

Improved Prediction of MHC II Antigen Presentation through Integration and Motif Deconvolution of Mass Spectrometry MHC Eluted Ligand Data

Birkir Reynisson, Carolina Barra, Saghar Kaabinejadian, William H. Hildebrand, Bjoern Peters, and Morten Nielsen*

Cite This: *J. Proteome Res.* 2020, 19, 2304–2315

Read Online

ACCESS |

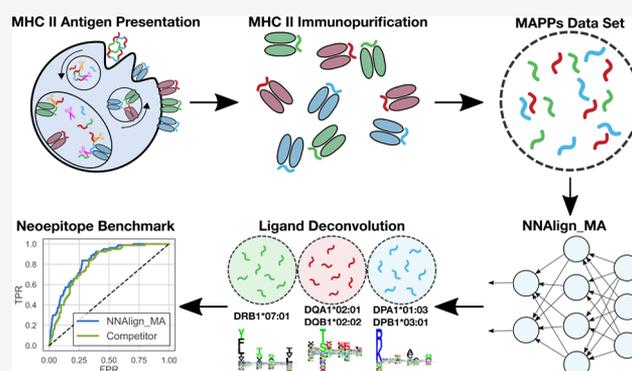
Metrics & More

Article Recommendations

Supporting Information

ABSTRACT: Major histocompatibility complex II (MHC II) molecules play a vital role in the onset and control of cellular immunity. In a highly selective process, MHC II presents peptides derived from exogenous antigens on the surface of antigen-presenting cells for T cell scrutiny. Understanding the rules defining this presentation holds critical insights into the regulation and potential manipulation of the cellular immune system. Here, we apply the NNAlign_MA machine learning framework to analyze and integrate large-scale eluted MHC II ligand mass spectrometry (MS) data sets to advance prediction of CD4+ epitopes. NNAlign_MA allows integration of mixed data types, handling ligands with multiple potential allele annotations, encoding of ligand context, leveraging information between data sets, and has pan-specific power allowing accurate predictions outside the set of molecules included in the training data. Applying this framework, we identified accurate binding motifs of more than 50 MHC class II molecules described by MS data, particularly expanding coverage for DP and DQ beyond that obtained using current MS motif deconvolution techniques. Furthermore, in large-scale benchmarking, the final model termed NetMHCIIpan-4.0 demonstrated improved performance beyond current state-of-the-art predictors for ligand and CD4+ T cell epitope prediction. These results suggest that NNAlign_MA and NetMHCIIpan-4.0 are powerful tools for analysis of immunopeptidome MS data, prediction of T cell epitopes, and development of personalized immunotherapies.

KEYWORDS: machine learning, bioinformatics, immunoinformatics, immunology, MHC II, antigen presentation, mass spectrometry, immunopeptidomics, neoepitopes



INTRODUCTION

Major histocompatibility complex class II (MHC II) molecules play a pivotal role in the adaptive immune system. Antigen-presenting cells (APCs) display MHC II molecules in complex with peptides¹ on their surface. These peptides are products of extracellular proteins internalized by APCs and proteolytically digested in endocytic compartments. During protein degradation, the MHC II binding cleft is loaded with peptide fragments of the antigen and peptides that bind stably (forming a pMHCII complex) are shuttled to the cell surface for presentation to T-helper (Th) cells of the immune system.² Th cells scrutinize the surface of APCs and if the T cell receptor recognizes a pMHCII complex, Th cells can become activated. Peptides that cause T cell activation are termed T cell epitopes. Regulation of Th cell activation is critical because they coordinate the activation of effector cells. Peptide–MHC presentation is a necessary and highly selective step in the process of T cell activation, and characterizing the rules

defining this presentation is essential for our understanding of cellular immunity.

Given this, characterizing the rules of MHC II peptide presentation has been the focus of considerable research efforts. MHC II is a heterodimer, the α and β chains of which together form the peptide-binding cleft. In humans, the human leukocyte antigen (HLA) of MHC class II is encoded by three different loci (HLA-DR, -DQ, and -DP).³ The corresponding HLA genes have numerous allelic variants with polymorphisms that are mostly clustered around the residue locations forming the peptide-binding cleft, resulting in a wide range of distinct

Received: December 20, 2019

Published: April 18, 2020



peptide-binding specificities. MHC II has an open binding cleft, allowing it to interact with peptides of a broad length range, most commonly of 13–25 amino acids.⁴ The binding cleft of MHC II interacts predominantly with a 9-mer register of the interacting peptide, termed the binding core.⁵ The placement of this binding core varies between peptides, resulting in peptide-flanking residues (PFRs) of differing lengths protruding out from the binding cleft of the MHC II molecule. PFRs influence both the interaction with MHC II and the activation of T cells.⁶ These facts together make the study and prediction of MHC II peptide interaction highly challenging.

Traditionally, *in vitro* peptide–MHC binding affinity (BA) assays have been used to generate data to characterize the specificity of MHC II molecules,⁷ and a range of machine-learning prediction models have been developed from this data to identify the rules of peptide–MHC binding (reviewed in refs 8 and 9). However, evidence suggests peptide–MHCII BA to be a relatively weak correlate of MHC antigen presentation¹⁰ and peptide immunogenicity.¹¹ Likewise, several studies have demonstrated that MHC-II peptide-binding prediction models can benefit from being trained on so-called immunopeptidome data obtained by liquid chromatography coupled mass spectrometry (LC–MS/MS) (reviewed in refs 8 and 12). In a typical MHC II immunopeptidome eluted ligand (EL) assay, MHC molecules are immunopurified from lysed APCs. The bound peptide ligands are next chromatographically eluted from MHC molecules and sequenced via MS/MS.^{13,14} The result of such an assay is a list of peptide sequences restricted to at least one of the MHC II molecules expressed by the interrogated cell line.

The biological relevance of EL data is a major advantage over BA data. EL data implicitly contain signals from steps of MHC II antigen presentation, such as antigen digestion, MHC loading of ligands, and cell surface transport. Amino acid preferences in termini of immunopeptidome ligands show clear evidence of proteolytic digestion, and ligand prediction has been improved by explicitly encoding this signal in prediction models.^{15–17} Models trained on EL data can learn the ligand length preference of MHC molecules because this preference is inherent to EL data but not to BA data (examples include refs 15, 17, and 18).

Although most prediction methods are trained either on BA or MHC MS EL data, the work by Jurtz et al.¹⁹ proposed an architecture where the two data types were effectively integrated into one machine-learning framework. This framework has later been applied with high success to train models with state-of-the-art performance for prediction of MHC class I antigen presentation and CD8 T cell epitopes.^{19,20} Recent framework refinements also cover MHC class II data and demonstrate how the integration of antigen-processing signals contained within and flanking the MS EL data boosts predictive power.¹⁵

The extensive polymorphism of MHC combined with the high experimental cost burden associated with characterizing the specificity of individual MHC molecules makes specificity characterization of all MHC molecules a prohibitively expensive undertaking. Given this, a proposed solution is pan-specific prediction models, which are trained on peptide interaction data covering large and diverse sets of MHCs with the purpose of learning the associations between the MHC protein sequence and its peptide specificity.^{21,22} The value of

these models lies in their ability to predict peptide–MHC binding for all alleles of the known sequence, including those characterized by limited or even no peptide-binding data.²³

Ligands eluted from cell lines expressing only one MHC molecule can be unambiguously annotated to that allele and are termed single allele (SA) ligands. Such data are generated from genetically engineered cell lines^{24,25} or careful experimental design, that is, matching an immunopurification antibody specific to an MHC locus with a cell line homozygous to said locus (as performed most often when conducting an HLA-DR specific LC–MS/MS experiment). However, this scenario is the exception in immunopeptidomics studies. EL data from patient samples will more often be composed of peptides of mixed specificities corresponding to the different MHC molecules expressed in the given APC. Such data are termed multi-allele (MA) ligands, and annotating MA ligands to their respective MHC in such data is not a trivial task. Ligands from MA data can be deconvoluted into separate motifs in an unsupervised manner with tools such as GibbsCluster²⁶ and MixMHCp,²⁷ but this still leaves the task of assigning motifs to their respective alleles.¹²

We recently proposed a solution to this critical challenge of MA data by an extension of the NNAlign algorithm termed NNAlign_MA.²⁸ NNAlign_MA handles MA data naturally by annotating MA data during training in a semi-supervised manner based on MHC co-occurrence, MHC exclusion, and pan-specific binding prediction. This algorithm builds on similar principles as the MixMHCp²⁷ and MoDec¹⁷ methods described earlier, however with the critical difference that NNAlign_MA allows to leverage information between data sets enabling accurate motif deconvolution also for data sets and alleles with limited ligand coverage (such as HLA-C for class I²⁸). In the work by Alvarez et al., the ability of NNAlign_MA to automatically deconvolute and learn binding motifs from complex EL MA data was showcased in several scenarios, and the algorithm was demonstrated to have a high potential for deconvoluting MA data and to construct accurate pan-specific predictors for MHC antigen presentation. For MHC II, this early work was conducted on a small MA data set covering mainly HLA-DR alleles, and also the NNAlign_MA modeling framework used did not incorporate signals of MHC II antigen processing. Here, we extend this work adapting NNAlign_MA to allow explicit encoding of ligand context to learn signals of antigen processing. Next, we apply the method to perform motif deconvolution and train a pan-specific prediction model of MHC II antigen presentation directly from an extensive data set of MA data covering HLA-DR, HLA-DQ, and HLA-DP. Benefitting from the unique abilities of NNAlign_MA to integrate mixed data types, handling ligands with multiple potential allele annotations, encoding signals of antigen processing, and leveraging information between data sets, we investigate its power for motif deconvolution, in particular, for DQ and DP molecules covered with few ligands in individual MA data sets. Also, the resulting pan-specific predictor, termed NetMHCIIpan-4.0, is benchmarked against the current state-of-the-art methods in terms of performance for ligand and CD4+ T cell epitope prediction.

MATERIALS AND METHODS

BA Data

BA data were gathered from the NetMHCIIpan-3.2 publication.²⁹ In line with the EL data set, only BA peptides of length 13–21 were included in the analysis. After filtering, the data set contained 131,077 peptides covering 59 HLA-II molecules. Peptide BA data are quantitative, and the measured BA was transformed to fall in the range [0, 1], as described in ref 30.

EL Data—Extraction

All MHC ligands in the immune epitope database (IEDB)³¹ were downloaded (January 28, 2019) and this table served as a guide to identify publications with MHC II EL data. Publications with full HLA class II typing and more than 1000 ligands were selected for data extraction. Ligand data were subsequently extracted from the individual publications.^{32–41} The Supporting Information of the publications was processed and ligands were extracted along with their source proteins and associated MHC molecules. Some additional data sets were found via literature search^{17,25,42–45} and a small in-house database of HLA-DR EL data was obtained from homozygous cell lines. Only cell lines with more than 250 measured ligands were included in the final data set.

The resulting data sets consist of tables of ligands, their source proteins, and a list of potential interacting MHC II molecules (the MHC molecules expressed in the given cell line). In cases of DQ or DP heterozygosity, all possible α and β chain combinations were included in the list of possible MHC II molecules. All post-translationally modified peptides were excluded, and only ligands of length 13–21 were included in the analysis (excluding 16.1% of the total extracted EL data), resulting in a total of 136 data sets, covering 381,066 MHC measurements, and 81 distinct MHC II alleles (3 DPA, 16 DPB, 13 DQA, 12 DQB, 34 DRB, and 3 H-2). For a complete overview of the data, refer to Table S1.

EL Data—Negative Peptide Generation

EL data contain only positive examples of MHC ligands. Negative data were defined as described earlier^{15,20} by randomly sampling peptides from the UniProt database.⁴⁶ Negatives were generated per data set and follow a uniform length distribution in the range 13–21. For each length, the number of negatives was equal to 5 times the number of ligands at the maximum of the ligand length distribution.

EL Data—Context

Terminal regions inside and outside of ligands were encoded to capture signals of proteolytic digestion, as described previously.¹⁵ A total of 12 residues were encoded for each ligand: six PFRs (three residues from the N-terminal and three from the C-terminal) and six ligand context residues from the source protein sequence (three residues upstream the ligand N-terminus and three residues downstream the C-terminus). Ligands that could not be mapped to a source protein were removed from the dataset (~1%). BA data received a context of Xs. Negative EL data in training sets were also assigned context from their sampled source protein.

Data Set Partitioning

NNAlign_MA was trained in a 5-fold cross-validation manner. Data were partitioned via common motif clustering⁴⁷ to ensure that no partition shared 9-mer subsequences. BA and EL data

were clustered simultaneously and next separated, resulting in a total of 10 partitions, 5 for EL data and 5 for BA data.

Epitopes

T cell epitopes were downloaded from the IEDB (downloaded: July 1, 2019) and filtered to include only positive MHC II peptide epitopes of length 13–21. Furthermore, epitopes measured by intracellular cytokine staining (ICS) or tetramer/multimer assays were included in the benchmark, resulting in a total of 1469 epitopes.

The *F*-rank evaluates the ability of prediction methods to prioritize epitope discovery. In short, epitopes are ranked by prediction scores among their length-matched peptides from the antigen, and the *F*-rank score is the percentage of false positives, that is, peptides with prediction scores higher than measured epitopes. An alternative *F*-rank scoring scheme incorporates the effects of antigen processing. Here, each antigen was in silico digested into *k*-mers (peptides of length 13–17 amino acids) and scores were predicted for each. Next, epitope length-matched peptides from an antigen were assigned a score by summing prediction scores of all *k*-mers whose predicted binding core overlaps completely with the given peptide. To remove noise in the data set, and filter out potential false positive epitopes, the data set was filtered to only include epitopes with *F*-rank values 10.0 for at least one of the methods benchmarked.

NNAlign_MA Training

The complete model consisted of an ensemble of 150 networks with 20, 40, and 60 hidden neurons with 10 random weight initializations for each of the 5 cross-validation folds (3 architectures, 10 seeds, and 5 folds). All models were trained using back-propagation with stochastic gradient descent, for 400 epochs, without early stopping, and a constant learning rate of 0.05. Only SA data were included in training for a burn-in period of 20 epochs. Subsequent training cycles included MA data.

NNAlign_MA Evaluation and MHC Restriction Deconvolution

MHC annotation for MA datasets is based on which MHC molecule expressed in a given cell line has the highest prediction score for a given ligand. Because of differences in prediction score distributions between MHC molecules imposed by the training data composition, basing assignment on raw prediction scores leads to a biased assignment, favoring highly represented data in the training set. This bias is addressed during the model training using rescaling,²⁸ and a similar approach was used for MA data evaluation. Here, normalized prediction scores for each molecule were generated by ranking them against a distribution of prediction scores of random natural peptides (independent of training data, 20k peptides each for length 13–21). Ranked prediction scores make the model output more interpretable and make relative comparison across molecules fairer.

Sequence Logos

Sequence logos were generated with Seq2Logo.⁴⁸ The amino acid background frequency for all logos was constructed from the set of ligand source proteins. MHC II binding motifs detected in EL data sets were generated from ligand-binding registers as predicted by NNAlign_MA, using the default Seq2Logo settings. To account for noise in EL data, ligands with a rank score greater than 20% were excluded and binding motif logos were generated for MHC-II molecules charac-

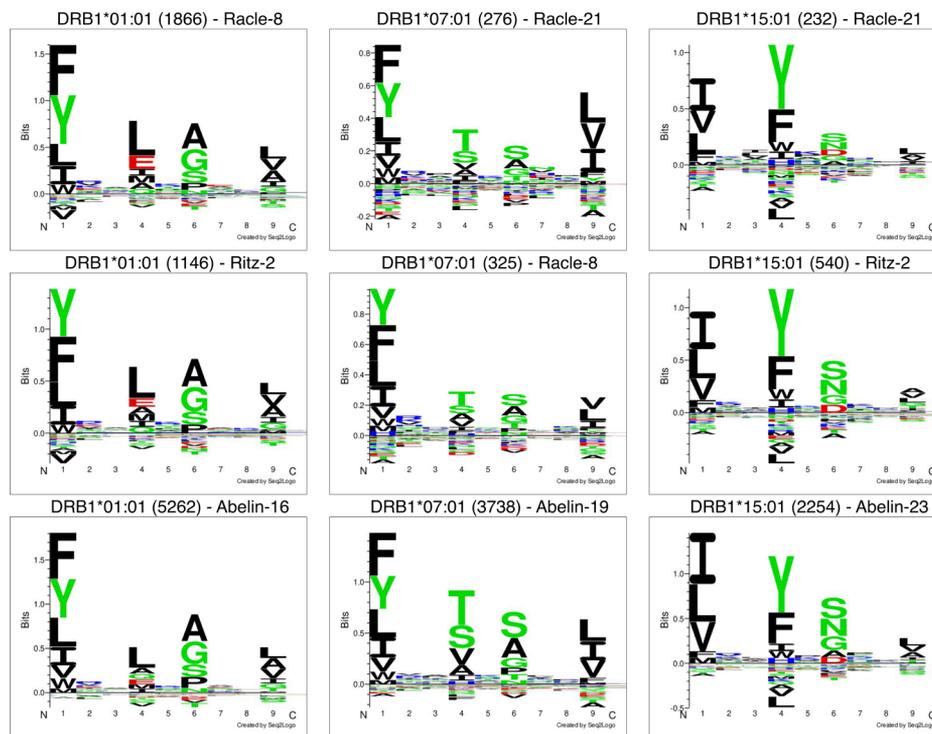


Figure 1. MA data motif deconvolution of 3 cell lines containing pairwise allele overlap compared to logos from 3 SA data sets. The figure shows the ability of NNAlign_MA to implicitly leverage allele overlap and the exclusion principle to annotate ligands from MA data and learn their binding motifs. The logos were generated with Seq2Logo, excluding ligands with a predicted rank >20%. Each logo is labeled with the allele, the number of ligands used to generate the logo and the identifier to the source data set.

terized with more than 100 ligands. For context, separate logos were made for the N- and C-terminal using weighted Kullback–Leibler, excluding sequence weighting. Logos from Gibbs Cluster deconvolution of context were generated with default settings of Seq2Logo. Note that a condition for inclusion in context motif visualization was that the PFR of the ligand is 3 residues or longer, in agreement with earlier work.¹⁵

Allele Sequence Distance Measure

The sequence distance between alleles is calculated based on the alignment score of their pseudo sequences with the following relation: $d = s(A, B) / \sqrt{s(A, A) \times s(B, B)}$, where $s(A, B)$ represents the BLOSUM50 alignment score of the two pseudo sequences.⁴⁹

Deconvolution Consistency Score

As a measure of NNAlign_MA's consistency in MHC motif deconvolution, the motifs for a given MHC as obtained from different MA data sets were compared in terms of the Pearson correlation coefficient of their position-specific scoring matrices (PSSMs). We defined the consistency score for an allele's deconvolution as the Pearson correlation coefficient averaged over all nonself comparisons.

RESULTS

The goal of this project is to build a pan-specific predictor of MHC II antigen presentation trained on EL data. Our approach is to train a model with NNAlign_MA²⁸ adapted to MHC II EL data by allowing for explicit encoding of ligand context. In large-scale benchmarks, we investigate the ability of NNAlign_MA to consistently deconvolute motifs from data sets of mixed DR, DP, and DQ peptide specificities. Furthermore, investigating the ligand context, we validate

earlier findings demonstrating an improved predictive performance for ligand discovery empowered by the signal of antigen processing and investigate this signal's consistency across MHC loci. Finally, independent CD4+ epitope benchmarks compare the developed model to recent, competing models.

NNAlign_MA Algorithm

MA data are numerous and diverse, ~82% of ligands gathered for this project stem from MA data sets (see Table S1 for a summary of EL data sets). When dealing with EL data from MHC II heterozygous cell lines, the ligands are of mixed specificities and it is not trivial to assign ligands to their respective restricting MHC II molecules. We have earlier proposed a simple machine-learning framework, NNAlign_MA, to resolve this task.²⁸ Using semi-supervised learning, NNAlign_MA leverages information from SA data to annotate MA data. This is achieved with a burn-in period in which only SA data are used for training, after which MA data are introduced, annotated, and used for training. MA data are annotated by predicting in every training iteration binding to all MHC molecules assigned to the given MA data set and assigning the restriction from the highest prediction value. With the allele assignment in place, the MA data become equivalent to SA data and can be used for training.

An essential strength of NNAlign_MA is that the MA data annotation is integral to the training process, circumventing the need for offline motif deconvolution and annotation. The fully automated process implicitly leverages data set allele overlap and exclusion principles to learn motifs and annotate ligands. An example of this functionality is shown in Figure 1, displaying the motifs identified from three MA EL data sets that pairwise share one MHC allele. This figure illustrates how

the model achieves allele annotation by comparing motifs and alleles shared between multiple data sets.

Comparison of Models Trained on SA and MA Data

NNalign_MA extends the training data by its ability to annotate MA data automatically. To assess if NNalign_MA can learn specificities uniquely covered by MA data without diluting learning from SA data, we compared three models in terms of cross-validation performance. One model, SA model, is trained exclusively on SA data (SA EL data and BA data) and another, MA model, is trained on all data (SA EL, MA EL, and BA data). The third model, NetMHCIIpan-3.2, is included as a representative model trained only on BA data. Figure 2 displays the performance of the different models on the two EL SA and MA data subsets (the performance of the SA and MA models evaluated using cross-validation).

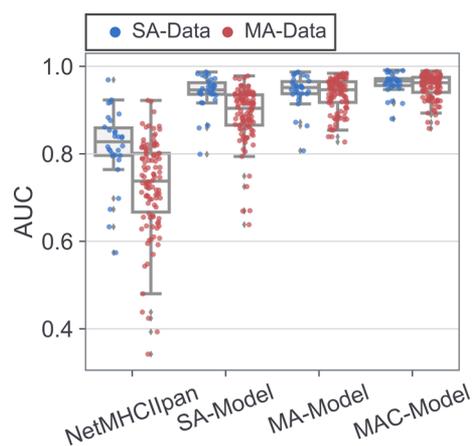


Figure 2. Cross-validation benchmark performance comparison of models with increasing complexity in terms of training data and encoding. SA model and MA model are trained on SA data and MA data, respectively. MAC model is trained on MA data, including ligand context encoding. SA, MA, and MAC models were evaluated on the complete data set by predicting each of the 5 test sets in a cross-validation setup, concatenating each test set result and calculating data set-wise AUC scores. The four models are compared separately for SA (blue dots, $n = 30$) and MA (red dots, $n = 106$) EL data sets. Every point in the figure represents an AUC value for a single ligand data set. Results of binomial statistical tests, excluding ties: SA data, comparing MA model and SA model, p -value < 0.01 . MA data, comparing MA model and SA model, p -value $\ll 0.001$. All data comparing MA model to NetMHCIIpan-3.2, p -value $\ll 0.001$. All data, comparing MA model and MAC model, p -value $\ll 0.001$.

These results demonstrate that the MA model significantly outperforms the SA model when evaluated on SA data ($p < 0.01$, binomial test excluding ties), supporting that NNalign_MA improves performance on SA data when including MA data in training. Moreover, the MA model significantly outperformed the SA model when evaluated on MA data (p -value $< 10^{-27}$, binomial test excluding ties). This gain in performance of the MA over the SA model correlated strongly with the average allele distance of the MA data sets to the alleles covered by SA data (PCC: 0.571, Figure S2). For details on this distance measure, refer to Materials and Methods. That is, the performance gain of the MA model is—as expected—most pronounced for MA data sets with alleles distant to the SA data. Finally, are both the SA and the MA models demonstrated to significantly outperform NetMHCpan-3.2 on both EL data sets ($p \ll 0.001$ in all cases). Taken together,

these results demonstrate that NNalign_MA can successfully deconvolute motifs from MA data and use this to boost the predictive power beyond that of methods trained on SA data. Likewise, these results confirm the earlier finding that methods trained on EL data outperform methods trained on BA data for the task of predicting MHC ELs.¹⁵

Source Protein Context Boosts Ligand Prediction

Immunopeptidome data inherently contain signals from steps leading to antigen presentation, and patterns of proteolytic digestion have been described in the terminal and context regions of ligands.¹⁵ Earlier work has suggested that models encoding this context information have superior performance when predicting ligand data.^{15–17} Here, we set out to validate this observation on an extensive data set.

The encoding of context data was performed as previously described¹⁵ (for details refer to Materials and Methods), and models with and without encoding of ligand context were trained and evaluated using cross-validation. The results of the benchmark are shown in Figure 2 and demonstrate a significantly improved performance of the model, MAC model, trained including the ligand context ($p \ll 0.001$ for both SA and MA data, binomial test excluding ties). This result thus confirms the earlier finding conducted on a data set covering only a limited set of HLA-DR alleles.¹⁵

Consistent Motif Deconvolution from MA Data

The posterior analyses demonstrated the high predictive power of NNalign_MA and suggested that accurate deconvolution of binding motifs in MA data is driving this performance. To further support this claim, we show in Figure S1, binding motifs generated from the predicted binding cores of ligands in each MA data set when deconvoluted by NNalign_MA. By visual inspection, this figure confirmed that for a vast majority of cases, NNalign_MA was able to successfully and consistently deconvolute motifs from MA data. For most data sets, motifs captured sharp information enrichment at well-defined anchor positions (P1, P4, P6, and P9 in most molecules).

A more quantitative analysis of deconvolution consistency was achieved by calculating pairwise correlations between PSSMs representing the motifs for MHC molecules shared between multiple data sets. That is, for every MHC molecule shared between 2 or more data sets, PSSMs were generated from the binding cores of ligands assigned to said allele for each data set and the PSSMs pairwise compared by Pearson correlation. Figure S3 visualizes this analysis in the form of heat maps. From these heat maps, a consistency score was calculated for a given allele from the average of these inter-data set correlation values. The result of this (Figure S4) demonstrated very high consistency values for most DR, DP, and H-2 alleles [median for each locus over 0.899 and overall lower values for DQ (median 0.841)].

To further quantify the accuracy of individual motif deconvolutions obtained by NNalign_MA, the positive predictive value (PPV) was calculated from the cross-validation evaluations. In short, using the cross-validated evaluations, all EL data (positive and random negative) for a given cell line are assigned to a given MHC molecule. Next, the PPV is calculated for each MHC in a given cell line from the proportion of true positives within the top N predictions in the context of the annotated natural negative peptides. To account for contaminants and other potential noise sources in the EL data, N was defined as the number of ligands assigned to a

given molecule with a rank score below 20 (discarding, on average, 7.7% of positive ligands). After filtering, the average ligand/negative ratio for each MHC molecule was 0.119, which is the expected PPV of a random predictor. The result of this analysis is presented in Table S2 and Figure S5 and overall confirmed a high performance with median PPV values for DR, DP, and H-2 alleles over 0.738. However, also here, the performance for DQ was found to be reduced (median PPV = 0.558).

Combined, these results support the claim that NNAlign_MA can accurately and consistently deconvolute binding motifs in EL MA data. We further quantify in Figure 3 the

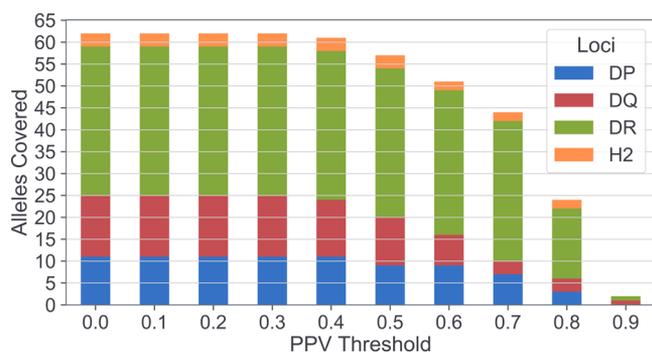


Figure 3. Number of MHC II molecules covered at different thresholds of PPV. Each molecule is represented by the average PPV, across data sets.

number of MHC-II molecules identified with a PPV above the conservative value of 0.5 (50% of the predicted positives are ligands, compared to the performance of 0.119 of a random predictor). Here, NNAlign_MA could identify accurate motifs for 34 HLA-DR, 9 HLA-DP, 11 HLA-DQ, and 3 H-2 molecules (see Figure S6 for an overview of the motifs for

these molecules identified between different data sets). In particular, for DR, these numbers are impressive, with a coverage of 100% of the alleles in the training data, and these numbers largely surpass the set of molecules described in the recent Racle¹⁷ and Abelin²⁵ publications, suggesting a superior capacity for motif deconvolution of NNAlign_MA compared to the MoDec method proposed by Racle et al., in particular, for molecules characterized by limited ligand annotations in individual MA data sets. The number of MHC molecules covered by accurate motif characterization, however, is in general lower for DQ (and DP) compared to DR reflecting the overall low number of ligands annotated to MHCs from these loci (on average less than 24% of ligands from pan-class II MA data sets are annotated to DP and DQ molecules) and the associated challenge this imposes when identifying accurate binding motifs.

Signal of Proteolytic Antigen Processing

We have previously demonstrated how the incorporation of ligand source protein context boosts the predictive power of NNAlign_MA (Figure 2). To further investigate the source of this and seek to relate the findings to specific proteolytic signals of antigen processing, we in Figure 4 show sequence logos representing the N and C terminal context signal contained within the EL data set. Here, all ligands were mapped to their antigen source protein to extract the ligand context of 6 residues (3 upstream and 3 downstream of the ligand). Along with the context, 3 residues at each terminus were extracted. A few observations can be made from this figure: first, the ligand contains more signal than the context (Figure 4a). Second, the data show a pronounced enrichment for proline in positions N + 2 and C - 2 in agreement with earlier findings.¹⁵ Finally, both termini are enriched for charged amino acids. An unsupervised deconvolution of the context motif (using GibbsCluster²⁶) supports the two latter observations (Figure 4b). The enrichment for P at positions

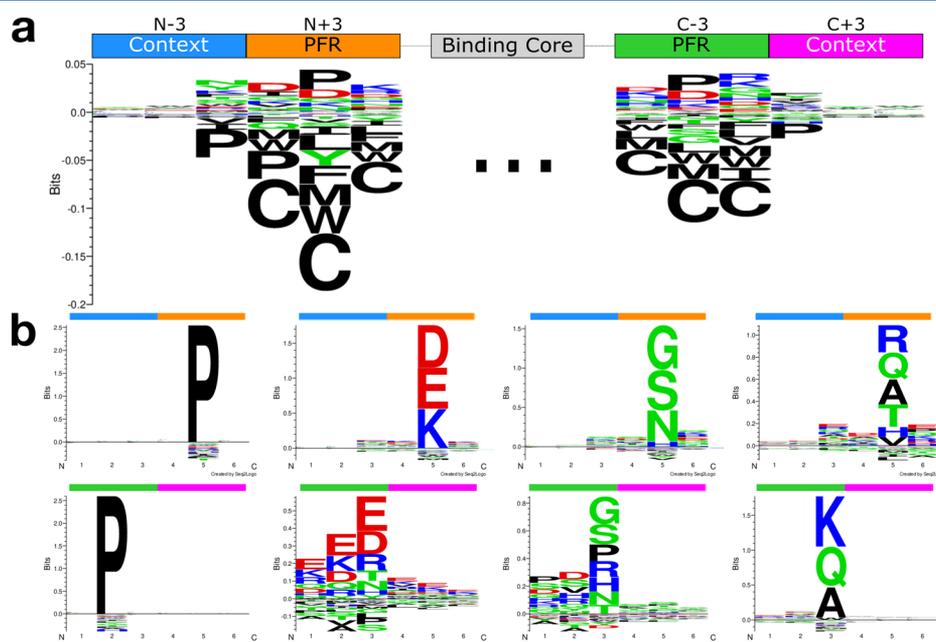


Figure 4. Sequence logos for N (top left) and C terminal (top right) and context sequences. (a) Logos for the N and C terminus of all ligands (with PFRs of length >2). The diagram indicates the placement of residues relative to the N and C terminus of ligands. (b) Upper and lower panel show unsupervised GibbsCluster deconvolution of N and C terminal contexts, respectively. Cluster motifs were ordered manually to align similar motifs for the N and C terminus, with the proline motif having the highest information content in both cases.

N + 2 and C - 2 is accentuated in the deconvolution, along with other well-defined motifs. This supports the notion that proteolytic enzymes of more than one specificity are at play in the degradation of MHC II antigens, in agreement with earlier findings.^{16,17} Furthermore, evaluating the predictive power for ligands assigned to each of the different clusters revealed a consistent high predictive performance across all clusters (median values above 0.975 in all cases) and a slightly improved performance for the clusters with P at position N + 2 and C - 2 both for N and C termini (AUCs computed data set-wise for each cluster, with negatives assigned to clusters proportional to ligand cluster assignment). Analyzing the data separately for DP, DQ, DR, and H2 molecules revealed very similar results (see Figures S7 and S8), again suggesting that the observed signal and identified motifs are related to antigen processing and not MHC binding.

CD4+ Epitope Evaluation

Prediction of CD4+ epitopes is the ultimate benchmark of peptide–MHC II binding predictors. Earlier work has suggested that prediction methods trained including EL data share improved predictive performance for this task compared to methods trained on BA data but also that the gained performance for prediction of MS ligands observed when including context information only to a minor degree—if at all—is transferred to epitope identification.^{15,16} Here, we set out to test if the models developed here align with these findings. That is, we benchmarked the predictive performance of the MA model on a large set of epitopes from the IEDB and compared performance to the state of the art methods NetMHCIIpan-3.2²⁹ (trained on BA data only), MHCnuggets⁵⁰ (trained on BA and EL data), and MixMHCIIpred¹⁷ (trained on EL data). Note that other methods for predicting MHC II binding trained on EL data have been proposed, including NeonMHC2²⁵ and MARIA.⁵¹ However, the NeonMHC2 only allows one to run max 20 predictions per day, and MARIA allows only 5000 predictions per submission, making it impractical to include in a benchmark covering more than 700,000 peptide–MHC combinations. Benchmark performance was evaluated using the *F*-rank score. In short, *F*-rank is the proportion of false-positive predictions within a given epitope source protein, that is, the percentage of peptides with a prediction scores higher than that of the epitope. An *F*-rank value of 0 corresponds to a perfect prediction (the known epitope is identified with the highest predicted binding value among all peptides found within the source protein) and a value of 50.0 to random prediction. Figure 5 gives the results of this experiment showing *F*-rank values obtained by in silico digesting the epitope source protein into overlapping peptides of the length of the epitopes. The benchmark is split into two subsets, one with epitopes restricted by MHCs covered by all methods (901 epitopes) and one with epitopes restricted by MHC-II molecules not covered by all methods except MixMHC2pred (76 epitopes covering 12 MHC II molecules). This benchmark demonstrated a significantly improved performance of the MA model compared to NetMHCIIpan-3.2, MixMHC2pred, MHCnuggets, and the MAC Model ($p < 0.012$, $p \ll 0.0001$, $p \ll 0.0001$, and $p \ll 0.0001$, respectively, binomial test excluding ties). Inspecting further epitopes of alleles covered by MixMHC2pred showed that NAlign_MA's gain in performance was not only driven by epitopes of a limited set of restrictions but generally hold across the vast majority of HLA molecules included in the benchmark

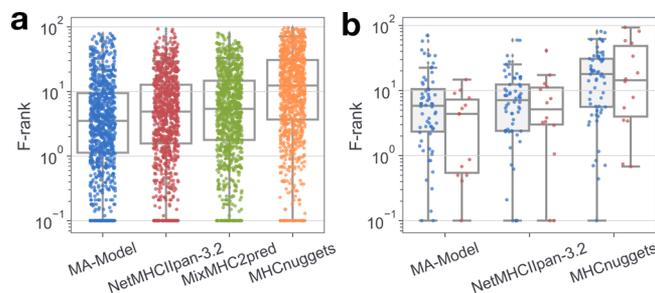


Figure 5. *F*-rank evaluation of NAlign_MA, NetMHCIIpan-3.2, and MixMHC2pred models on ICS + tetramer epitopes from the IEDB. (a) MA model (blue) *F*-rank is compared to NetMHCIIpan-3.2 (red), MixMHC2pred (green), and MHCnuggets (orange) for epitopes restricted to molecules covered by the MixMHC2pred method (901 epitopes). (b) Epitope *F*-rank of MA model, NetMHCIIpan-3.2, and MHCnuggets is compared for epitopes restricted to molecules covered by the MA EL training data (blue) (62 epitopes) and epitopes restricted to molecules outside the training data (red) (14 epitopes). (a,b) Each point in the plot represents the *F*-rank of one epitope, *F*-rank being the percentage of false-positive length-matched peptides from the epitope source protein, as described in Materials and Methods. Only results for epitopes for which at least one prediction method had *F*-rank below 10.0 are presented. For visualization, *F*-rank values of 0 are presented as 0.1005.

(improved median *F*-rank for 18 out of 20 restrictions, p -value < 0.005 in a binomial test, Figure S9). Likewise, the MA model showed improved performance compared to NetMHCIIpan-3.2 and MHCnuggets for the subset of epitopes with MHC restriction covered by the training data (this result was only significant in the case of MHCnuggets, p -value = 0.03, binomial test excluding ties). This gain in performance over MHCnuggets was further maintained ($p = 0.012$, binomial test excluding ties) when analyzing the subset of epitope restricted to MHC molecules not included in the training data (NAlign and NetMHCIIpan-3.2 share comparable predictive performance on this data set). These “novel” MHC molecules all share a close distance to the training data (less than 0.13 in all cases), suggesting that the NAlign_MA modeling framework has pan-specific prediction power at least within this distance space.

The epitope benchmark further confirmed the earlier finding that signals of antigen processing contained within the ligand context did not benefit the predictive power for epitope discovery (Figure S10). As discussed earlier,^{15,16} this observation might not come as a surprise because peptides tested for T cell immunogenicity most commonly are generated as overlapping peptides spanning a source protein and hence are not expected to follow any rules of antigen processing. However, even attempting to account for this bias applying a scoring scheme where the score of a given peptide was assigned from the sum of the individual prediction values from all 13–17-mer peptides with the predicted binding core overlapping the given peptide (for details, see Materials and Methods) maintained the improved predictive performance of the MA model and hence did not alter this conclusion.

These findings thus align with earlier results^{15,16} demonstrating an improved performance for CD4+ epitope identification of models trained including EL data compared to methods trained on BA data only (NetMHCIIpan-3.2) and that adding context information to such predictors does not

boost performance for prediction of epitopes despite this being observed for prediction of MHC ligands.

Neopeptide Evaluation

A final benchmark compared the performance of NNAlign_MA and MixMHC2pred for prediction of CD4+ neopeptides. A recent study¹⁷ compiled CD4+ neopeptide immunogenicity measures from patients with full DP, DQ, and DR typing, including both positive and negative results. Here, we present a comparison of prediction performance on this epitope set in terms of AUC and AUC0.1 (Figure 6). Excluding peptides

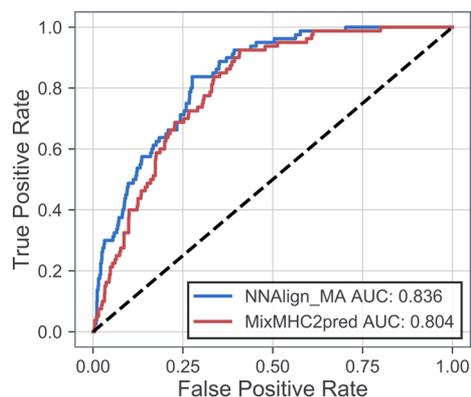


Figure 6. Neopeptide benchmark comparing NNAlign_MA and MixMHC2pred. The receiver operator curve for the prediction of a CD4+ neopeptide benchmark data set ($n = 928$).

with a length of less than 13 amino acids resulted in a benchmark of 928 peptides. For predictions, to make a fair comparison between the two methods, only HLA molecules covered by MixMHC2pred were included. MixMHC2pred prediction was performed as described in ref 17 by assigning a prediction value to each peptide as the high prediction score of all overlapping 15-mers to all HLA molecules of the patient covered by MixMHC2pred. For NNAlign_MA, the sub-k-mer approach described earlier¹⁶ was applied. Here predictions were made for all 13–21-mers subsequences contained within a given peptide to all HLA molecules of the patient covered by MixMHC2pred, with HLA annotation based on the rank. A score for each peptide was assigned as the average k-mer prediction score. The result of this analysis demonstrated a higher AUC of NNAlign_MA compared to MixMHC2pred (Figure 6). Furthermore, focusing on the earlier part of the ROC curve, the difference becomes more substantial with AUC0.1 values of 0.284 (NNAlign_MA) and 0.182 (MixMHC2pred), suggesting that NNAlign_MA reliably can identify a more substantial proportion of the neopeptides. By way of example, the two methods identify 49% (NNAlign_MA) and 37% (MixMHC2pred) of the neopeptides at a false positive rate of 10%.

NetMHCIIpan-4.0

A webserver implementation of the NNAlign_MA and NNAlign_MAC models is available at <http://www.cbs.dtu.dk/services/NetMHCIIpan-4.0>.

DISCUSSION

In this work, we have trained and evaluated a pan-specific MHC II binding predictor. This predictor integrates recent advances toward full utilization of EL data for prediction of MHC antigen presentation: two output neuron architectures

allowing the integration of multiple data types, the encoding of proteolytic ligand context, and use of the NNAlign_MA machine-learning framework allowing training on MA EL data. Training on an extensive EL data set, we have in a stepwise manner shown how these features come together to construct a prediction model that significantly improves upon current state-of-the-art methods, including its predecessor NetMHCIIpan-3.2 when it comes to predicting both ligands and CD4+ epitopes.

First, we provided examples of how NNAlign_MA consistently deconvolutes MHC binding motifs from separate MA data sets confirming that NNAlign_MA provides accurate solutions to the challenge of annotating binding motifs and learning from MA data.

Second, MA models outperformed both SA models and NetMHCIIpan in cross-validation evaluation of MA data and SA data, supporting the claim that MA data did not dilute the learning and performance of NNAlign_MA. Relatedly, the largest gain of MA over SA models was for data sets that contained alleles distant to alleles covered by SA data. This result underlines how NNAlign_MA successfully increases the coverage of predictable alleles by deconvoluting binding motifs in and training on MA data. This observation is critical for our goal of training a pan-specific MHC II predictor on EL data, as the success of pan-specific predictors is contingent on training data with broad allelic coverage.^{21,52}

Third, and in line with earlier work,^{15,17} encoding proteolytic context in models improved performance for prediction of ligands significantly. We found that the proteolytic context of ligands is a general feature that holds across MHC II loci and species. Clustering context information resulted in motifs with a strong signal of amino acid enrichment/depletion in the peptide-flanking regions. These motifs are likely a signature of proteolytic cleavage processes taking place in the MHC II antigen presentation pathway.¹⁶ Across loci and species, we observed a motif of proline enrichment in the N + 2 and C – 2 positions of ligands. This is consistent with the hypothesis of blocked endopeptidase digestion by the secondary α -amine of proline. This proline enrichment has been associated with peptides processed in the endo-lysosomal pathway of MHC II antigen presentation.⁵³ Likewise, we consistently observed motifs with enrichment for aspartic and glutamic acid. We suggest that this enrichment reflects cathepsin disfavoring negatively charged amino acids near cleavage sites,⁵⁴ leading to decreased digestion of ligands with said signature. With lesser information content, a motif of the N-terminal context deconvolution displayed enrichment of glycine, serine, and threonine. A similar enrichment has been observed downstream of the cleavage site for cathepsin S, a protease known to participate in the MHC II antigen presentation pathway.⁵⁵ The identified processing motifs are thus to a high degree associated with decreased potential for antigen processing. Given this, we speculate that they could serve as novel MHC II agnostic ways for engineering enhanced/reduced peptide and protein immunogenicity; that is, amino acid mutations away from these motifs would lead to enhanced potential for cleavage, leading to a lower level of antigen presentation and hence a reduction in potential immunogenicity. Likewise, could mutations into these motifs lower the potential for antigen processing, potentially inducing increased immunogenicity?

The proposed model improved CD4+ epitope prediction when compared to the state-of-the-art predictors NetMHCII-

pan-3.2, MHCnuggets, and MixMHC2pred. In contrast to the findings for MS ligands, in this benchmark, the predictive power of the models did not benefit from the integration of context encoding. This was observed even when accounting for bias in the data imposed by T cell epitope peptides most often being generated synthetically and thus not reflecting rules of natural antigen processing. This result is in line with earlier observations^{15,16} and suggests that the cleavage signal, even though present in individual ligands, are diluted when analyzing T cell epitopes because of the vast volume of possible ligands with binding cores overlapping the epitope.

Finally, the performance of NNAlign_MA for the identification of CD4+ cancer neoepitopes was benchmarked against MixMHC2pred. Here, NNAlign_MA demonstrated an overall improved predictive performance compared to MixMHC2pred, identifying close to half of the neoepitopes at a false-positive rate of 10%. From a clinical and applied perspective, this result suggests NNAlign_MA (and the resulting predictor NetMHCIIpan-4.0) to be a useful tool as a guide also for identification of neoepitopes.

A key factor determining the power of a predictive model for MHC antigen presentation is the volume and MHC diversity of the training data. Here, we have demonstrated how the ability of NNAlign_MA to deconvolute motifs and integrate MA EL data into the training framework allows for an effective increase in the MHC diversity of the training data. For example, the training data used in the current study are limited to 19 (14 DR, 2 DP, 2 DQ, and 1 H2) molecules if focusing only on SA EL data. By integrating MA data, this number is expanded to 116 (34 DR, 25 DP, 54 DQ, and 3 H2) (accounting for all possible α and β chain combinations in heterozygous DQ and DP data sets). However, this HLA diversity is only relevant if associated with large volumes of peptide data. Analyzing the results of the motif deconvolution of NNAlign_MA on the training data for data sets revealed a high bias in the EL data with very limited ligands annotated to DQ (and to a lesser degree DP) molecules. This bias is more substantial than what was expected from the difference in the relative expression of MHC II between the different loci and is more likely caused by differences in the specificities of the antibodies used for immunoprecipitation (IP) when purifying DR versus DQ/DP MHC complexes before running the MS. Recently, Abelin et al.²⁵ have proposed a framework to experimentally generate SA MHC MS EL data sets to resolve this issue (and the general issues of handling MA data sets) and demonstrated how this approach could be used to generate high volumes of SA EL data also for DQ and DP molecules. Alternative approaches would include working with antibodies with improved DQ (and DP) specificities and with these conduct sequential IPs allowing for identification of larger volumes of ligands for DQ and DP.

Another complicating factor for analyzing and interpreting MA EL data for DQ and DP stems from the fact that both the α and β chains are polymorphic, making it nontrivial to assess which of the four possible combinations of alpha and two β chains are presented as HLA molecules in a given cell. However, carefully selecting homozygous cell lines or cell lines with only one expressed α chain could help resolve this, and in this context, EL data covering DQ and DP can without any further complication be integrated into the NNAlign_MA modeling framework.

A central power of the NNAlign_MA modeling framework is its ability to leverage information between data sets, boosting

performance for molecules with limited ligand coverage combined with its pan-specificity allowing us to make accurate predictions outside the set of molecules included in the training data. The first part was illustrated above with the expanded number of molecules characterized by accurate motifs in comparison to the method by Racle et al. Although the pan-specific prediction power of the original NNAlign model has been documented previously using different kinds of leave-one-allele-out experiments,^{18,21} this is only the case to a minimal extent for NNAlign_MA.²⁸ The reason for this is that leave-one-out experiments cannot readily be conducted on MA data sets because molecules in this situation, in general, are shared between multiple data sets making it nontrivial to define what to leave-out. Given this, we here limited the leave-one-allele-out experiments to a set of CD4+ epitopes restricted to HLA molecules not included in the training data (Figure 5b). The predictive performance of NNAlign_MA was maintained for this set of epitopes suggesting that at least within the limits of this benchmark, NNAlign_MA does achieve pan-specific prediction power, allowing coverage of MHC molecules not included in the training data.

As illustrated here, the power allows the method to leverage information between data sets enabling the method to identify accurate binding motifs also for molecules characterized by limited ligand counts in individual data sets. This ability allowed us to expand the coverage of MHC molecules in the training data with accurate binding motifs compared to that obtained with dataset-specific motif deconvolution methods such as MoDec. Beyond identifying motifs for a higher number of MHC molecules, this pan-specific power has large impacts on the application of the prediction models developed. Here, the allele-specific models, such as MixMHC2pred developed using the MoDec deconvolution framework, has predictive power limited to the few HLA molecules with accurate motif deconvolution covered in the training data.

In contrast, the pan-specific method proposed here allows accurate predictions also outside this limited space of HLA molecules. For example, MixMHC2pred covers 33 HLA II molecules corresponding to cases where accurate motifs could be identified from the training data. By use of NNAlign_MA, the number of molecules with characterized motifs is increased to 57, and by expanding using the pan-specific power of the trained prediction model, the coverage is extended to 1917 (including HLA-II molecules with a pseudo-sequence distance <0.05 to any of the molecules with accurate binding motifs).

Properties other than antigen processing and HLA binding, such as protein expression and differential access to the MHC-II presentation pathway, contribute to the likelihood of antigen presentation. Recent papers (including Abelin et al.²⁵ and Chen et al.⁵¹) have suggested modeling frameworks integrating a panel of such properties, suggesting considerable improvement for the prediction of HLA ligandomes. Further work remains to be performed to validate the generality of these findings and their potential impact on general rational epitope discovery.

In conclusion, we have shown how the relatively simple NNAlign_MA machine-learning framework can be applied to deconvolute binding motifs in MHC class II MA EL data sets and how this deconvolution allows for identification of signals associated with antigen processing and construction of a pan-specific prediction model with significantly improved performance compared to state of the art for prediction of both ELs and CD4+ T cell epitopes.

■ ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jproteome.9b00874>.

Predicted binding core motif logos deconvoluted from all MA data sets; evidence of the benefit of extending training data scope of MA data; quantitative analysis of MA data deconvolution consistency; Kernel density estimate plot of consistency scores, colored by loci; PPV distributions of MHC deconvolutions; MHC II logos of alleles confidently learned by the trained model; context logos from different MHC loci: DP, DQ, DR, and H2; unsupervised deconvolution of context motifs via Gibbs clustering; *F*-rank results of ICS + tetramer benchmark split up by MHC restriction; *F*-rank results of the ICS + tetramer benchmark for MA and MAC models with *k*-mer evaluation; metadata for EL data sets; and PPV of MHC deconvolutions for cell lines (PDF)

■ AUTHOR INFORMATION

Corresponding Author

Morten Nielsen – Department of Health Technology, Technical University of Denmark, Lyngby 2800, Denmark; Instituto de Investigaciones Biotecnológicas, Universidad Nacional de San Martín, CP1650, Argentina; orcid.org/0000-0001-7885-4311; Email: morni@dtu.dk

Authors

Birkir Reynisson – Department of Health Technology, Technical University of Denmark, Lyngby 2800, Denmark; orcid.org/0000-0002-7016-4938

Carolina Barra – Department of Health Technology, Technical University of Denmark, Lyngby 2800, Denmark

Saghar Kaabinejadian – Pure MHC LLC, Oklahoma City, Oklahoma 73104, United States

William H. Hildebrand – Department of Microbiology and Immunology, University of Oklahoma Health Sciences Center, Oklahoma City, Oklahoma 73104, United States

Bjoern Peters – Division of Vaccine Discovery, La Jolla Institute for Allergy and Immunology, La Jolla, California 92037, United States; Department of Medicine, University of California, San Diego, San Diego, California 92093, United States

Complete contact information is available at: <https://pubs.acs.org/doi/10.1021/acs.jproteome.9b00874>

Notes

The authors declare no competing financial interest.

■ REFERENCES

- (1) Unanue, E. R.; Turk, V.; Neefjes, J. Variations in MHC Class II Antigen Processing and Presentation in Health and Disease. *Annu. Rev. Immunol.* **2016**, *34*, 265–297.
- (2) Wiczorek, M.; Abualrous, E. T.; Sticht, J.; Álvaro-Benito, M.; Stolzenberg, S.; Noé, F.; Freund, C. Major Histocompatibility Complex (MHC) Class I and MHC Class II Proteins: Conformational Plasticity in Antigen Presentation. *Front. Immunol.* **2017**, *8*, 292.
- (3) Traherne, J. A. Human MHC Architecture and Evolution: Implications for Disease Association Studies. *Int. J. Immunogenet.* **2008**, *35*, 179–192.
- (4) Chicz, R. M.; Urban, R. G.; Lane, W. S.; Gorga, J. C.; Stern, L. J.; Vignali, D. A. A.; Strominger, J. L. Predominant Naturally Processed Peptides Bound to HLA-DR1 Are Derived from MHC-Related

Molecules and Are Heterogeneous in Size. *Nature* **1992**, *358*, 764–768.

- (5) Andreatta, M.; Jurtz, V. I.; Kaefer, T.; Sette, A.; Peters, B.; Nielsen, M. Machine Learning Reveals a Non-Canonical Mode of Peptide Binding to MHC Class II Molecules. *Immunology* **2017**, *152*, 255–264.

- (6) Godkin, A. J.; Smith, K. J.; Willis, A.; Tejada-Simon, M. V.; Zhang, J.; Elliott, T.; Hill, A. V. S. Naturally Processed HLA Class II Peptides Reveal Highly Conserved Immunogenic Flanking Region Sequence Preferences That Reflect Antigen Processing Rather than Peptide-MHC Interactions. *J. Immunol.* **2001**, *166*, 6720–6727.

- (7) Buus, S.; Sette, A.; Colon, S.; Miles, C.; Grey, H. The Relation between Major Histocompatibility Complex (MHC) Restriction and the Capacity of Ia to Bind Immunogenic Peptides. *Science* **1987**, *235*, 1353–1358.

- (8) Gfeller, D.; Bassani-Sternberg, M. Predicting Antigen Presentation—What Could We Learn From a Million Peptides? *Front. Immunol.* **2018**, *9*, 1716.

- (9) Nielsen, M.; Lund, O.; Buus, S.; Lundegaard, C. MHC Class II Epitope Predictive Algorithms. *Immunology* **2010**, *130*, 319–328.

- (10) Bassani-Sternberg, M.; Bräunlein, E.; Klar, R.; Engleitner, T.; Sinitcyn, P.; Audehm, S.; Straub, M.; Weber, J.; Slotta-Huspenina, J.; Specht, K.; Martignoni, M. E.; Werner, A.; Hein, R.; Busch, D. H.; Peschel, C.; Rad, R.; Cox, J.; Mann, M.; Krackhardt, A. M. Direct Identification of Clinically Relevant Neoepitopes Presented on Native Human Melanoma Tissue by Mass Spectrometry. *Nat. Commun.* **2016**, *7*, 13404.

- (11) Backert, L.; Kohlbacher, O. Immunoinformatics and Epitope Prediction in the Age of Genomic Medicine. *Genome Med.* **2015**, *7*, 119.

- (12) Alvarez, B.; Barra, C.; Nielsen, M.; Andreatta, M. Computational Tools for the Identification and Interpretation of Sequence Motifs in Immunopeptidomes. *Proteomics* **2018**, *18*, 1700252.

- (13) Dudek, N. L.; Croft, N. P.; Schittenhelm, R. B.; Ramarathnam, S. H.; Purcell, A. W. A Systems Approach to Understand Antigen Presentation and the Immune Response. *Methods Mol. Biol.* **2016**, *1394*, 189–209.

- (14) Caron, E.; Kowalewski, D. J.; Chiek Koh, C.; Sturm, T.; Schuster, H.; Aebbersold, R. Analysis of Major Histocompatibility Complex (MHC) Immunopeptidomes Using Mass Spectrometry. *Mol. Cell. Proteomics* **2015**, *14*, 3105–3117.

- (15) Barra, C.; Alvarez, B.; Paul, S.; Sette, A.; Peters, B.; Andreatta, M.; Buus, S.; Nielsen, M. Footprints of Antigen Processing Boost MHC Class II Natural Ligand Predictions. *Genome Med.* **2018**, *10*, 84.

- (16) Paul, S.; Karosiene, E.; Dhanda, S. K.; Jurtz, V.; Edwards, L.; Nielsen, M.; Sette, A.; Peters, B. Determination of a Predictive Cleavage Motif for Eluted Major Histocompatibility Complex Class II Ligands. *Front. Immunol.* **2018**, *9*, 1795.

- (17) Racle, J.; Michaux, J.; Rockinger, G. A.; Arnaud, M.; Bobisse, S.; Chong, C.; Guillaume, P.; Coukos, G.; Harari, A.; Jandus, C.; Bassani-Sternberg, M.; Gfeller, D. Robust Prediction of HLA Class II Epitopes by Deep Motif Deconvolution of Immunopeptidomes. *Nat. Biotechnol.* **2019**, *37*, 1283.

- (18) Garde, C.; Ramarathnam, S. H.; Jappe, E. C.; Nielsen, M.; Kringelum, J. V.; Trolle, T.; Purcell, A. W. Improved Peptide-MHC Class II Interaction Prediction through Integration of Eluted Ligand and Peptide Affinity Data. *Immunogenetics* **2019**, *71*, 445–454.

- (19) Jurtz, V.; Paul, S.; Andreatta, M.; Marcatili, P.; Peters, B.; Nielsen, M. NetMHCpan-4.0: Improved Peptide-MHC Class I Interaction Predictions Integrating Eluted Ligand and Peptide Binding Affinity Data. *J. Immunol.* **2017**, *199*, 3360–3368.

- (20) Nielsen, M.; Connelley, T.; Ternette, N. Improved Prediction of Bovine Leucocyte Antigens (BoLA) Presented Ligands by Use of Mass-Spectrometry-Determined Ligand and in Vitro Binding Data. *J. Proteome Res.* **2018**, *17*, 559–567.

- (21) Hoof, I.; Peters, B.; Sidney, J.; Pedersen, L. E.; Sette, A.; Lund, O.; Buus, S.; Nielsen, M. NetMHCpan, a Method for MHC Class I Binding Prediction beyond Humans. *Immunogenetics* **2009**, *61*, 1–13.

- (22) Karosiene, E.; Rasmussen, M.; Blicher, T.; Lund, O.; Buus, S.; Nielsen, M. NetMHCIIpan-3.0, a Common Pan-Specific MHC Class II Prediction Method Including All Three Human MHC Class II Isotypes, HLA-DR, HLA-DP and HLA-DQ. *Immunogenetics* **2013**, *65*, 711–724.
- (23) Nielsen, M.; Lundegaard, C.; Blicher, T.; Lamberth, K.; Harndahl, M.; Justesen, S.; Røder, G.; Peters, B.; Sette, A.; Lund, O.; Buus, S. NetMHCpan, a Method for Quantitative Predictions of Peptide Binding to Any HLA-A and -B Locus Protein of Known Sequence. *PLoS One* **2007**, *2*, No. e796.
- (24) Abelin, J. G.; Keskin, D. B.; Sarkizova, S.; Hartigan, C. R.; Zhang, W.; Sidney, J.; Stevens, J.; Lane, W.; Zhang, G. L.; Eisenhaure, T. M.; Clauser, K. R.; Hacohen, N.; Rooney, M. S.; Carr, S. A.; Wu, C. J. Mass Spectrometry Profiling of HLA-Associated Peptidomes in Mono-Allelic Cells Enables More Accurate Epitope Prediction. *Immunity* **2017**, *46*, 315–326.
- (25) Abelin, J. G.; Harjanto, D.; Malloy, M.; Suri, P.; Colson, T.; Goulding, S. P.; Creech, A. L.; Serrano, L. R.; Nasir, G.; Nasrullah, Y.; McGann, C. D.; Velez, D.; Ting, Y. S.; Poran, A.; Rothenberg, D. A.; Chhangawala, S.; Rubinsteyn, A.; Hammerbacher, J.; Gaynor, R. B.; Fritsch, E. F.; Greshock, J.; Oslund, R. C.; Barthelme, D.; Addona, T. A.; Arieta, C. M.; Rooney, M. S. Defining HLA-II Ligand Processing and Binding Rules with Mass Spectrometry Enhances Cancer Epitope Prediction. *Immunity* **2019**, *51*, 766–779.
- (26) Andreatta, M.; Alvarez, B.; Nielsen, M. GibbsCluster: Unsupervised Clustering and Alignment of Peptide Sequences. *Nucleic Acids Res.* **2017**, *45*, W458–W463.
- (27) Bassani-Sternberg, M.; Gfeller, D. Unsupervised HLA Peptidome Deconvolution Improves Ligand Prediction Accuracy and Predicts Cooperative Effects in Peptide-HLA Interactions. *J. Immunol.* **2016**, *197*, 2492–2499.
- (28) Alvarez, B.; Reynisson, B.; Barra, C.; Buus, S.; Ternette, N.; Connelley, T.; Andreatta, M.; Nielsen, M. NNAlign_MA; MHC Peptidome Deconvolution for Accurate MHC Binding Motif Characterization and Improved T Cell Epitope Predictions. *Mol. Cell. Proteomics* **2019**, *18*, 2459.
- (29) Jensen, K. K.; Andreatta, M.; Marcatili, P.; Buus, S.; Greenbaum, J. A.; Yan, Z.; Sette, A.; Peters, B.; Nielsen, M. Improved Methods for Predicting Peptide Binding Affinity to MHC Class II Molecules. *Immunology* **2018**, *154*, 394–406.
- (30) Nielsen, M.; Lundegaard, C.; Worning, P.; Lauemøller, S. L.; Lamberth, K.; Buus, S.; Brunak, S.; Lund, O. Reliable Prediction of T-Cell Epitopes Using Neural Networks with Novel Sequence Representations. *Protein Sci.* **2003**, *12*, 1007–1017.
- (31) Vita, R.; Mahajan, S.; Overton, J. A.; Dhanda, S. K.; Martini, S.; Cantrell, J. R.; Wheeler, D. K.; Sette, A.; Peters, B. The Immune Epitope Database (IEDB): 2018 Update. *Nucleic Acids Res.* **2019**, *47*, D339–D343.
- (32) Heyder, T.; Kohler, M.; Tarasova, N. K.; Haag, S.; Rutishauser, D.; Rivera, N. V.; Sandin, C.; Mia, S.; Malmström, V.; Wheelock, Å. M.; Wahlström, J.; Holmdahl, R.; Eklund, A.; Zubarev, R. A.; Grunewald, J.; Ytterberg, A. J. Approach for Identifying Human Leukocyte Antigen (HLA)-DR Bound Peptides from Scarce Clinical Samples. *Mol. Cell. Proteomics* **2016**, *15*, 3017–3029.
- (33) Bergseng, E.; Dørum, S.; Arntzen, M. Ø.; Nielsen, M.; Nygård, S.; Buus, S.; de Souza, G. A.; Sollid, L. M. Different Binding Motifs of the Celiac Disease-Associated HLA Molecules DQ2.5, DQ2.2, and DQ7.5 Revealed by Relative Quantitative Proteomics of Endogenous Peptide Repertoires. *Immunogenetics* **2015**, *67*, 73–84.
- (34) Sofron, A.; Ritz, D.; Neri, D.; Fugmann, T. High-Resolution Analysis of the Murine MHC Class II Immunopeptidome. *Eur. J. Immunol.* **2016**, *46*, 319–328.
- (35) Clement, C. C.; Becerra, A.; Yin, L.; Zolla, V.; Huang, L.; Merlin, S.; Follenzi, A.; Shaffer, S. A.; Stern, L. J.; Santambrogio, L. The Dendritic Cell Major Histocompatibility Complex II (MHC II) Peptidome Derives from a Variety of Processing Pathways and Includes Peptides with a Broad Spectrum of HLA-DM Sensitivity. *J. Biol. Chem.* **2016**, *291*, 5576–5595.
- (36) Karunakaran, K. P.; Yu, H.; Jiang, X.; Chan, Q.; Goldberg, M. F.; Jenkins, M. K.; Foster, L. J.; Brunham, R. C. Identification of MHC-Bound Peptides from Dendritic Cells Infected with Salmonella Enterica Strain SL1344: Implications for a Nontyphoidal Salmonella Vaccine. *J. Proteome Res.* **2017**, *16*, 298–306.
- (37) Ooi, J. D.; Petersen, J.; Tan, Y. H.; Huynh, M.; Willett, Z. J.; Ramarathinam, S. H.; Eggenhuizen, P. J.; Loh, K. L.; Watson, K. A.; Gan, P. Y.; Alikhan, M. A.; Dudek, N. L.; Handel, A.; Hudson, B. G.; Fugger, L.; Power, D. A.; Holt, S. G.; Coates, P. T.; Gregersen, J. W.; Purcell, A. W.; Holdsworth, S. R.; La Gruta, N. L.; Reid, H. H.; Rossjohn, J.; Kitching, A. R. Dominant Protection from HLA-Linked Autoimmunity by Antigen-Specific Regulatory T Cells. *Nature* **2017**, *545*, 243–247.
- (38) Wang, Q.; Drouin, E. E.; Yao, C.; Zhang, J.; Huang, Y.; Leon, D. R.; Steere, A. C.; Costello, C. E. Immunogenic HLA-DR-Presented Self-Peptides Identified Directly from Clinical Samples of Synovial Tissue, Synovial Fluid, or Peripheral Blood in Patients with Rheumatoid Arthritis or Lyme Arthritis. *J. Proteome Res.* **2017**, *16*, 122–136.
- (39) Nelde, A.; Kowalewski, D. J.; Backert, L.; Schuster, H.; Werner, J.-O.; Klein, R.; Kohlbacher, O.; Kanz, L.; Salih, H. R.; Rammensee, H.-G.; Stevanović, S.; Walz, J. S. HLA Ligandome Analysis of Primary Chronic Lymphocytic Leukemia (CLL) Cells under Lenalidomide Treatment Confirms the Suitability of Lenalidomide for Combination with T-Cell-Based Immunotherapy. *Oncoimmunology* **2018**, *7*, No. e1316438.
- (40) Ritz, D.; Sani, E.; Debiec, H.; Ronco, P.; Neri, D.; Fugmann, T. Membranal and Blood-Soluble HLA Class II Peptidome Analyses Using Data-Dependent and Independent Acquisition. *Proteomics* **2018**, *18*, 1700246.
- (41) Ting, Y. T.; Petersen, J.; Ramarathinam, S. H.; Scally, S. W.; Loh, K. L.; Thomas, R.; Suri, A.; Baker, D. G.; Purcell, A. W.; Reid, H. H.; Rossjohn, J. The Interplay between Citrullination and HLA-DRB1 Polymorphism in Shaping Peptide Binding Hierarchies in Rheumatoid Arthritis. *J. Biol. Chem.* **2018**, *293*, 3236–3251.
- (42) Scally, S. W.; Petersen, J.; Law, S. C.; Dudek, N. L.; Nel, H. J.; Loh, K. L.; Wijeyewickrema, L. C.; Eckle, S. B. G.; van Heemst, J.; Pike, R. N.; McCluskey, J.; Toes, R. E.; La Gruta, N. L.; Purcell, A. W.; Reid, H. H.; Thomas, R.; Rossjohn, J. A Molecular Basis for the Association of the HLA-DRB1 Locus, Citrullination, and Rheumatoid Arthritis. *J. Exp. Med.* **2013**, *210*, 2569–2582.
- (43) Khodadoust, M. S.; Olsson, N.; Wagar, L. E.; Haabeth, O. A. W.; Chen, B.; Swaminathan, K.; Rawson, K.; Liu, C. L.; Steiner, D.; Lund, P.; Rao, S.; Zhang, L.; Marceau, C.; Stehr, H.; Newman, A. M.; Czerwinski, D. K.; Carlton, V. E. H.; Moorhead, M.; Faham, M.; Kohrt, H. E.; Carette, J.; Green, M. R.; Davis, M. M.; Levy, R.; Elias, J. E.; Alizadeh, A. A. Antigen Presentation Profiling Reveals Recognition of Lymphoma Immunoglobulin Neoantigens. *Nature* **2017**, *543*, 723–727.
- (44) Álvaro-Benito, M.; Morrison, E.; Abualrous, E. T.; Kuroopka, B.; Freund, C. Quantification of HLA-DM-Dependent Major Histocompatibility Complex of Class II Immunopeptidomes by the Peptide Landscape Antigenic Epitope Alignment Utility. *Front. Immunol.* **2018**, *9*, 872.
- (45) Nanaware, P. P.; Jurewicz, M. M.; Leszyk, J. D.; Shaffer, S. A.; Stern, L. J. HLA-DO Modulates the Diversity of the MHC-II Self-Peptidome. *Mol. Cell. Proteomics* **2019**, *18*, 490–503.
- (46) UniProt Consortium. UniProt: A Worldwide Hub of Protein Knowledge. *Nucleic Acids Res.* **2019**, *47*, D506–D515.
- (47) Nielsen, M.; Lundegaard, C.; Lund, O. Prediction of MHC Class II Binding Affinity Using SMM-Align, a Novel Stabilization Matrix Alignment Method. *BMC Bioinf.* **2007**, *8*, 238.
- (48) Thomsen, M. C. F.; Nielsen, M. Seq2Logo: A Method for Construction and Visualization of Amino Acid Binding Motifs and Sequence Profiles Including Sequence Weighting, Pseudo Counts and Two-Sided Representation of Amino Acid Enrichment and Depletion. *Nucleic Acids Res.* **2012**, *40*, W281–W287.
- (49) Nielsen, M.; Lundegaard, C.; Blicher, T.; Peters, B.; Sette, A.; Justesen, S.; Buus, S.; Lund, O. Quantitative Predictions of Peptide

Binding to Any HLA-DR Molecule of Known Sequence: NetMHCIIpan. *PLoS Comput. Biol.* **2008**, *4*, No. e1000107.

(50) Shao, X. M.; Bhattacharya, R.; Huang, J.; Sivakumar, I. K. A.; Tokheim, C.; Zheng, L.; Hirsch, D.; Kaminow, B.; Omdahl, A.; Bonsack, M.; Riemer, A. B.; Velculescu, V. E.; Anagnostou, V.; Pagel, K. A.; Karchin, R. High-Throughput Prediction of MHC Class I and II Neoantigens with MHCnuggets. *Cancer Immunol. Res.* **2020**, *8*, 396.

(51) Chen, B.; Khodadoust, M. S.; Olsson, N.; Wagar, L. E.; Fast, E.; Liu, C. L.; Muftuoglu, Y.; Swarder, B. J.; Diehn, M.; Levy, R.; Davis, M. M.; Elias, J. E.; Altman, R. B.; Alizadeh, A. A. Predicting HLA Class II Antigen Presentation through Integrated Deep Learning. *Nat. Biotechnol.* **2019**, *37*, 1332–1343.

(52) Karosiene, E.; Lundegaard, C.; Lund, O.; Nielsen, M. NetMHCcons: A Consensus Method for the Major Histocompatibility Complex Class I Predictions. *Immunogenetics* **2012**, *64*, 177–186.

(53) Ciudad, M. T.; Sorvillo, N.; van Alphen, F. P.; Catalán, D.; Meijer, A. B.; Voorberg, J.; Jaraquemada, D. Analysis of the HLA-DR Peptidome from Human Dendritic Cells Reveals High Affinity Repertoires and Nonconventional Pathways of Peptide Generation. *J. Leukocyte Biol.* **2017**, *101*, 15–27.

(54) Choe, Y.; Leonetti, F.; Greenbaum, D. C.; Lecaille, F.; Bogyo, M.; Brömme, D.; Ellman, J. A.; Craik, C. S. Substrate Profiling of Cysteine Proteases Using a Combinatorial Peptide Library Identifies Functionally Unique Specificities. *J. Biol. Chem.* **2006**, *281*, 12824–12832.

(55) Biniossek, M. L.; Nägler, D. K.; Becker-Pauly, C.; Schilling, O. Proteomic Identification of Protease Cleavage Sites Characterizes Prime and Non-Prime Specificity of Cysteine Cathepsins B, L, and S. *J. Proteome Res.* **2011**, *10*, 5363–5373.