

Computational Tools for the Identification and Interpretation of Sequence Motifs in Immunopeptidomes

Bruno Alvarez, Carolina Barra, Morten Nielsen, and Massimo Andreatta*

Recent advances in proteomics and mass-spectrometry have widely expanded the detectable peptide repertoire presented by major histocompatibility complex (MHC) molecules on the cell surface, collectively known as the immunopeptidome. Finely characterizing the immunopeptidome brings about important basic insights into the mechanisms of antigen presentation, but can also reveal promising targets for vaccine development and cancer immunotherapy. This report describes a number of practical and efficient approaches to analyze immunopeptidomics data, discussing the identification of meaningful sequence motifs in various scenarios and considering current limitations. Guidelines are provided for the filtering of false hits and contaminants, and to address the problem of motif deconvolution in cell lines expressing multiple MHC alleles, both for the MHC class I and class II systems. Finally, it is demonstrated how machine learning can be readily employed by non-expert users to generate accurate prediction models directly from mass-spectrometry eluted ligand data sets.

1. Introduction

The comprehensive set of peptides presented on the cell surface by major histocompatibility complex (MHC) molecules, collectively referred to as the immunopeptidome, represents a unique fingerprint of the health of a cell. T lymphocytes routinely scan this pool of MHC-associated peptides, and can help eliminating infected or cancerous cells that present abnormal peptides on their surface. MHC class I molecules mainly bind peptides derived from intracellular pathogens (such as viruses and some bacteria) and present them to cytotoxic T lymphocytes; MHC class II epitopes are mainly derived from extracellular proteins and are presented to T-helper lymphocytes.

Recent technological advances in the field of mass spectrometry (MS) have brought about a revolution in the study of immunopeptidomes (reviewed by Caron et al.^[1]), with several thousands of peptides that can be detected in a single experiment. Large data sets of naturally presented peptides have

been beneficial to define more accurately the rules of peptide-MHC binding^[2–4] and have also a tremendous potential in defining pathogen-derived T-cell epitopes^[5,6] and neo-epitopes unique to cancerous cells.^[7–10] Part of the appeal of MS-based approaches is that they do not require prior knowledge of MHC motifs, and there is no human intervention in defining a library of candidate sequences to be tested. Therefore, MS provides a large but relatively unbiased sampling of the population of processed and presented peptides available for T-cell recognition.^[3]

In most MS-based pipelines, spectra from eluted peptides are matched against a reference database of natural proteins using algorithms like MaxQuant^[11] or PEAKS,^[12,13] and filtered against a decoy database to limit the false discovery rate

(FDR). Strict FDR filters (typically in the order of 1%) should ensure that most spectra are correctly assigned to bona fide ligands, but often leads to discarding a large portion of the spectra. Several approaches have been proposed to increase the yield of spectral assignment. For example, Mascot Percolator performs machine learning on high-confidence matches to rescore database search results for lower-confidence peptides.^[14] Instead of matching spectra to an entire protein database, SpectMHC constructs reduced, targeted databases of potential MHC ligands, effectively reducing the amount of spurious decoy hits.^[15] Recent work has also suggested that a portion of the unassigned spectra may also be explained by proteasome-generated spliced peptides, which would require the inclusion of spliced variants in the target database.^[16,17]

After spectral assignment to amino acid sequences, peptides must often be aligned and/or clustered to extract meaningful sequence motifs of antigen presentation. The analysis protocols here will generally differ depending on the type of receptor (MHC I vs MHC class II) and type of sample used (cellular vs soluble MHC molecules and mono- vs poly-allelic cell lines). On the one hand, MHC I ligands have a limited range of lengths, typically 8–11 amino acids long, and are characterized by very conserved amino acid preferences at the positions interacting with the MHC binding groove (anchor positions). On the other hand, MHC II ligands are normally longer, with only a portion, the binding core, directly interacting with the MHC groove^[18]; in this case a more sophisticated alignment process is needed to extract conserved binding preferences. In transgenic cells expressing a single MHC molecule (mono-allelic), only one specificity is

B. Alvarez, Dr. C. Barra, Dr. M. Nielsen, Dr. M. Andreatta
Instituto de Investigaciones Biotecnológicas
Universidad Nacional de San Martín
Argentina
E-mail: mandreatta@iibintech.com.ar
Dr. M. Nielsen
Department of Bio and Health Informatics
Technical University of Denmark
Denmark

DOI: 10.1002/pmic.201700252

expected to be present in the data and motif identification is relatively straightforward. Conversely, unmodified cells will naturally present peptides bound to multiple MHC alleles (up to six for HLA class I), with generally different binding preferences; in this case, the multiple specificities contained in the data must be deconvoluted, either by assigning MHC restriction with predictive methods, or by unsupervised clustering.

A popular tool for the unsupervised identification of sequence motifs in immunopeptidomes is GibbsCluster,^[19,20] a web-based and downloadable method that has been included into numerous pipelines for the deconvolution of ligand motifs in the MHC class I^[10,21–23] and MHC class II^[24–26] systems. The GibbsCluster algorithm takes as input a list of peptide sequences (potentially of variable length), and uses a heuristic search to group them into information-rich groups. Besides the sequence motif defining each group, additional properties such as the ligand length distribution of each cluster can be analyzed. A similar method, MixMHCp,^[2,27] has shown performance comparable to GibbsCluster, with the limitation that it can only handle peptides of uniform length. A useful feature of GibbsCluster is the “trash cluster,” a check on internal motif consistency that can filter out outliers that cannot be assigned to any clusters. In the context of MS eluted ligand data, spurious data points can originate both from LC-MS/MS contaminants and from erroneous spectral matches. As a noise filter, GibbsCluster can be beneficial also for mono-allelic data sets where no motif deconvolution is required.

While sequence motifs are generated by GibbsCluster in an unsupervised manner, the method cannot directly assign the MHC restriction of each ligand; this must be done by comparing the unsupervised motifs with published binding motifs of the MHC molecules in the sample.^[28] While this comparative approach is in most cases feasible for human MHC, whose most prevalent alleles have been well characterized and documented, it will fall short for samples containing uncharacterized specificities. Aiming to overcome this limitation, Bassani-Sternberg et al.^[27] suggested a strategy for automatic, unbiased annotation of MHC restriction by comparing motifs detected in multiple data sets with known haplotypes. Exploiting the co-occurrence of MHC alleles across different data sets, they were able to assign motifs to individual alleles without relying on a priori assumptions on their binding specificity, also for alleles without previously documented ligands.

Over the past decades, many efforts have been dedicated to the development of computational methods for the prediction of peptide binding to MHC class I molecules. Most of these T-cell epitope prediction methods have been traditionally trained solely on *in vitro* data of peptide-MHC binding affinity. Although peptide-MHC affinity is arguably the most selective step in antigen presentation, other factors influence the likelihood of a peptide being presented on the cell surface for T-cell recognition.^[29,30] *In vitro* binding affinity data does not address the fact that antigen presentation is a complex, integrative physiological process that combines antigen processing, transport and binding affinity/stability of the peptide-MHC complex. Finally, *in vitro* data fails to reflect any peptide length preference of different MHC-I alleles. Because naturally eluted ligands incorporate information about these additional properties of antigen presentation, large MS-derived sets of peptides can potentially enable

the generation of more accurate prediction models. Recent studies have suggested that models trained on MHC class I ligand data outperform binding affinity-based predictors when it comes to identification of eluted ligands and T-cell epitopes, both in an allele-specific setting^[2,31] as well as with pan-allelic coverage.^[4] Generic tools for machine learning from peptide sequences such as NNAlign^[31,32] can be applied to individual MS data sets to generate custom-made prediction models, which can in turn be employed for further downstream analyses of the immunopeptidome.

The rapidly expanding collection of naturally eluted ligands revealed by MS and the analysis toolkits developed in its wake hold great promise in understanding the structure of the immunopeptidome and the rules of antigen presentation. However, because of the complexities inherent to MS eluted ligand data, it is not a trivial task to analyze and interpret the information these data sets contain. In this report, we seek to address some common issues and describe strategies to analyze MS ligand data and derive sequence motifs in the various scenarios outlined above (MHCI vs MHCII; mono-allelic vs poly-allelic cell lines), with guidelines and examples on publicly available datasets.

2. MHC Class I, Mono-Allelic Cells

In a recent publication, Abelin et al.^[3] described the development of transgenic cells that express a single human MHC class I allele (HLA), and used them to generate a large set of MHC ligands covering 16 HLA class I alleles. There are obvious advantages in using mono-allelic cells to characterize MHC ligands: firstly, no deconvolution/clustering is required to define motifs at the single-allele resolution; secondly, the assignment of individual peptides to their allele does not have to depend on binding predictions or prior knowledge of the motifs. Apart from technical difficulties in the cell generation, a possible drawback is that the relative level of expression of different MHC alleles in a given cell, and the amount of ligands they present, is lost in a mono-allelic setting. The amount of ligands presented by different alleles may also depend on competition between MHC molecules, where the newly available digested peptides from an unfolding antigen fragment would presumably be captured by MHCs with the highest affinity.^[33]

Although most software for MS spectra mapping uses a strict false discovery rate (FDR) threshold, incorrect ligands may still be present among the matches that pass the FDR check. These may consist of common contaminants such as keratin or histone proteins, as well as residual peptides from previous runs of the LC-MS/MS instruments used for sample preparation.^[34,35] GibbsCluster is a useful tool to detect and remove such contaminants and false hits. For each allele in the Abelin data set,^[3] we applied GibbsCluster-2.0 with default preset options for “MHC class I ligands of lengths 8–13,” specifying a single cluster. Between 0.4 and 16% of the peptides (mean 4%) of lengths 8–13 were inconsistent with the motif identified by GibbsCluster-2.0 and were removed by the program as noise. While distinct motifs can be discerned before trash cluster filtering (see three representative alleles in **Figure 1A**), the post-filtering motifs have higher information content and more well-defined anchor residues (**Figure 1B**). Peptides in the “trash cluster” may sometimes hint at the origin of

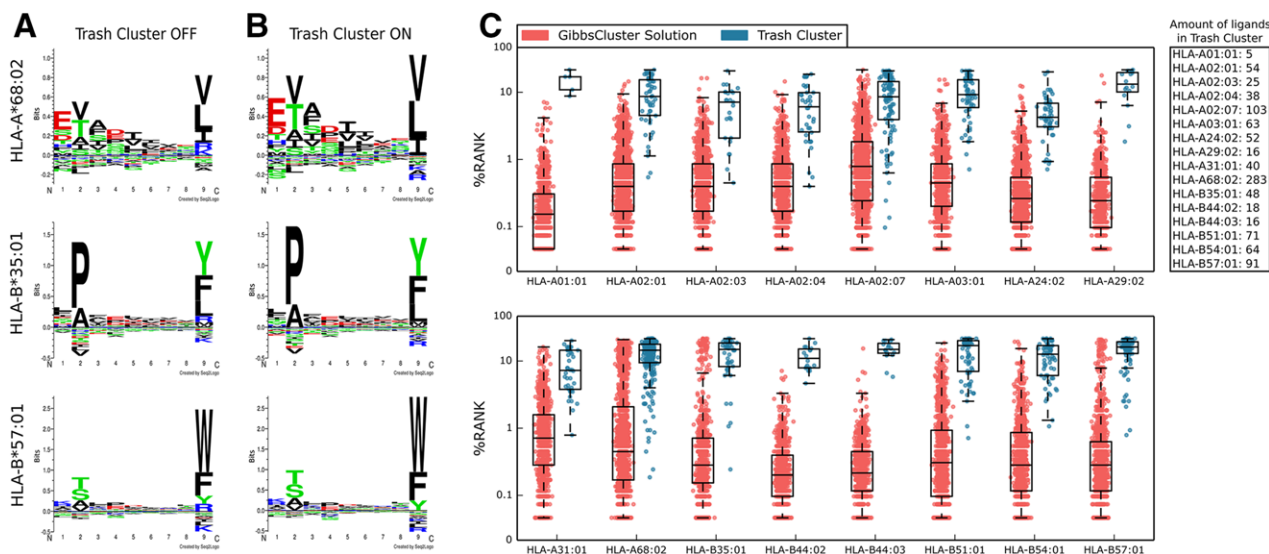


Figure 1. Visualizing motifs and removing contaminants with GibbsCluster. Sequence motifs of three representative alleles (A) before trash cluster filtering and (B) after filtering. The post-filtering motifs have higher information content and lack the putative K/R contamination at P9. C) Distribution of NetMHCpan-3.0 percentile rank scores for peptides in the main cluster (red) and in the trash cluster (blue).

the contamination: for example, the observation of terminal arginine/lysine preferences at the C-terminus in several of the 16 alleles points towards tryptic peptides polluting the mixtures (Supporting Information, Figure 1). The ligands in the Abelin data set have in general very good correspondence to known MHC binding preferences, with an average NetMHCpan-3.0 percentile rank^[36] well below 1% for most alleles (Figure 1C, red boxplots). In contrast, peptides in the trash cluster match very poorly the preferences of their MHC and are assigned high NetMHCpan rank scores (Figure 1C, blue boxplots).

3. MHC Class I, Poly-Allelic Cells

Unmodified antigen-presenting cells will generally express up to six different MHC class I alleles (two each for HLA-A, HLA-B, and HLA-C). The immunopeptidome of these cells therefore consists of multiple specificities mixed together, where the global haplotype is known but the restriction of each individual ligand is unknown. For example, Bassani-Sternberg et al.^[21] described the LC-MS/MS analysis of peptides eluted from seven different cancer cell lines and primary cells, which had been HLA-typed at high resolution, and demonstrated how the GibbsCluster approach could be used to deconvolute the individual peptide restrictions. Here we illustrate the application of GibbsCluster to one of the cell lines from the Bassani-Sternberg study, HCC1143, which expresses the five alleles HLA-A*31:01, HLA-B*35:08, HLA-B*37:01, HLA-C*04:01, HLA-C*06:02.

GibbsCluster finds an optimal solution of four clusters, with a close correspondence to all but one of the HCC1143 alleles (Figure 2), failing to separate HLA-C*04:01 ligands. HLA-C molecules have low expression levels and rather degenerate binding preferences,^[27,37] making the deconvolution of their motifs more challenging. The motifs determined by unsupervised clustering show a remarkable correspondence with the binding preferences predicted by NetMHCpan-3.0.^[36] There are, in some

instances, subtle differences between the NetMHCpan and GibbsCluster motifs, as in the case of additional secondary anchors (e.g., a positively charged P5 for HLA-B*37:01). This suggests that motifs directly derived from eluted ligands may carry an additional level of information on peptide presentation (for instance, secondary anchors conferring improved peptide-MHC complex stability) compared to the NetMHCpan motifs, which were constructed from in vitro binding affinity data. The sizes of the clusters give an indication of the relative level of expression of the different alleles, with the largest group corresponding to the homozygous HLA-A*31:01 (1253 peptides), followed by the two HLA-B alleles (610 and 460 peptides, respectively) and by the lowly expressed HLA-C*06:02 (409 peptides). Finally, 45 peptides were collected by the trash cluster. Interestingly, for six out of seven cell lines in the Bassani-Sternberg data set, we noted a C-terminal enrichment for arginine/lysine in peptide discarded in the trash cluster (Supporting Information, Figure 2). A similar observation was made for the Abelin data set discussed previously, and hints that residual peptides derived from trypsin digestion may often be present in the LC column.

As an alternative approach to unsupervised clustering, one can assign each peptide to a MHC allele using peptide-MHC binding prediction methods; then deriving sequence motifs from the resulting groups of peptides. We applied NetMHCpan^[36] to the peptides eluted from the HCC1143 cell line, assigning peptides to the MHC molecule in the haplotype with the lowest predicted NetMHCpan percentile rank. If a peptide could not be assigned to any MHC molecule with rank $\leq 2\%$, then it was discarded in a trash cluster. While this setup mimics the GibbsCluster strategy, it has the very important difference that NetMHCpan utilizes known motif preferences of the MHC molecules to make the assignments, whereas GibbsCluster is unsupervised and requires no prior knowledge of the motifs. In the case of the HCC1143 cell line, the MHC molecules are all well characterized and the solutions found by the two approaches are remarkably similar (Supporting Information, Figure 3). Assignment by NetMHCpan has

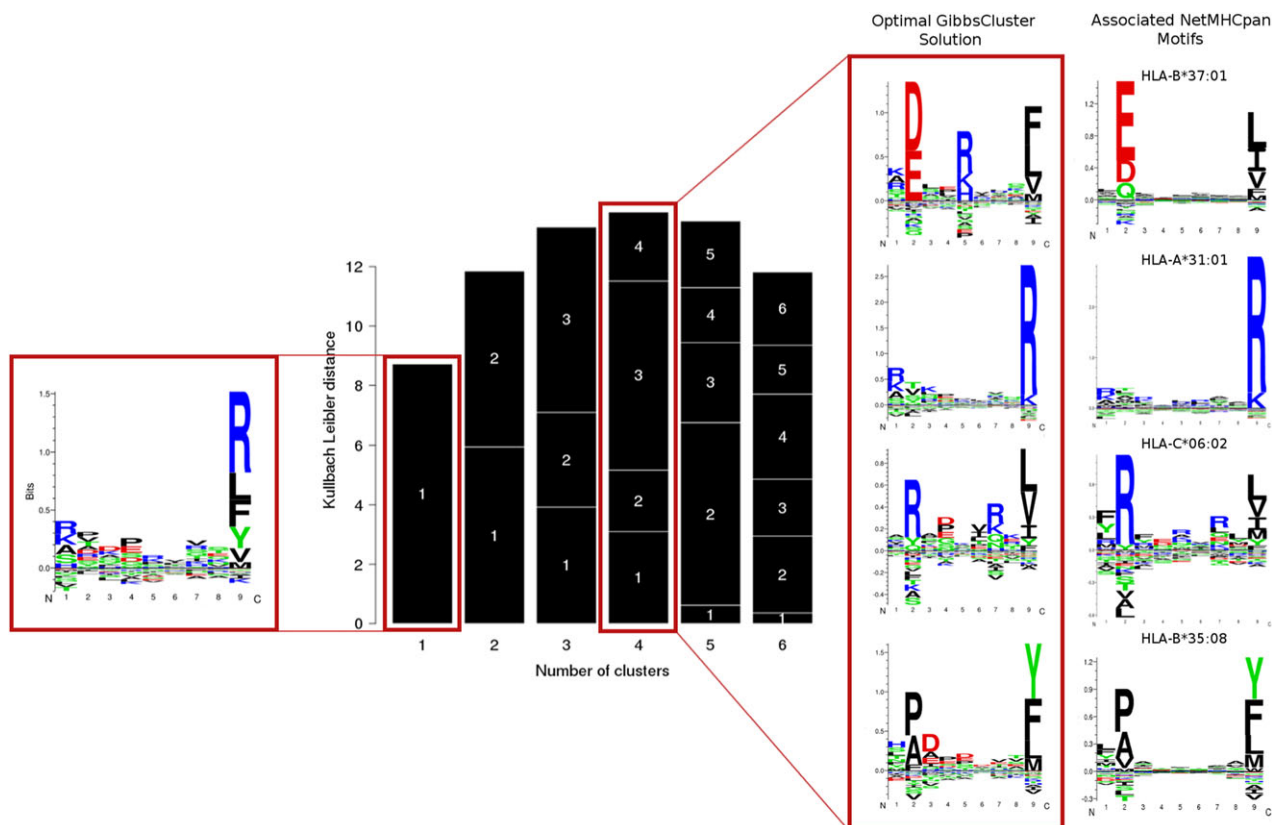


Figure 2. Clustering results for the HCC1143 cell line. The single cluster solution (left) is a mixture of multiple specificities, dominated by the most abundant alleles. The solution with highest information content corresponds to four clusters, with motifs highlighted in the red box (center). The motifs identified by unsupervised clustering show a remarkable correspondence with those predicted by NetMHCpan-3.0 (right). The GibbsCluster method was run using the default preset parameters for “MHC class I ligands of lengths 8–13,” except for the number of iterations which was set to 100 (slower but more accurate), and number of groups, which was allowed to vary between 1 and 6. NetMHCpan logos were obtained from the NetMHCpan-3.0 website (<http://www.cbs.dtu.dk/services/NetMHCpan-3.0/logos.php>) and were constructed from the top 1% scoring peptides from a large set of natural random peptides.

the potential advantage that at least a fraction of peptides could be assigned to HLA-C*04:01, a specificity that was not detected by unsupervised clustering. However, in cases where the haplotype is not fully characterized, or when the known MHC alleles have poorly studied motifs, the assignment by NetMHCpan would fail. This is exemplified by a recent study of bovine MHC ligands,^[38] for which the motifs derived by GibbsCluster differed dramatically from the assignments made by NetMHCpan due to paucity of training data available to NetMHCpan for these alleles. Note also, that the number of ligands discarded to the trash cluster using this approach was more than ten times higher compared to those discarded by GibbsCluster (463 versus 45).

4. MHC Class II, Mono-Allelic Cells

Analyzing MHC class II binding data is for many reasons more complex compared to MHC class I. First and foremost, the HLA class II binding groove is open at both ends, accommodating peptides of a wide range of length by letting them protrude at either terminus of the nonamer binding core. Sophisticated alignment methods are therefore required to identify the conserved binding preferences of MHC class II molecules.^[39–41]

Secondly, the binding motifs for MHC class II are in general more degenerate compared to the highly conserved MHC class I motifs.^[42,43] These observations make the analysis and interpretation of MHC class II binding data, including MS ligands, highly challenging.

In a recent paper by Ooi et al.,^[44] MS eluted ligand data were used to investigate how patients expressing different HLA class II alleles have different susceptibility to autoimmune diseases. To characterize the specificity for each allele, they generated transgenic mice bearing the human HLA-DR1 MHC class II allele. On these data, we illustrate how the GibbsCluster method can be used to identify the binding motif of MHC class II molecules from mono-allelic MS ligand data and at the same time remove potential outliers. The 5740 non-redundant raw eluted peptide sequences were uploaded to the GibbsCluster web server, setting the recommended preset parameters for MHC class II peptides, except for the number of iterations per sequence per temperature step (set to 100) and the number of temperature steps (set to 50); these parameters entail a slower, but more accurate, motif search. The method recovered the binding motif for allele HLA-DRB1*01:01, with strong amino acid preferences at anchor residues at P1, P4, P6, and P9 (Figure 3A). These preferences were observed both without (Figure 3A, left panel) or with a trash

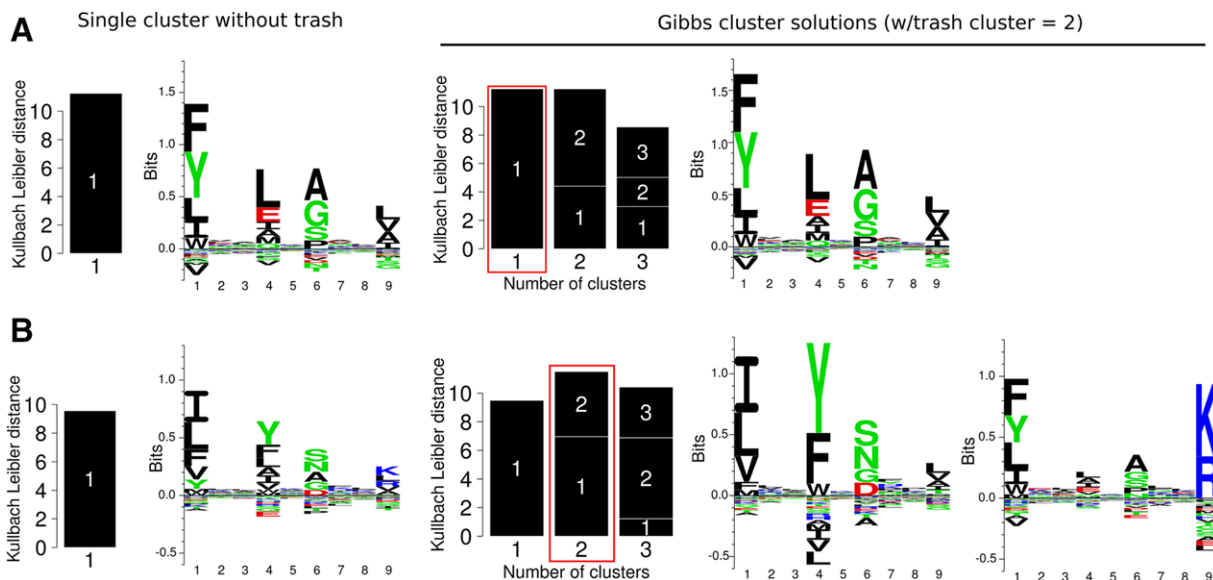


Figure 3. Sequence motifs identified by GibbsCluster-2.0 on MHC class II ligand data. The method identifies distinct amino acid preferences at the anchor positions P1, P4, P6, and P9 both without (left panels) and with (right panels) the trash cluster activated. A) Visualizing the motif and removing outliers from the mono-allelic human-DR1 mouse-transfected cell lines. B) Motif identification on mixed allelic data of DR15-DR51-EBV transformed cell lines.

cluster activated (Figure 3A, right panel). By activating the trash cluster option with a threshold of 2, 179 peptides (3% of data) were removed, and the logo showed a 20% increase in information content (Figure 3A, right panel).

5. MHC Class II, Poly-Allelic Cells

Another data set obtained from the Ooi et al. study^[44] consists of peptides eluted from HW09013 cells that express the HLA-DR15/DR51 class II alleles. On this poly-allelic data set of MS eluted ligands, we set out to demonstrate how the GibbsCluster can be used to separate multiple specificities in MHC class II ligand data. The set of 2782 unique eluted peptides was submitted to GibbsCluster, using the recommended preset parameters for MHC class II and allowing the program to search up to three clusters. The unfiltered, single-cluster solution shows a motif with the correct P1, P4, P6, and P9 anchor positions, but with low information content and preferences that are a mixture of the two alleles in the sample (Figure 3B, left panel). Activating the trash cluster with a threshold of 2, the maximum information content is observed for the solution with two clusters (Figure 3B, right panel). The amino acid preferences identified by GibbsCluster resemble previously published motifs derived from binding affinity data for HLA-DRB1*15:01 and HLA-DRB5*01:01,^[31,45] and closely overlap with the global peptidome of DR15/51 characterized in a recent study.^[46] Specifically, cluster 1 was composed of 1610 peptides (57.9%) and its motif resembles the HLA-DR15 binding preferences; cluster 2 comprised 1050 peptides (37.7%) and corresponds to the HLA-DR51 alleles; 122 peptides (4.4%) did not match to either group and were collected by the trash cluster.

In order to validate the solutions generated by GibbsCluster, we examined the composition of the clusters in terms of binding

potential predicted by NetMHCIIpan-3.1.^[47] Both for the mono-allelic DR1 and poly-allelic DR15/51 serotypes discussed above, we obtained predicted percentile rank scores for all peptides in the cluster solutions and in their relative trash cluster (Figure 4). The predicted median rank score for HLA-DRB1*01:01 in the DR1 cluster was 4% (first quartile (Q1) = 0.9, third quartile (Q3) = 12), whereas the trash cluster had a median rank score of 41% (Q1 = 20.5, Q3 = 75). In the poly-allelic data, cluster 1 was associated with HLA-DRB1*15:01, and showed a median rank score of 13% (Q1 = 5, Q3 = 30); cluster 2 was associated to HLA-DRB5*01:01 and obtained a median rank




Figure 4. NetMHCIIpan percentile rank score for GibbsCluster solutions in the DR1 and DR15/51 data sets. Percentile rank scores were predicted by netMHCIIpan-3.1 for each GibbsCluster group with matching alleles present in MS data samples. In the case of the mixed allele dataset DR15/51, peptides in the trash cluster were scored by NetMHCIIpan to both DRB1*15:01 and DRB5*01:01, selecting the lowest rank score of the two.

score of 4% ($Q1 = 1.1$, $Q3 = 11$); peptides in the trash cluster were evaluated against both alleles, assigning the best rank of the two, which resulted in an average rank score of 41% ($Q1 = 23$, $Q3 = 75$) (Figure 4). Overall, the NetMHCIIpan percentile score distributions suggest that the trash cluster could successfully collect peptides with very poor correspondence to the known preferences of the MHC class II molecules, and that probably derived either from incorrect spectral matches or from contaminants. The relatively high predicted rank values for the peptides mapped to the HLA-DRB1*15:01 cluster further suggest that the binding motif for this molecule predicted by NetMHCIIpan-3.1, which was trained on binding affinity data, shared a rather weak overlap with the binding motif contained within the MS ligand data. This observation underlines the high potential of MS ligand data to complement our knowledge on peptide characteristics required for MHC antigen presentation, as previously remarked for MHC class I.^[2–4,21,23]

6. Generating Prediction Models from MS Ligand Data

The approaches described so far in this report are mainly concerned with extracting and visualizing meaningful patterns within complex, often noisy, mixtures of peptides sequences. A further step is the generalization of the motifs identified in the data at hand, by constructing prediction models. Machine learning algorithms such as NNAlign,^[32] when provided with training examples suitably labeled (e.g., ligands vs non-ligands), can be instructed to automatically learn the features that distinguish positive from negative examples. Such models can then be applied on external data sets to discover more occurrences of the patterns learned on the training data. In the context of peptide-MHC interactions, a good prediction model should have the ability to capture the binding preferences contained in the training data, both in terms of sequence motifs and peptide length distribution. In the next two sections, we illustrate some simple examples of prediction models directly constructed from MHC class I and class II eluted ligands.

6.1. MHC Class I Prediction Model

As an example application, we continue with the Abelin ligand elution dataset previously analyzed and filtered using GibbsCluster-2.0 (Figure 1). For each of the representative alleles HLA-A*68:02, HLA-B*35:01, and HLA-B*57:01, we prepared a training set consisting of post-filtering ligands (positive instances) and random natural peptides (negative instances). Positive instances were labeled with a target value of 1, negatives with a target value of 0. In line with earlier work,^[4] the amount of random negatives was imposed to be the same for each length 8–13, and corresponded for each length to five times the amount of positives for the most abundant peptide length. This uniform length distribution of the random negatives was adopted as a background against which machine learning can be employed to learn the amino acid and length preference of the natural binders.

On each of the three data sets, we trained a prediction model with the NNAlign-2.0 web server, using the recommended preset options for MHC class I ligands of variable length. In a cross-validation experiment, the three models returned an area under the Receiver Operating Characteristic curve (AUC) of 0.961, 0.984, and 0.979, respectively. In order to derive the amino acid and peptide length preferences learned by the model, we used it to evaluate a large set of 900 000 random natural peptides with a flat length distribution, and extracted the top 0.1% scoring peptides. The composition of these high-scoring peptides should reflect the main preferences identified by the method to distinguish positive from negative instances. Indeed, the binding motif drawn from the top 0.1% peptides closely reflects the amino acid preferences of the training data (Figure 5A, B). Moreover, all three methods could capture the preference for 9mer peptides over other peptide lengths; 10mers were moderately allowed, 8mers and 11mers were observed more infrequently (Figure 5C).

6.2. MHC Class II Prediction Model

To illustrate how the NNAlign framework can be used to construct MHC class II prediction models, we go back to the DR1 and DR15/51 data sets from Ooi et al.^[44] previously filtered and clustered with GibbsCluster (Figure 3). To enrich the positive instances with artificial negative examples, a set of natural random negatives of lengths 11–19 amino acids was added to each eluted ligands data set. Positive instances were labeled with a target value of 1, negatives with a target value of 0. Similarly to the training set preparation described above for MHC class I, the amount of random negatives for each length corresponded to five times the amount of positives for the most abundant peptide length. For each of the three specificities deconvoluted by GibbsCluster in the DR1 and DR15/51 cells, we applied NNAlign-2.0 to generate a prediction model, using the preset parameters for MHC class II recommended by the NNAlign server. For the mono-allelic DR1 serotype, all ligands except those removed by the trash cluster were used to train a model. For the DR15/51 cells, for which the clustering analysis revealed two separate specificities, we generated a separate model from the ligands contained in each of the two clusters.

The three models revealed high internal consistency, with cross-validated performance of AUC = 0.952, 0.974, and 0.952, respectively. NNAlign automatically generates a matrix (and logo) representation of the motif learned by the method, constructed from the top 1% scoring predictions from a large set of random natural peptides. We may compare the motifs learned by NNAlign to: i) the binding preferences in the MS training data, identified by GibbsCluster, ii) the GibbsCluster motifs identified in tetramer-validated epitopes extracted from the IEDB for the three DR molecules, iii) the binding preferences predicted by NetMHCIIpan-3.1 for these DR molecules. In general, the motifs learned by the NNAlign models share a remarkable overall correspondence to the preferences found by GibbsCluster for the MS ligand data, with similar amino acid enrichments at the anchor positions P1, P4, and P6, as well as the strong P9 for the DR51-associated ligands (Figure 6, first and second columns). Likewise, the binding motifs constructed from

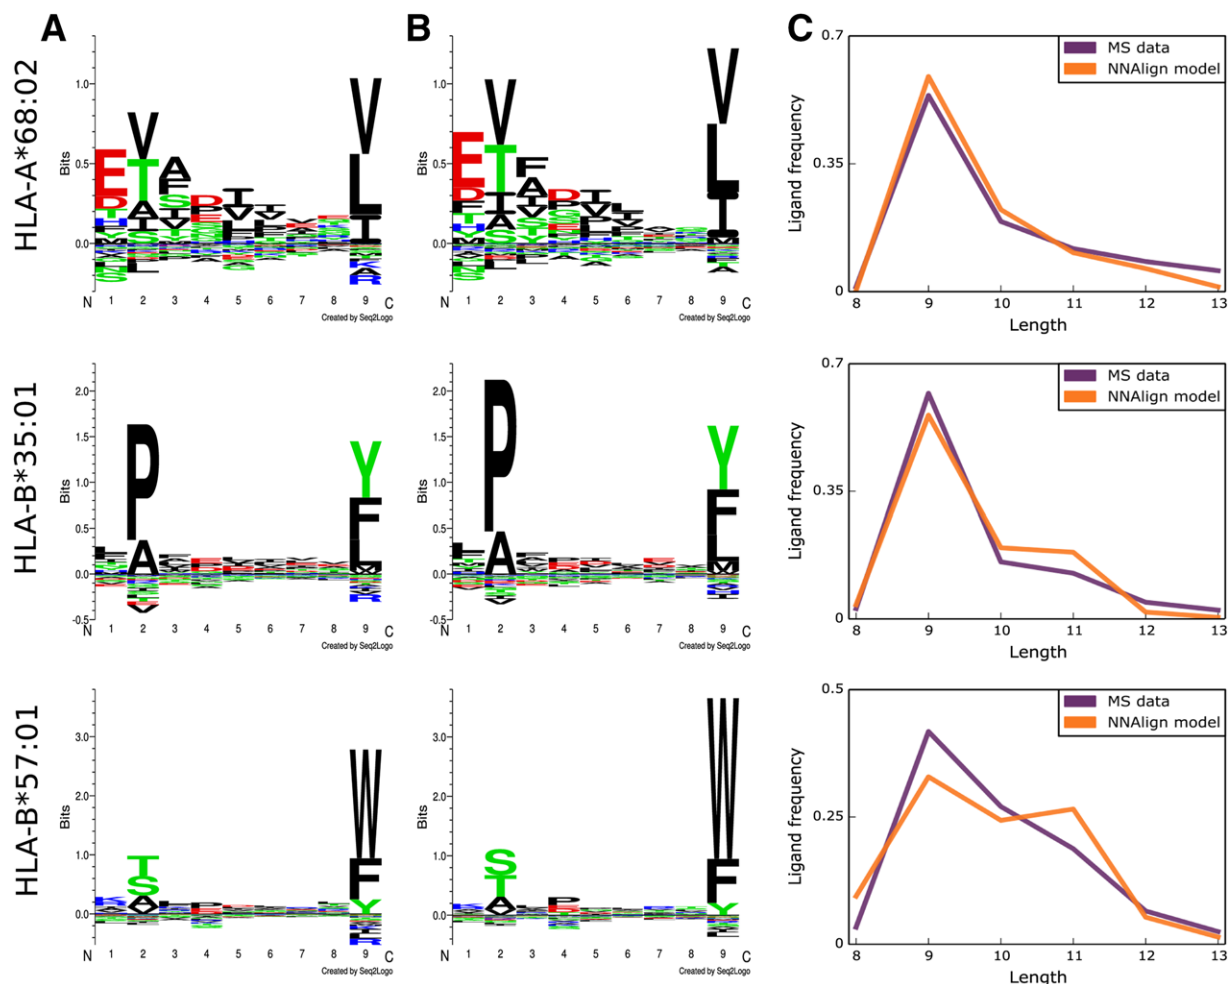


Figure 5. Generating prediction models from MS ligand data. A) Sequence motifs of the training data for three MHC class I alleles, aligned and filtered by GibbsCluster; B) sequence motifs captured by NNAlign-2.0; C) ligand length preferences in the training MS data compared to length preferences learned by the NNAlign model.

the rather small amount of tetramer-validated epitopes obtained from the Immune Epitope Database (IEDB)^[48] for the three DR molecules (231 for HLA-DRB1*01:01, 129 for HLA-DRB1*15:01, 73 for HLA-DRB5*01:01) correspond well with the motifs of the NNAlign models, and the MS ligand data (Figure 6, third column). In contrast, the logos derived from in vitro binding affinity data (NetMHCIIpan) in all cases show substantial differences to both the MS- and epitope-derived motifs (Figure 6, fourth column). These discrepancies are most pronounced for HLA-DRB1*15:01, where the NetMHCIIpan motif has weakly defined preferences at the anchor residues, and an enrichment of arginine (R) throughout the binding motif: a preference that is completely absent from the MS and epitope-derived motif. Another, more subtle difference is the enrichment of glutamic acid (E) at P4 in the MS and epitope motifs for HLA-DRB1*01:01; this preference is absent in the NetMHCIIpan motif. Finally, NetMHCIIpan displays a preference for R/K at position P8 for HLA-DRB5*01:01; this anchor is completely absent in the motif derived from MS and tetramer-validated epitope data. Taken together, these results show that ligand elution is a stronger correlate of epitope presentation than

peptide-MHC binding affinity, suggesting that epitope prediction models may greatly benefit from incorporating MS eluted ligand data.

7. Final Remarks

The binding specificities of MHC molecules have been traditionally characterized using in vitro assays of binding affinity. The peptide-MHC binding data amassed through decades of painstakingly low-throughput experiments have had a tremendous contribution to the characterization of the binding preference for the most prevalent MHC molecules, and more generally to the understanding of the peptide repertoire available for T-cell recognition. However, because of the extreme polymorphism of the MHC-encoding genes, with up to several thousand allelic variants per locus, the full characterization of their specificities remains infeasible. Tandem mass spectrometry has emerged in the past decade as a powerful, high-throughput alternative for the identification of peptides eluted on the surface of antigen-presenting cells.

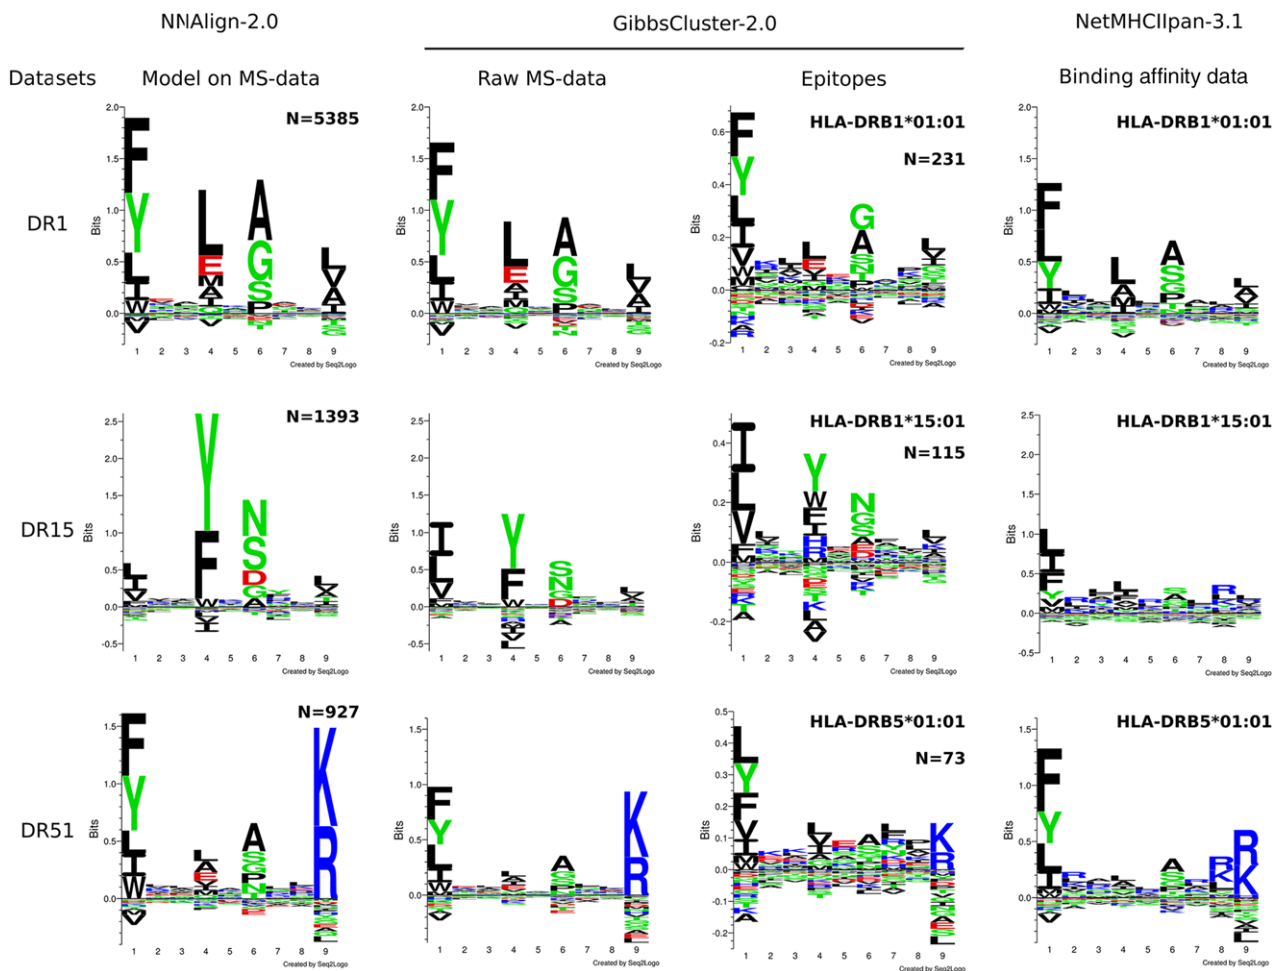


Figure 6. Comparison of motifs generated by different approaches for three HLA-DR alleles. NNAAlign-2.0 motifs were obtained by training artificial neural networks on each MS data set, and evaluating 100 000 random peptides. The top scoring 1% peptides were used to build logos. Raw MS data were aligned, clustered, and filtered in an unsupervised manner using GibbsCluster, with a trash cluster threshold = 2. The same procedure was applied to tetramer-positive data downloaded from the IEDB. Note that due to small data set size, epitope logos are shown in a different y-axis scale. Binding motifs for NetMHCIIpan-3.1 were determined by evaluating 100 000 random peptides, and visualizing the core motif of the top 1% scoring sequences.

The appeal of MS-based techniques does not only reside in the sheer amount of ligand data that can be detected in a single experiment. Because MS ligands are derived from a biological system that incorporates all properties of antigen presentation including binding affinity, binding stability, proper peptide processing and translocation, and impact of MHC binding chaperones, these techniques should capture additional signals besides the binding affinity measurable by *in vitro* assays. Accurate tools for the identification of sequence motifs in eluted ligand datasets are essential to interpret the patterns underlying the immunopeptidome and to benefit from this data deluge.

In this report, we described some straightforward, efficient approaches to extract motifs from immunopeptidomes in a number of scenarios commonly encountered in the field. We outlined analyses for MHC class I and class II, both in cell lines expressing a single MHC allele and in unmodified cells with multiple MHC allelic variants. GibbsCluster^[20] is our tool of choice because it can effectively remove residual contaminants after FDR filtering,

deconvolute multiple motifs in a mixture of peptides of variable length, and because it works both for MHC class I and class II ligands. In general, MHC class I molecules have strong, well-defined motifs, and even in samples containing several specificities it is often feasible to separate them into individual clusters. Unambiguously associating each cluster to individual MHC molecules remains an unresolved problem, especially for alleles with unknown binding motifs. So far only Bassani et al.^[27] have attempted to tackle this question, exploiting the co-occurrence of MHC class I alleles across different data sets of known haplotype to assign motifs to individual alleles. More work along these lines is needed to automatically annotate the MHC restriction of peptides in poly-allelic datasets. The current implementation of GibbsCluster assumes that each peptide is restricted to one and only one MHC molecule. When cells express different alleles with similar binding motifs, or in the case of MHC class II ligands binding to multiple alleles in different alignment frames, it is likely that an individual peptide can act as ligand for

multiple MHCs in a mixture. Future improvements to the algorithm should aim to address this limitation and account for potential multiple restrictions of individual ligands.

Ultimately, prediction methods can only be as good as the data used to train them. While MHC ligands sequences obtained by mass spectrometry show remarkable reproducibility and produce binding motifs consistent with those derived with more low-throughput assays, there remain several potential sources of error and bias in MS-based pipelines for ligand sequencing. For example, there is a documented underrepresentation of cysteine in MHC ligand data sets, as this amino acid interferes with MS precursor fragmentation.^[3,27] Different software tools for spectrum-peptide mapping use different functions to score candidate sequences, and they will generally identify non-identical sets of ligands. Post-translational modifications (PTMs) have also been shown to have a role in shaping the MHC ligand repertoire.^[49] However, accounting for such modified residues further complicates accurate spectrum-peptide matching and PTMs are often not comprehensively considered in MS pipelines. Finally, common contaminants such as keratin and histone proteins are often co-eluted with MHC ligands and add a level of noise to the sequenced immunopeptidome.^[34,35] Reducing biases and sources of error in the data-generation pipelines will also inevitably affect in a positive way the data interpretation and the prediction tools constructed on these data.

A number of recent reports have described the first prediction methods trained on MHC class I ligand elution data from mass spectrometry.^[3,4,28,50] Their results indicate that methods trained on naturally presented peptides largely outperform prediction methods trained solely on in vitro binding affinity data when it comes to the identification of MHC ligands and epitopes. No reports have yet been published on models directly trained on MHC class II eluted ligands. Because the performance of MHC class II prediction methods still lags far behind their class I counterparts for epitope prediction, antigen processing factors are likely to play a major role in the generation of MHC class II ligands. Incorporating naturally processed ligand from MS experiments in the training pipelines of MHC class II prediction methods is an exciting and yet unexplored opportunity to close that gap. A simple but powerful approach to generate prediction models from ligand data is the NNAlign method.^[32] We illustrated the construction of models from MS eluted ligands both for MHC class I and MHC class II, and showed that they capture the preferences of the training data both in terms of binding motif and ligand length distribution. Taken together, these computational tools allow researchers to interpret motifs contained in immunopeptidomes and generate prediction models to scan protein databases for epitope candidates.

Supporting Information

Supporting Information is available from the Wiley Online Library or from the author.

Acknowledgements

B.A. and C.B. contributed equally to this work. This work was supported by federal funds from the National Institute of Allergy and

Infectious Diseases, National Institutes of Health, Department of Health and Human Services, under Contract No. HHSN272201200010C, and by the Agencia Nacional de Promoción Científica y Tecnológica, Argentina (PICT-2016-0089).

Conflict of Interest

The authors have declared no conflict of interest.

Keywords

GibbsCluster, mass spectrometry, MHC, prediction models, sequence motifs

Received: October 6, 2017

Revised: December 15, 2017

- [1] E. Caron, D. J. Kowalewski, C. Chiek Koh, T. Sturm, H. Schuster, R. Aebersold, *Mol. Cell. Proteomics* **2015**, *14*, 3105.
- [2] M. Bassani-Sternberg, D. Gfeller, *J. Immunol. Baltim. Md 1950* **2016**, *197*, 2492.
- [3] J. G. Abelin, D. B. Keskin, S. Sarkizova, C. R. Hartigan, W. Zhang, J. Sidney, J. Stevens, W. Lane, G. L. Zhang, T. M. Eisenhaure, K. R. Clauser, N. Hacohen, M. S. Rooney, S. A. Carr, C. J. Wu, *Immunity* **2017**, *46*, 315.
- [4] V. I. Jurtz, S. Paul, M. Andreatta, P. Marcatili, B. Peters, M. Nielsen, *bioRxiv* **2017**.
- [5] N. Ternette, H. Yang, T. Partridge, A. Llano, S. Cedeño, R. Fischer, P. D. Charles, N. L. Dudek, B. Mothe, M. Crespo, W. M. Fischer, B. T. M. Korber, M. Nielsen, P. Borrow, A. W. Purcell, C. Brander, L. Dorrell, B. M. Kessler, T. Hanke, *Eur. J. Immunol.* **2016**, *46*, 60.
- [6] J. C. Yaciuk, M. Skaley, W. Bardet, F. Schafer, D. Mojsilovic, S. Cate, C. J. Stewart, C. McMurtrey, K. W. Jackson, R. Buchli, A. Olvera, S. Cedeño, M. Plana, B. Mothe, C. Brander, J. T. West, W. H. Hildebrand, *J. Virol.* **2014**, *88*, 12992.
- [7] C. Berlin, D. J. Kowalewski, H. Schuster, N. Mirza, S. Walz, M. Hande, B. Schmid-Horch, H. R. Salih, L. Kanz, H.-G. Rammensee, S. Stevanović, J. S. Stickel, *Leukemia* **2015**, *29*, 647.
- [8] M. Bassani-Sternberg, G. Coukos, *Curr. Opin. Immunol.* **2016**, *41*, 9.
- [9] S. Kalaora, E. Barnea, E. Merhavi-Shoham, N. Qutob, J. K. Teer, N. Shimony, J. Schachter, S. A. Rosenberg, M. J. Besser, A. Admon, Y. Samuels, *Oncotarget* **2016**, *7*, 5110.
- [10] M. Bassani-Sternberg, E. Bräunlein, R. Klar, T. Engleitner, P. Sinitcyn, S. Audehm, M. Straub, J. Weber, J. Slotta-Huspenina, K. Specht, M. E. Martignoni, A. Werner, R. Hein, D. H. Busch, C. Peschel, R. Rad, J. Cox, M. Mann, A. M. Krackhardt, *Nat. Commun.* **2016**, *7*, 13404.
- [11] J. Cox, M. Mann, *Nat. Biotechnol.* **2008**, *26*, 1367.
- [12] B. Ma, K. Zhang, C. Hendrie, C. Liang, M. Li, A. Doherty-Kirby, G. Lajoie, *Rapid Commun. Mass Spectrom. RCM* **2003**, *17*, 2337.
- [13] J. Zhang, L. Xin, B. Shan, W. Chen, M. Xie, D. Yuen, W. Zhang, Z. Zhang, G. a. Lajoie, B. Ma, *Mol. Cell. Proteomics* **2012**, *11*, M111.010587.
- [14] M. Brosch, L. Yu, T. Hubbard, J. Choudhary, *J. Proteome Res.* **2009**, *8*, 3176.
- [15] J. P. Murphy, P. Konda, D. J. Kowalewski, H. Schuster, D. Clements, Y. Kim, A. M. Cohen, T. Sharif, M. Nielsen, S. Stevanovic, P. W. Lee, S. Gujar, *J. Proteome Res.* **2017**, *16*, 1806.
- [16] T. Delong, T. A. Wiles, R. L. Baker, B. Bradley, G. Barbour, R. Reisdorph, M. Armstrong, R. L. Powell, N. Reisdorph, N. Kumar, C. M.

- Elso, M. DeNicola, R. Bottino, A. C. Powers, D. M. Harlan, S. C. Kent, S. I. Mannering, K. Haskins, *Science* **2016**, 351, 711.
- [17] J. Liepe, F. Marino, J. Sidney, A. Jeko, D. E. Bunting, A. Sette, P. M. Kloetzel, M. P. H. Stumpf, A. J. R. Heck, M. Mishto, *Science* **2016**, 354, 354.
- [18] H. G. Rammensee, T. Friede, S. Stevanović, *Immunogenetics* **1995**, 47, 178.
- [19] M. Andreatta, O. Lund, M. Nielsen, *Bioinformatics* **2013**, 29, 8.
- [20] M. Andreatta, B. Alvarez, M. Nielsen, *Nucleic Acids Res.* **2017**, DOI 10.1093/nar/gkx248.
- [21] M. Bassani-Sternberg, S. Pletscher-Frankild, L. J. Jensen, M. Mann, *Mol. Cell. Proteomics MCP* **2015**, 14, 658.
- [22] P. Probst, J. Kopp, A. Oxenius, M. P. Colombo, D. Ritz, T. Fugmann, D. Neri, *Cancer Res.* **2017**, 77, 3644.
- [23] D. Ritz, A. Gloger, B. Weide, C. Garbe, D. Neri, T. Fugmann, *Proteomics* **2016**, 16, 1570.
- [24] A. Sofron, D. Ritz, D. Neri, T. Fugmann, *Eur. J. Immunol.* **2016**, 46, 319.
- [25] T. Fugmann, A. Sofron, D. Ritz, F. Bootz, D. Neri, *J. Immunol. Baltim. Md 1950* **2017**, 198, 1357.
- [26] G. P. M. Mommen, F. Marino, H. D. Meiring, M. C. M. Poelen, J. A. M. van Gaans-van den Brink, S. Mohammed, A. J. R. Heck, C. A. C. M. van Els, *Mol. Cell. Proteomics MCP* **2016**, 15, 1412.
- [27] M. Bassani-Sternberg, C. Chong, P. Guillaume, M. Solleder, H. Pak, P. O. Gannon, L. E. Kandalaf, G. Coukos, D. Gfeller, *PLoS Comput. Biol.* **2017**, 13, e1005725.
- [28] M. Nielsen, T. Connelley, N. Ternette, *bioRxiv* **2017**, DOI 10.1101/195016.
- [29] S. Tenzer, B. Peters, S. Bulik, O. Schoor, C. Lemmel, M. M. Schatz, P.-M. Kloetzel, H.-G. Rammensee, H. Schild, H.-G. Holzhütter, *Cell. Mol. Life Sci. CMLS* **2005**, 62, 1025.
- [30] M. Harndahl, M. Rasmussen, G. Roder, I. Dalgaard Pedersen, M. Sørensen, M. Nielsen, S. Buus, *Eur. J. Immunol.* **2012**, 42, 1405.
- [31] M. Andreatta, C. Schafer-Nielsen, O. Lund, S. S. S. Buus, M. Nielsen, *PLoS ONE* **2011**, 6, e26781.
- [32] M. Nielsen, M. Andreatta, *Nucleic Acids Res.* **2017**, DOI 10.1093/nar/gkx276.
- [33] E. E. Sercarz, E. Maverakis, *Nat. Rev. Immunol.* **2003**, 3, 621.
- [34] K. Hodge, S. T. Have, L. Hutton, A. I. Lamond, *J. Proteomics* **2013**, 88, 92.
- [35] D. Mellacheruvu, Z. Wright, A. L. Couzens, J.-P. Lambert, N. A. St-Denis, T. Li, Y. V. Miteva, S. Hauri, M. E. Sardu, T. Y. Low, V. A. Halim, R. D. Bagshaw, N. C. Hubner, A. Al-Hakim, A. Bouchard, D. Faubert, D. Fermin, W. H. Dunham, M. Goudreault, Z.-Y. Lin, B. G. Badillo, T. Pawson, D. Durocher, B. Coulombe, R. Aebersold, G. Superti-Furga, J. Colinge, A. J. R. Heck, H. Choi, M. Gstaiger, S. Mohammed, I. M. Cristea, K. L. Bennett, M. P. Washburn, B. Raught, R. M. Ewing, A.-C. Gingras, A. I. Nesvizhskii, *Nat. Methods* **2013**, 10, 730.
- [36] M. Nielsen, M. Andreatta, *Genome Med.* **2016**, 8, 33.
- [37] M. Rasmussen, M. Harndahl, A. Stryhn, R. Boucherma, L. L. Nielsen, F. A. Lemonnier, M. Nielsen, S. Buus, *J. Immunol. Baltim. Md 1950* **2014**, 193, 4790.
- [38] M. Nielsen, T. Connelley, N. Ternette, *J. Proteome Res.* **2017**, DOI 10.1021/acs.jproteome.7b00675.
- [39] M. Nielsen, O. Lund, S. Buus, C. Lundegaard, *Immunology* **2010**, 130, 319.
- [40] M. Andreatta, V. I. Jurtz, T. Kaever, A. Sette, B. Peters, M. Nielsen, *Immunology* **2017**, 152, 255.
- [41] M. Nielsen, C. Lundegaard, P. Worning, C. S. Hvid, K. Lamberth, S. Buus, S. Brunak, O. Lund, *Bioinforma. Oxf. Engl.* **2004**, 20, 1388.
- [42] T. Sturniolo, E. Bono, J. Ding, L. Radrizzani, O. Tuercu, U. Sahin, M. Braxenthaler, F. Gallazzi, M. P. Protti, F. Sinigaglia, J. Hammer, *Nat. Biotechnol.* **1999**, 17, 555.
- [43] M. Andreatta, M. Nielsen, *Immunology* **2012**, DOI 10.1111/j.1365-2567.2012.03579.x.
- [44] J. D. Ooi, J. Petersen, Y. H. Tan, M. Huynh, Z. J. Willett, S. H. Ramarathinam, P. J. Eggenhuizen, K. L. Loh, K. A. Watson, P. Y. Gan, M. A. Alikhan, N. L. Dudek, A. Handel, B. G. Hudson, L. Fugger, D. A. Power, S. G. Holt, P. T. Coates, J. W. Gregersen, A. W. Purcell, S. R. Holdsworth, N. L. La Gruta, H. H. Reid, J. Rossjohn, A. R. Kitching, *Nature* **2017**, 545, 243.
- [45] A. B. Vogt, H. Kropshofer, H. Kalbacher, M. Kalbus, H. G. Rammensee, J. E. Coligan, R. Martin, *J. Immunol. Baltim. Md 1950* **1994**, 153, 1665.
- [46] E. M. Scholz, M. Marcilla, X. Daura, D. Arribas-Layton, E. A. James, I. Alvarez, *Front. Immunol.* **2017**, 8, 984.
- [47] M. Andreatta, E. Karosiene, M. Rasmussen, A. Stryhn, S. S. Buus, M. Nielsen, *Immunogenetics* **2015**, 67, 641.
- [48] R. Vita, J. A. Overton, J. A. Greenbaum, J. Ponomarenko, J. D. Clark, J. R. Cantrell, D. K. Wheeler, J. L. Gabbard, D. Hix, A. Sette, B. Peters, *Nucleic Acids Res.* **2015**, 43, D405.
- [49] C. A. C. M. van Els, V. Corbière, K. Smits, J. A. M. van Gaans-van den Brink, M. C. M. Poelen, F. Mascart, H. D. Meiring, C. Loch, *Front. Immunol.* **2014**, 5, 361.
- [50] S. Giguère, A. Drouin, A. Lacoste, M. Marchand, J. Corbeil, F. Laviolette, *J. Immunol. Methods* **2013**, 400–401, 30.