Because MS²PIP does not account for CE in its predictions, MS²PIP and consequently MS²Rescore could potentially be biased toward spectra obtained with certain CE values. Therefore, search results of replicate mass spectrometry runs with CE settings varying from 25 to 35 were postprocessed with MS²Rescore. For each CE value, MS²Rescore shows a significant increase in identification rate. However, for larger (less optimal) CE values, the overall identification rate decreases (supplemental Fig. S9A). This is most likely due to a reduced quality in fragmentation spectra, which is reflected in

the decreased explained ion current and in line with the b- and y-ion MS²PIP PCC distributions for accepted PSMs when using suboptimal CE values (supplemental Fig. S9, D–F). Most interestingly, the relative gain in unique identified peptides for MS²Rescore increases for higher (and therefore less optimal) CE values, approaching a 60% increase for CE 35, by slightly shifting the feature weights away from fragmentation features in favor of DeepLC retention time features (supplemental Fig. S9, B and C). Consequently, MS²Rescore is able to recover peptides that would otherwise be lost due to lower-quality fragmentation spectra.

A similar effect is observed for low abundant peptides. Indeed, the largest relative gain achieved by MS²Rescore in terms of number of unique identified peptides is seen for the lowest precursor intensities (supplemental Fig. S10, A and B), where traditional rescoring fails to recover most identifications. MS²Rescore is thus not only able to increase the amount of identifications for immunopeptides in general, it can recover peptides previously lost due to low precursor intensities and thus lower-quality spectra or nonoptimal instrument settings.

MS²Rescore is Unbiased to Different HLA Classes

To evaluate the performance of MS²Rescore on HLA class II peptides, another publicly available data set was reanalyzed. However, while for the HLA class I data set human immunopeptides and MaxQuant search engine (13) results were used, for this HLA class II data set, mouse data was searched with PEAKS DB (14). As was the case for HLA class I peptides, MS²Rescore significantly increases the identification rate for HLA class II peptides with 15% and 57% for the 1% and 0.1% FDR threshold, respectively (supplemental Fig. S11). These increases are, however, slightly lower than for the HLA class I data set. Moreover, where previously conventional rescoring showed a significant increase in comparison to search engine rescoring, here, the gain in comparison with no rescoring is lower for both identification rate as well as number of unique identified peptides. This is likely due to (i) the less extensive search engine features that are calculated for the PEAKS DB pipeline in MS²Rescore (supplemental Table S3) and (ii) to the fact PEAKS DB is likely better equipped to identify immunopeptides than MaxQuant due to its *de novo*-assisted database search (14). Nevertheless, the full MS²Rescore feature set, including peak intensity and retention time predictions, still results in a significantly higher identification rate. Altogether, these results show that MS²Rescore generalizes well across HLA class I and class II immunopeptides, across different species, and can boost the performance from different search engines.

DISCUSSION

By training new peak intensity prediction models, we were able to greatly enhance immunopeptide identification rate

through PSM rescoring. While all newly trained MS²PIP models greatly enhance peak intensity predictions for immunopeptides, the model trained solely on immunopeptides performed best. Even though the immuno-chymotrypsin model contained the same immunopeptide train set, the addition of the chymotrypsin-digested peptides did lower the performance slightly. Similarly, not including chymotrypsin-digested peptides in the training data resulted in lower accuracies for the chymotrypsin-digested peptides. Indeed, immunopeptides are generally much smaller and consequently carry a lower charge state; as a consequence, these immunopeptide-specific MS²PIP models are not able to predict the behavior of longer and higher charged peptides in the mass spectrometer. While both immuno- and chymotrypsin-digested peptides are considered nontryptic, their properties can be very different, leading to reduced accuracy of peak intensity of MS²PIP when applied on a different type of nontryptic peptides. Surprisingly, while immuno- and chymotrypsin-digested peptides are antagonistic, immunopeptides and tryptic peptides seem synergistic in terms of training data. This comes as no surprise as almost 50% of the immunopeptide training data consists of tryptic peptides. Immunopeptides are thus not necessarily nontryptic. However, the actual occurrence of tryptic peptides in immunopeptidomics samples is most likely much lower. This unrepresentatively sized tryptic portion most likely originated from the tryptic bias in current immunopeptidomics workflows. Indeed, in previous studies, tryptic MHC peptide coverage could rise to 70% (42). By training new, nontryptic models of MS²PIP, we take a first step in decreasing this tryptic bias to ultimately be able to analyze an unbiased immunopeptide landscape.

Moreover, by integrating the new immunopeptide model with retention time predictions and search engine features into MS²Rescore, we greatly enhanced the ability of Percolator to rescore immunopeptide PSMs, resulting in a much-improved immunopeptide identification workflow. Furthermore, rescoring drastically increases the number of unique identified peptides, which is of crucial importance for the discovery of potential neo-epitopes for cancer vaccination or xenopeptides for anti-bacterial and to a lesser extent, anti-viral vaccines. Moreover, while previously almost no identifications were found at a more confident 0.1% FDR threshold, MS²Rescore allows a lowering of the FDR threshold to 0.1%, while retaining 83% of the peptides identified at 1% FDR. This illustrates the large increase in confidence of the identified PSMs MS²Rescore introduces. Besides the increase in both PSM confidence and identification rate, MS²Rescore has shown to be unbiased with regard to HLA patterns and CE settings. Most importantly, the relative identification gain introduced by MS²Rescore is even larger for HLA patterns that initially had fewer identifications, showing that MS²Rescore is able to increase the view on the immunopeptide landscape for traditionally harder-to-identify HLA patterns. Moreover, MS²Rescore is able to recover peptide identifications that

would have been lost due to lower-quality spectra by making use of DeepLC retention time features and can therefore recover substantial additional identifications for low abundant peptides. This potentially enables the recovery of biologically relevant neo- or xeno-epitopes that occur less frequently in the sample. Furthermore, MS²Rescore is able to gain immunopeptide identifications regardless of the search engine used, for both HLA class I and class II peptides, and across different species.

Additionally, MS²Rescore with DeepLC and the new immunopeptide MS²PIP models shows an improved identification rate over the recently published Prosit effort, especially for lower FDR thresholds. As Prosit has shown to provide more accurate predictions compared to previous MS²PIP models, it is unlikely that MS²Rescore's higher performance can be attributed to superior peak intensity predictions. Indeed, the peak intensity prediction accuracies of the new MS²PIP models and of Prosit are highly similar for immunopeptides (supplemental Fig. S7) even when Prosit has been optimized for the right CE. These negligible differences in peak intensity prediction correlations are therefore likely not the reason for the higher performance of MS²Rescore in favor of Prosit. Instead, it is more likely that the main difference in rescoring performance is the result of the generation of more relevant MS²PIP-, DeepLC, and search engine-derived features. Indeed, when the majority of the search engine features and all DeepLC retention time features were omitted, reflecting the more limited Prosit feature set, the performance of MS²Rescore drops as well (supplemental Fig. S6). By providing a more extensive feature set, MS²Rescore creates a unique feature space that allows Percolator to separate true from false identifications much better than when provided with limited features without retention time or peak intensity information (Fig. 3). The combination of all these calculated features is therefore likely to be driver of MS²Rescore's superior performance.

MS²Rescore is freely available under the permissive Apache 2.0 open-source license on GitHub (<https://github.com/compomics/ms2rescore>) and can easily be installed locally through the cross-platform PyPI Python package as well as with a standalone windows install script. Both a command line interface and a graphical user interface are available, various identification files from different search engines are accepted, and both MS²PIP and DeepLC can handle a variety of modifications, eliminating the need to filter identification files before rescoring. Altogether, these new models show great promise to greatly extend the immunopeptide landscape in existing and future immunopectidomics experiments.

DATA AVAILABILITY

MS²Rescore is available at <https://github.com/compomics/ms2rescore>. All additional code used in this work is available at [compomics/ms2rescore-immunopectidomics-manuscript](https://github.com/compomics/ms2rescore-immunopectidomics-manuscript) (github.com). The data used for training and evaluation of

the newly trained models, the models themselves as well as the MS²Rescore output is available on Zenodo at <https://doi.org/10.5281/zenodo.6532013>. Additional mass spectrometry proteomics data have been deposited to the ProteomeXchange Consortium via the PRIDE (30) partner repository with the data set identifier PXD033868.

Supplemental data—This article contains supplemental data (6, 37, 43–50).

Acknowledgments—NanoLC-MS/MS instruments were supported by the French Proteomic Infrastructure (ProFI FR2048; ANR-10-INBS-08-03).

Funding and additional information—R. B. acknowledges funding from the Vlaams Agentschap Innoveren en Ondernemen under project number HBC.2020.2205.; R. G. received funding from the Research Foundation Flanders (FWO) [1S50918N]. S. D. and L. M. acknowledge funding from the European Union's Horizon 2020 Programme (H2020-INFRAIA-2018-1) [823839]; L. M. acknowledges funding from the Research Foundation Flanders (FWO) [G028821N] and from Ghent University Concerted Research Action [BOF21/GOA/033]. A. D. received funding from the Research Foundation Flanders (FWO) [1SE3722].

Author contributions—A. D., A. H., and R. G. methodology; A. D. software; A. D., R. B., and R. G. validation; A. D. and R. G. formal analysis; A. D. and R. G. writing—original draft; A. D., R. B., A. H., C. C., S. D., L. M., and R. G. writing—review and editing; A. D., R. B., C. C., S. D., L. M., and R. G. funding acquisition; A. D., R. B., S. D., L. M., and R. G. conceptualization; A. D., A. H., C. C., and R. G. investigation; C. C. and L. M. resources; L. M. and R. G. project administration; R. G. supervision.

Conflict of interest—The authors declare that they have no conflict of interest with the contents of the article.

Abbreviations—The abbreviations used are: CE, collision energy; FDR, false discovery rate; HCD, higher-energy collision-induced dissociation; MS2, tandem mass spectrometry; MGF, mascot generic format; PCC, Pearson correlation coefficient; PSM, peptide-spectrum match..

Received December 8, 2021, and in revised form, June 30, 2022
Published, MCPRO Papers in Press, July 6, 2022, <https://doi.org/10.1016/j.mcpro.2022.100266>

REFERENCES

1. Sattler, S. (2017). In: *Advances in Experimental Medicine and Biology*, Springer New York LLC, New York: 3–14
2. Raoofi, E., Hemmati, M., Eftekhari, S., Khaksaran, K., Mahmodi, Z., Farajollahi, M. M., et al. (2020) Epitope prediction by novel immunoinformatics approach: a state-of-the-art Review. *Int. J. Pept. Res. Ther.* **26**, 1155–1163
3. Mayer, R. L., and Impens, F. (2021) Immunopectidomics for next-generation bacterial vaccine development. *Trends Microbiol.* **29**, 1034–1045

4. Larsen, M. V., Lundegaard, C., Lamberth, K., Buus, S., Lund, O., and Nielsen, M. (2007) Large-scale validation of methods for cytotoxic T-lymphocyte epitope prediction. *BMC Bioinformatics* **8**, 1–12
5. Zhang, H., Lundegaard, C., and Nielsen, M. (2009) Pan-specific MHC class I predictors: a benchmark of HLA class I pan-specific prediction methods. *Bioinformatics* **25**, 83–89
6. Bassani-Sternberg, M., Pletscher-Frankild, S., Jensen, L. J., and Mann, M. (2015) Mass spectrometry of human leukocyte antigen class I peptidomes reveals strong effects of protein abundance and turnover on antigen presentation. *Mol. Cell Proteomics* **14**, 658–673
7. Solleder, M., Guillaume, P., Racle, J., Michaux, J., Pak, H. S., Müller, M., et al. (2020) Mass spectrometry based immunopeptidomics leads to robust predictions of phosphorylated HLA class I ligands. *Mol. Cell Proteomics* **19**, 390–404
8. Faridi, P., Purcell, A. W., and Croft, N. P. (2018) In Immunopeptidomics we need a sniper instead of a shotgun. *Proteomics* **18**, e1700464
9. Pfammatter, S., Bonnell, E., Lanoix, J., Vincent, K., Hardy, M.-P. P., Courcelles, M., et al. (2020) Extending the comprehensiveness of immunopeptidome analyses using isobaric peptide labeling. *Anal. Chem.* **92**, 9194–9204
10. Purcell, A. W., Ramarathnam, S. H., and Ternet, N. (2019) Mass spectrometry-based identification of MHC-bound peptides for immunopeptidomics. *Nat. Protoc.* **14**, 1687–1707
11. Eng, J. K., McCormack, A. L., and Yates, J. R. (1994) An approach to correlate tandem mass spectral data of peptides with amino acid sequences in a protein database. *J. Am. Soc. Mass Spectrom.* **5**, 976–989
12. Craig, R., and Beavis, R. C. (2004) TANDEM: matching proteins with tandem mass spectra. *Bioinformatics* **20**, 1466–1467
13. Cox, J., Neuhauser, N., Michalski, A., Scheltema, R. A., Olsen, J. V., and Mann, M. (2011) Andromeda: a peptide search engine integrated into the MaxQuant environment. *J. Proteome Res.* **10**, 1794–1805
14. Zhang, J., Xin, L., Shan, B., Chen, W., Xie, M., Yuen, D., et al. (2012) Peaks DB: de novo sequencing assisted database search for sensitive and accurate peptide identification. *Mol. Cell Proteomics* **11**, M111.010587
15. Jiang, J., Natarajan, K., and Margulies, D. H. (2019). In: *Advances in Experimental Medicine and Biology*, Springer New York LLC, New York: 21–62
16. Faridi, P., Li, C., Ramarathnam, S. H., Vivian, J. P., Illing, P. T., Mifsud, N. A., et al. (2018) A subset of HLA-I peptides are not genomically templated: evidence for cis- and trans-spliced peptide ligands. *Sci. Immunol.* **3**, eaar3947
17. Colaert, N., Degroove, S., Helsens, K., and Martens, L. (2011) Analysis of the resolution limitations of peptide identification algorithms. *J. Proteome Res.* **10**, 5555–5561
18. Verheggen, K., Røder, H., Berven, F. S., Martens, L., Barsnes, H., and Vaudel, M. (2020) Anatomy and evolution of database search engines—a central component of mass spectrometry based proteomic workflows. *Mass Spectrom. Rev.* **39**, 292–306
19. Bichmann, L., Nelde, A., Ghosh, M., Heumos, L., Mohr, C., Peltzer, A., et al. (2019) MHCquant: automated and reproducible data analysis for immunopeptidomics. *J. Proteome Res.* **18**, 3876–3884
20. Dorfer, V., Maltsev, S., Winkler, S., and Mechtler, K. (2018) CharmerT: boosting peptide identifications by chimeric spectra identification and retention time prediction. *J. Proteome Res.* **17**, 2581–2589
21. Silva, A. S. C., Bouwmeester, R., Martens, L., and Degroove, S. (2019) Accurate peptide fragmentation predictions allow data driven approaches to replace and improve upon proteomics search engine scoring functions. *Bioinformatics* **35**, 1401–1403
22. Li, K., Jain, A., Malovannaya, A., Wen, B., and Zhang, B. (2020) DeepRescore: leveraging deep learning to improve peptide identification in immunopeptidomics. *Proteomics* **20**, e1900334
23. Wilhelm, M., Zolg, D. P., Graber, M., Gessulat, S., Schmidt, T., Schnatbaum, K., et al. (2021) Deep learning boosts sensitivity of mass spectrometry-based immunopeptidomics. *Nat. Commun.* **12**, 3346
24. Bouwmeester, R., Gabriels, R., Hulstaert, N., Martens, L., and Degroove, S. (2021) DeepLC can predict retention times for peptides that carry as-yet unseen modifications. *Nat. Methods*, 1–7
25. Gabriels, R., Martens, L., and Degroove, S. (2019) Updated MS²PIP web server delivers fast and accurate MS² peak intensity prediction for multiple fragmentation methods, instruments and labeling techniques. *Nucl. Acids Res.* **47**, W295–W299
26. Degroove, S., Maddelein, D., and Martens, L. (2015) MS2PIP prediction server: compute and visualize MS2 peak intensity predictions for CID and HCD fragmentation. *Nucl. Acids Res.* **43**, W326–W330
27. Degroove, S., and Martens, L. (2013) MS2PIP: a tool for MS/MS peak intensity prediction. *Bioinformatics* **29**, 3199–3203
28. Ruiz Cuevas, M. V., Hardy, M.-P., Holly, J., Bonnell, E., Durette, C., Courcelles, M., et al. (2020) Most non-canonical proteins uniquely populate the proteome or immunopeptidome. *Cell Rep.* **34**, 108815
29. Martens, L., Hermjakob, H., Jones, P., Adamsk, M., Taylor, C., States, D., et al. (2005) PRIDE: the proteomics identifications database. *Proteomics* **5**, 3537–3545
30. Perez-Riverol, Y., Csordas, A., Bai, J., Bernal-Llinares, M., Hewapathirana, S., Kundu, D. J., et al. (2019) The PRIDE database and related tools and resources in 2019: improving support for quantification data. *Nucl. Acids Res.* **47**, D442–D450
31. Hulstaert, N., Shofstahl, J., Sachsenberg, T., Walzer, M., Barsnes, H., Martens, L., et al. (2020) ThermoRawFileParser: modular, scalable, and cross-platform RAW file conversion. *J. Proteome Res.* **19**, 537–542
32. Deutsch, E. W., Perez-Riverol, Y., Carver, J., Kawano, S., Mendoza, L., Van Den Bossche, T., et al. (2021) Universal spectrum identifier for mass spectra. *Nat. Methods* **18**, 768–770
33. Chen, T., and Guestrin, C. (2016). In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, New York, NY
34. Bergstra, J., Yamins, D., and Cox, D. D. (2013) 30th international conference on machine learning. *ICML* **2013**, 115–123
35. Granholm, V., Kim, S., Navarro, J. C. F., Sjölund, E., Smith, R. D., and Käll, L. (2014) Fast and accurate database searches with MS-GF+percolator. *J. Proteome Res.* **13**, 890–897
36. Tyanova, S., Temu, T., and Cox, J. (2016) The MaxQuant computational platform for mass spectrometry-based shotgun proteomics. *Nat. Protoc.* **11**, 2301–2319
37. Sarkizova, S., Klaeger, S., Le, P. M., Li, L. W., Oliveira, G., Keshishian, H., et al. (2020) A large peptidome dataset improves HLA class I epitope prediction across most of the human population. *Nat. Biotechnol.* **38**, 199–209
38. Andreatta, M., Alvarez, B., and Nielsen, M. (2017) GibbsCluster: unsupervised clustering and alignment of peptide sequences. *Nucl. Acids Res.* **45**, W458–W463
39. Hunter, J. D. (2007) Matplotlib: a 2D graphics environment. *Comput. Sci. Eng.* **9**, 90–95
40. Waskom, M. (2021) seaborn: statistical data visualization. *J. Open Source Softw.* **6**, 3021
41. Bittremieux, W. (2020) Spectrum-utils: a Python package for mass spectrometry data processing and visualization. *Anal. Chem.* **92**, 659–661
42. Chen, R., Fauteux, F., Foote, S., Stupak, J., Tremblay, T. L., Gumani, K., et al. (2018) Chemical derivatization strategy for extending the identification of MHC class I immunopeptides. *Anal. Chem.* **90**, 11409–11416
43. Racle, J., Michaux, J., Rockinger, G. A., Arnaud, M., Bobisse, S., Chong, C., et al. (2019) Robust prediction of HLA class II epitopes by deep motif deconvolution of immunopeptidomes. *Nat. Biotechnol.* **37**, 1283–1286
44. Chong, C., Marino, F., Pak, H., Racle, J., Daniel, R. T., Müller, M., et al. (2018) High-throughput and sensitive immunopeptidomics platform reveals profound interferon γ -mediated remodeling of the human leukocyte antigen (HLA) ligandome. *Mol. Cell. Proteomics* **17**, 533–548
45. Gfeller, D., Guillaume, P., Michaux, J., Pak, H.-S., Daniel, R. T., Racle, J., et al. (2018) The length distribution and multiple specificity of naturally presented HLA-I ligands. *J. Immunol.* **201**, 3705–3716
46. Bassani-Sternberg, M., Bräunlein, E., Klar, R., Engleitner, T., Sinitcyn, P., Audehm, S., et al. (2016) Direct identification of clinically relevant neoepitopes presented on native human melanoma tissue by mass spectrometry. *Nat. Commun.* **7**, 1–16
47. Wang, D., Eraslan, B., Wieland, T., Hallström, B., Hopf, T., Zolg, D. P., et al. (2019) A deep proteome and transcriptome abundance atlas of 29 healthy human tissues. *Mol. Syst. Biol.* **15**, e8503

48. Bassani-Sternberg, M., Chong, C., Guillaume, P., Solleder, M., Pak, H. S., Gannon, P. O., *et al.* (2017) Deciphering HLA-I motifs across HLA peptidomes improves neo-antigen predictions and identifies allosteric regulating HLA specificity. *PLOS Comput. Biol.* **13**, 1–28
49. Marino, F., Semilietof, A., Michaux, J., Pak, H. S., Coukos, G., Müller, M., *et al.* (2020) Biogenesis of HLA ligand presentation in immune cells upon activation reveals changes in peptide length preference. *Front. Immunol.* **11**. doi:1981
50. Gravina, F., Sanchuki, H. S., Rodrigues, T. E., Gerhardt, E. C. M., Pedrosa, F. O., Souza, E. M., *et al.* (2018) Proteome analysis of an *Escherichia coli* ptsN-null strain under different nitrogen regimes. *J. Proteomics* **174**, 28–35