

THE POWER TO SEE



CytoFLEX FLOW
CYTOMETER

Advanced Sensitivity
and Resolution



Pan-Specific Prediction of Peptide–MHC Class I Complex Stability, a Correlate of T Cell Immunogenicity

This information is current as of July 10, 2016.

Michael Rasmussen, Emilio Fenoy, Mikkel Harndahl, Anne Bregnballe Kristensen, Ida Kallehauge Nielsen, Morten Nielsen and Søren Buus

J Immunol published online 8 July 2016

<http://www.jimmunol.org/content/early/2016/07/08/jimmunol.1600582>

-
- Supplementary Material** <http://www.jimmunol.org/content/suppl/2016/07/08/jimmunol.1600582.DCSupplemental.html>
- Subscriptions** Information about subscribing to *The Journal of Immunology* is online at: <http://jimmunol.org/subscriptions>
- Permissions** Submit copyright permission requests at: <http://www.aai.org/ji/copyright.html>
- Email Alerts** Receive free email-alerts when new articles cite this article. Sign up at: <http://jimmunol.org/cgi/alerts/etoc>

The Journal of Immunology is published twice each month by
The American Association of Immunologists, Inc.,
9650 Rockville Pike, Bethesda, MD 20814-3994.
Copyright © 2016 by The American Association of
Immunologists, Inc. All rights reserved.
Print ISSN: 0022-1767 Online ISSN: 1550-6606.



Pan-Specific Prediction of Peptide–MHC Class I Complex Stability, a Correlate of T Cell Immunogenicity

Michael Rasmussen,^{*,1} Emilio Fenoy,^{†,1} Mikkel Harndahl,^{*} Anne Bregnballe Kristensen,^{*} Ida Kallehauge Nielsen,^{*} Morten Nielsen,^{‡,‡} and Søren Buus^{*}

Binding of peptides to MHC class I (MHC-I) molecules is the most selective event in the processing and presentation of Ags to CTL, and insights into the mechanisms that govern peptide–MHC-I binding should facilitate our understanding of CTL biology. Peptide–MHC-I interactions have traditionally been quantified by the strength of the interaction, that is, the binding affinity, yet it has been shown that the stability of the peptide–MHC-I complex is a better correlate of immunogenicity compared with binding affinity. In this study, we have experimentally analyzed peptide–MHC-I complex stability of a large panel of human MHC-I allotypes and generated a body of data sufficient to develop a neural network–based pan-specific predictor of peptide–MHC-I complex stability. Integrating the neural network predictors of peptide–MHC-I complex stability with state-of-the-art predictors of peptide–MHC-I binding is shown to significantly improve the prediction of CTL epitopes. The method is publicly available at <http://www.cbs.dtu.dk/services/NetMHCstabpan>. *The Journal of Immunology*, 2016, 197: 000–000.

Binding and presentation of antigenic peptides by MHC molecules is a central event in the initiation of an adaptive immune response. MHC class I (MHC-I) molecules bind peptides derived from the intracellular protein metabolism and present them at the cell surface for scrutiny by effector cells of the immune system. The resulting presentation of peptides of pathogenic origin (e.g., from virus or tumors) can trigger the activation of immune effector cells such as CD8⁺ CTLs. Peptide binding to MHC-I molecules is the single most selective of the many events leading to Ag presentation, and considerable efforts have been dedicated to unravel the rules of this event. For a peptide to induce an immune response, it must bind with sufficient affinity and stability to the restricting MHC-I element to allow translocation of intact peptide–MHC-I complexes (pMHC) to the cell surface where it should remain associated long enough for circulating CD8⁺ CTLs of appropriate specificity to arrive and recognize it. Although the law of mass action establishes a very close and direct relationship between affinity and stability, the earlier logic would

suggest that stability might be the better immunogenicity correlate of the two parameters. As previously noted by us (1), several others have suggested that the stability of a pMHC complex correlates with immunogenicity (both for MHC-I [2–8] and for MHC-II [9, 10]), and it has even been suggested that stability correlates better with immunogenicity than affinity does (both for MHC-I [11–14] and MHC-II [15]). Unfortunately, these claims have been supported by limited and to some extent biased experimental data. The former is due to the experimental challenges associated with measuring complex stability, which effectively has meant that affinity has remained the most frequently established correlate of immunogenicity. The latter is due to the frequent use of predicting algorithms to select peptides of interest, which effectively precludes the use of immunogenicity data obtained from such predictions to compare the impact of affinity versus stability.

Using an HLA-A*02:01-transgenic mouse model and examining antivaccinia immune responses in an unbiased and comprehensive manner, Assarsson et al. (16) showed that as much as 90% of the peptides that were found to bind with high affinity to HLA-A*02:01 nonetheless were not recognized postinfection. Asking whether lack of stability could be part of the explanation of this high rate of false discovery, Assarsson et al. (16) examined a small selection of 12 of these peptides spanning dominant, subdominant, cryptic, and nonimmunogenic events and concluded that “*a suggestive, but not statistically significant, trend was noted for off-rates and dominance.*” In a follow-up study (1), we examined the stability of a larger collection of the peptide–HLA-A*02:01 interactions reported by Assarsson et al. (16), and in increasing the sample size we indeed found that stability is a significantly better correlate of immunogenicity than is affinity. We suggested that 30% of the peptides, which are nonimmunogenic even though they experimentally can be verified as binders to HLA-A*02:01, lack immunogenicity because they form MHC-I complexes of low stability. Given this we concluded that pMHC-I complex stability is important in the generation of T cell responses.

Over the last few decades, accurate and reliable *in silico* methods capable of predicting the binding affinity of peptide binding to MHC-I have been developed (17–22) and these have been highly successful in narrowing down the search for T cell

^{*}Laboratory of Experimental Immunology, Department of International Health, Immunology and Microbiology, Faculty of Health and Medical Sciences, University of Copenhagen, 2200 N Copenhagen, Denmark; [†]Instituto de Investigaciones Biotecnológicas, Universidad Nacional de San Martín, San Martín, B 1650 HMP, Buenos Aires, Argentina; and [‡]Center for Biological Sequence Analysis, Technical University of Denmark, 2800 Kongens Lyngby, Denmark

¹M.R. and E.F. share first authorship.

ORCID: 0000-0001-9121-7149 (E.F.); 0000-0003-0791-8925 (I.K.N.); 0000-0001-8363-1999 (S.B.).

Received for publication April 6, 2016. Accepted for publication June 8, 2016.

This work was supported by the National Institute of Allergy and Infectious Diseases, National Institutes of Health, Department of Health and Human Services under Contracts HHSN272201200010C and HHSN27200900045C and the Agencia Nacional de Promoción Científica y Tecnológica, Argentina (Grant PICT-2012-0115).

Address correspondence and reprint requests to Prof. Søren Buus, Laboratory of Experimental Immunology, Department of International Health, Immunology and Microbiology, Faculty of Health and Medical Sciences, University of Copenhagen, 2200 N Copenhagen, Denmark. E-mail address: sbuus@sund.ku.dk

The online version of this article contains supplemental material.

Abbreviations used in this article: ANN, artificial neural network; AUC, area under the receiver operating characteristic curve; IEDB, Immune Epitope database; MHC-I, MHC class I; PCC, Pearson's correlation coefficient; pMHC, peptide–MHC-I.

Copyright © 2016 by The American Association of Immunologists, Inc. 0022-1767/16/\$30.00

epitopes. However, the high false discovery rate in terms of immunogenicity (i.e., that only a few of the peptides, which are predicted and subsequently shown to bind to MHC-I, are actually recognized by T cells postinfection) remains a problem. Recently, we have shown that accurate predictors of pMHC-I complex stability can be developed and that these predict immunogenic peptides to form pMHC-I complexes of higher stability compared with nonimmunogenic peptides (1, 23). These studies and corresponding predictors covered a limited set of human MHC-I allotypes. In this study, we have generated a large body of pMHC-I complex stability data using a large repository of human MHC-I molecules and synthetic peptides with the aim of developing pan-specific predictors of MHC-I complex stability. Using these data, we report the development of an artificial neural network (ANN)-based pan-specific predictor of pMHC-I complex stability and evaluate its accuracy in predicting pMHC-I complex stability. By integrating this novel predictor with state-of-the-art predictors of pMHC-I binding affinity, we find that constructing a model with 80% weight on affinity and 20% weight on stability, the prediction of MHC-I ligands and T cell epitopes is significantly improved beyond that of any of the two models alone. Current predictors of pMHC-I binding affinity are trained on a significantly larger body of data compared with available pMHC-I complex stability data. To assess whether the relative low contribution of stability predictions was influenced by this unbalance in data volume, we evaluated the performance of predictors trained on balanced pMHC-I binding affinity data and pMHC-I complex stability data. Evaluating such predictors, we find that the integrative model with similar weight on affinity and stability achieves optimal performance, and that this significantly outperforms the two individual models for prediction of both MHC ligands and T cell epitopes.

Materials and Methods

Peptide–MHC-I complex stability

A total of 9023 in-house nonameric peptides were predicted for binding to 76 HLA-I molecules using the NetMHCpan 2.8 prediction server (17, 18). Peptides with a predicted high-affinity binding (<500 nM or 2% rank) to multiple of the available MHC-I molecules were grouped into 12 peptide sets of 383 peptides each. Each peptide set was screened for pMHC-I complex stability on multiple MHC-I allotypes according to Table I using a scintillation proximity-based peptide–MHC-I dissociation assay (24). The reported half-life of pMHC-I complexes is the geometric mean of the half-life from two independent experiments. This data set was combined with an in-house small-scale data set from earlier studies arriving at a data set of 28,939 measurements covering 80 HLA-A and -B molecules.

Data sets

For the training of the stability predictor, only data from HLA molecules characterized by at least one stable binding and two nonstable binding measurements (using a half-life threshold at 37°C of 2 h) were included. This left a data set of pMHC-I complex half-lives with 28,166 measurements covering 75 HLA molecules. The data were enriched with 1000 random natural 9mer peptides for each allele. These peptides were selected to have predicted affinities (using *NetMHCpan-2.8*) weaker than 20,000 nM and were assigned a stability value of 0 h. For retraining of the *NetMHCpan* method, the Immune Epitope database (IEDB) data set corresponding to version 2.8 was used. This data set consists of 136,153 peptide affinity measurements covering 152 distinct MHC molecules.

To generate balanced data sets where an equal number of HLA allotypes were represented for both stability and affinity, and for each allotype an equal number of peptide data points and peptide binders were represented, we applied the following procedures: for stability, the peptides were classified into binders and nonbinders using a threshold of a stability half-life of 1 h. For affinity, the same classification was made using an affinity threshold of 500 nM. Next, the subset of 58 allotypes covered by at least one binder and two nonbinders for both affinity and stability was identified. For each of these allotypes, the number of binders/nonbinders was

identified as the lowest number from the stability and affinity data set, respectively, and this number of binders/nonbinders was selected from each data set. This selection was done prioritizing peptides with measured binding to few HLA alleles to ensure the highest degree of peptide diversity. After this procedure, a balanced data set was constructed for both affinity and stability, each containing 17,998 data points covering 58 alleles. As for the complete data set, these data sets were enriched with 1000 random natural 9mer for each allotype. As before, these peptides were selected to have predicted affinities (using *NetMHCpan-2.8*) weaker than 20,000 nM.

ANN training

ANNs were trained as a feed-forward ANN method (25) using either Blosom50 encoding with a normalization factor of 5 or sparse encoded with 1 of the 20 inputs being 0.9 and the remaining 19 being 0.05 with 40, 50, or 60 neurons in the hidden layer. In each case, the pool of unique peptides was split into five sets in a typical 5-fold cross-validation scheme with all peptide–HLA-I stability data for a given peptide placed in the same group (26) (in this way, no peptide can belong to more than one group); 4/5 of the data was used for training and the remaining 1/5 was left for testing and early stop. Before ANN training, half-life values were transformed to a value between 0 and 1. The transformation used was $s = 2^{-t_0/th}$, where s is the transformed value, t_h is the measured pMHC-I complex half-life (in hours), and t_0 is a threshold value that is fitted to obtain a suitable distribution of the data points for training purpose. From data shown in Table II, it is apparent that the HLA-I allotypes display highly variable ability to form stable complexes; by way of example, the 95th percentile (the half-life for the top 5% most stable peptide) for HLA-A*03:01 was 50 times greater than the corresponding value for HLA-B*08:01 (see Table II). This large range in half-life values makes rescaling of data before ANN training not trivial, yet essential. Given the data, it is likely that an allotype-specific threshold for rescaling would result in the best performing network, because information about the stability range of each specific allotype would be encoded in the network. However, as discussed later, this is not an applicable solution when developing a pan-specific predictor, and as an alternative to allotype-specific rescaling, we therefore probed different global rescaling thresholds in the range 0.5–2 h to optimize the predictive performance of the networks.

For comparison with pan-specific ANNs, the Pearson's correlation coefficient (PCC) was calculated for each allotype included in the training set from the test set predictions. The PCC of each allotype was then compared between networks using a binomial test counting the number of times a network outperformed the other network excluding ties.

Prediction of T cell epitopes and MHC ligands

The predictive performance of the different prediction methods was evaluated on a large set of 9mer T cell epitopes and MHC-I ligands. T cell epitopes were extracted from both the IEDB (27) and the SYFPEITHI (28) databases, whereas MHC-I ligands were obtained from the SYFPEITHI database only. The evaluation data sets were filtered for predicted binding affinity to the reported MHC restriction element, peptide length, and presence in the training data set as described by Jørgensen et al. (23). In brief, all T cell epitopes and ligands were filtered for predicted binding affinity to the restricting MHC using *NetMHCpan-2.8* with a rank threshold for binding of $\leq 10\%$ rank to exclude erroneously annotated epitopes/ligands. The evaluation set was further restricted to exclude HLA-peptide pairs present in the training data set. This filtered data set consisted of 1268 ligands representing 35 allotypes and 1230 epitopes representing 28 allotypes, respectively. Because the data were unevenly distributed among the allotypes, a balanced data set was extracted limiting the number of ligands/epitope per allotype to 100 and filtering out those allotypes with <10 data points. This final evaluation set consisted of 1058 ligands and 598 epitopes, covering 31 and 23 different allotypes, respectively. For evaluation of methods trained on balanced binding data sets, the epitope and ligand data set consisted, after filtering and size restriction, of 1085 MHC ligands and 856 T cell epitopes, covering 31 and 23 different allotypes, respectively. The increase in size of the evaluation data are explained by the dramatic reduction in the size of affinity data in the balanced training data set (from 136,153 to 17,998 data points).

The performance of the ANN methods was evaluated as described earlier (17, 29) using the area under the receiver operating characteristic curve (AUC) performance measure: the source protein of each ligand/epitope was fragmented into 9mers, the reported ligand/epitope was annotated as positive and every other 9mer as negative. In this way, one AUC value is calculated for each HLA-ligand/epitope pair, and the performance of a method is reported as the average of the AUC over all HLA-ligand/epitope pairs.

Table I. Grouping of HLA-I allotypes by predicted binding motif

Group	Allotype											
	A1	A2	A3	A24	A26	B7	B8	B27	B39	B44	B58	B62
	HLA-A*01:01	HLA-A*02:01	HLA-A*03:01	HLA-A*23:01	HLA-A*25:01	HLA-B*07:02	HLA-B*08:01	HLA-B*27:02	HLA-B*15:09	HLA-B*18:01	HLA-B*15:17	HLA-B*15:01
	HLA-A*29:02	HLA-A*02:03	HLA-A*11:01	HLA-A*24:02	HLA-A*26:01	HLA-B*35:01	HLA-B*08:02	HLA-B*27:03	HLA-B*15:10	HLA-B*40:01	HLA-B*57:01	HLA-B*15:02
	HLA-A*30:02	HLA-A*02:10	HLA-A*30:01	HLA-A*24:03	HLA-A*26:02	HLA-B*35:03	HLA-B*08:03	HLA-B*27:05	HLA-B*38:01	HLA-B*40:13	HLA-B*57:02	HLA-B*15:42
	HLA-A*43:01	HLA-A*02:11	HLA-A*31:01	HLA-A*24:07	HLA-A*26:03	HLA-B*39:10	HLA-B*14:01	HLA-B*27:20	HLA-B*39:01	HLA-B*41:01	HLA-B*57:03	HLA-B*46:01
	HLA-A*80:01	HLA-A*02:12	HLA-A*33:03	HLA-A*24:19	HLA-A*66:01	HLA-B*42:01	HLA-B*14:02		HLA-B*39:02	HLA-B*44:05	HLA-B*58:01	
		HLA-A*02:16	HLA-A*68:01	HLA-A*32:07	HLA-A*68:23	HLA-B*42:02			HLA-B*39:06	HLA-B*44:05		
		HLA-A*02:19	HLA-A*74:01			HLA-B*51:01			HLA-B*48:01			
		HLA-A*02:50				HLA-B*54:01						
		HLA-A*03:19				HLA-B*55:01						
						HLA-B*56:01						
						HLA-B*81:01						
						HLA-B*83:01						

pMHC-I binding was predicted using the *NetMHCpan-2.8* prediction algorithm for 76 MHC-I allotypes, which were subsequently grouped according to overlapping binding motifs resulting in 12 groups. Binding data for the seven allotypes denoted in italics were inconclusive.

Results

Large-scale analysis of peptide–HLA-I stability

We have recently reported the generation of a predictor trained on pMHC-I stability data and have shown that such a tool can accurately predict immunogenic MHC-I ligands and that combining this tool with state-of-the-art predictors of pMHC-I binding (*NetMHCpan*) can improve the predictive performance of T cell epitopes and ligands (1, 23). These studies were conducted on a representative set of human MHC-I allotypes, and the resulting predictive tool has limited allotype coverage. In this study, we have sought to reinforce these findings and generate sufficient pMHC-I stability data to develop pan-specific predictors of pMHC-I stability. Having 76 HLA-I allotypes and >9000 nonameric peptides at our disposal, we devised a strategy to allow rapid screening of pMHC-I stability. All nonameric peptides were predicted for binding to each HLA-I allotype using *NetMHCpan-2.8*. Next, peptides predicted to bind to multiple HLA-I allotypes were grouped resulting in 12 peptide sets that could be used to screen for stability on 4–12 different allotypes (Table I). This allowed us to screen the same peptide set on multiple HLA-I allotypes, increasing the throughput of stability assay. Each peptide group was limited to contain 383 peptides to fit the experimental layout. Having compiled the 12 peptide sets, each HLA-I allotype was screened for pMHC-I stability to the representative peptide set using a scintillation proximity assay-based pMHC-I dissociation assay (24). For 69 of the 76 allotypes, we were able to obtain conclusive data (data for 7 allotypes denoted in italics in Table I were inconclusive). Combining this data set with a small-scale in-house data set resulted in a data set of 28,939 individual pMHC-I half-life data points covering 80 HLA-I allotypes. Peptides forming stable pMHC-I complexes ($t_{1/2} > 1$ h) were found for most HLA-I allotypes. Of these, only data from allotypes covered by at least one binding and two nonbinding peptides (defined using a half-life threshold of 2 h) were included in the training data for the pan-specific predictor, arriving at the data set of 28,166 9mer measurements covering 75 HLA molecules (Table II).

Development of ANN predictors of pHLA-I complex stability

ANNs constitute a powerful machine-learning framework for data mining, and ANNs have successfully been applied to develop highly efficient predictors of peptide–MHC-I binding (17, 18). To obtain a pan-specific predictor, we applied the *NetMHCpan* approach earlier developed for pan-specific prediction of peptide–MHC binding affinity. In this study, an MHC-I pseudosequence was extracted from the polymorphic MHC-I residues in potential contact with the bound peptide (17), and this pseudosequence was encoded together with each peptide in the training. The data were partitioned and the pan-specific stability predictor was trained as described in *Materials and Methods*. Before training, the half-life values were rescaled as described in *Materials and Methods*. Both allotype-specific and global rescaling thresholds ranging from 0.5 to 2 h (here denoted t_0) were used to identify the optimal rescaling. The highest performance was obtained using allotype-specific rescaling values (PCC = 0.693) (Fig. 1). However, using allotype-specific rescaling is not a viable option because our current aim is to develop a pan-specific prediction method, which by its very nature should be applicable also to allotypes without experimental data; that is, such a method cannot be made dependent on the availability of specific allotype data. Moving forward, we therefore selected the global rescaling threshold with the highest performance ($t_0 = 1$ h, PCC = 0.676). This performance is slightly lower than the performance of the network using allotype-specific rescaling, yet statistically comparable ($p = 0.901$, binomial test excluding ties).

Table II. Summary of pHLA-I stability data

Group	Allotype	Data Points	$t_{1/2}$ Percentile		Group	Allotype	Data Points	$t_{1/2}$ Percentile	
			95th	5th				95th	5th
A1	HLA-A*01:01	220	40.3	0	B7	HLA-B*07:02	647	9.8	0
	HLA-A*29:02	383	20.4	0.1		HLA-B*35:01	650	7.4	0
	HLA-A*30:02	367	14.1	0		HLA-B*35:03	376	7.0	0.2
	HLA-A*43:01	359	6.7	0		HLA-B*39:10	376	32.5	0
	HLA-A*80:01	351	2.0	0		HLA-B*42:01	350	34.1	0.4
A2	HLA-A*02:01	1023	28.7	0.1	HLA-B*42:02	376	13.0	0	
	HLA-A*02:03	379	46.7	0.3	HLA-B*51:01	353	15.0	0	
	HLA-A*02:10	370	31.3	0.1	HLA-B*54:01	357	10.8	0	
	HLA-A*02:11	379	28.9	0.8	HLA-B*55:01	378	9.7	0	
	HLA-A*02:12	372	37.4	0.4	HLA-B*56:01	368	7.8	0	
	HLA-A*02:16	373	36.5	0.4	HLA-B*81:01	367	2.7	0	
	HLA-A*02:19	354	39.6	0.2	B8	HLA-B*08:01	465	1.2	0
	HLA-A*02:50	375	23.2	0.7		HLA-B*08:03	350	1.9	0
				HLA-B*14:01 (C67S)		374	0.6	0	
A3	HLA-A*03:01	861	52.2	0	HLA-B*14:02 (C67S)	382	4.5	0	
	HLA-A*11:01	577	71.7	0.1	B27	HLA-B*27:02	359	34.9	0
	HLA-A*30:01	386	37.4	0		HLA-B*27:03	368	25.1	0
	HLA-A*31:01	382	77.3	0		HLA-B*27:05	366	25.3	0
	HLA-A*33:03	371	56.0	0		HLA-B*27:20	368	8.3	0
	HLA-A*68:01	387	56.0	0	B39	HLA-B*15:10	363	7.7	0
	HLA-A*74:01	361	94.5	0		HLA-B*39:01	680	9.3	0
				HLA-B*39:02		375	14.9	0	
A24	HLA-A*23:01	372	37.3	0.4	HLA-B*39:06 (C67S)	379	0.5	0	
	HLA-A*24:02	573	30.2	0.2	B44	HLA-B*18:01	359	8.0	0
	HLA-A*24:03	512	37.6	0.2		HLA-B*40:01	302	2.3	0
	HLA-A*24:07	352	24.4	0.2		HLA-B*41:01	368	3.5	0
	HLA-A*24:19	378	2.7	0		HLA-B*44:05	352	12.1	0
	HLA-A*32:07	364	43.5	0		HLA-B*45:01	368	5.3	0
				B58		HLA-B*15:17	362	6.3	0
A26	HLA-A*25:01	363	21.7		0	HLA-B*57:01	349	12.4	0
	HLA-A*26:01	272	23.8		0	HLA-B*57:02	363	7.2	0
	HLA-A*26:02	343	35.6		0	HLA-B*57:03	369	7.8	0.6
	HLA-A*26:03	341	15.3	0	HLA-B*58:01	379	21.3	0	
	HLA-A*66:01	382	3.6	0	B62	HLA-B*15:01	1070	12.5	0
HLA-A*68:23	369	9.9	0	HLA-B*15:02		349	5.7	0	
Others	HLA-A*02:05	21	26.3	0.6		HLA-B*46:01	361	5.3	0
	HLA-A*32:01	32	57.2	0					
	HLA-A*68:02	16	15.8	0.1					
	HLA-A*69:01	15	4.1	0					
	HLA-B*13:02	7	6.0	0					
HLA-B*35:08	27	8.7	0						
HLA-B*40:02	19	24.4	0.5						

Seventy-five HLA-I allotypes were grouped into 12 supertype groups by predicted binding motif and screened for pMHC-I complex stability to peptide sets representing each supertype binding preference. The extra group, Others, contains molecules with few binding data obtained from an in-house database. For each molecule the total number of peptide binding measurements and 95th and 5th percentile half-life values are indicated. The 95th percentile value corresponds to the half-time of the top 5% most stable peptide and 5th percentile value corresponds to the bottom 5% least stable peptide for each HLA molecule.

Integrating stability- and affinity-based predictions of pMHC-I interaction

We have previously reported that combining predictors of pMHC-I affinity with predictors of pMHC-I stability improves the prediction of T cell epitopes and MHC-I ligands (23). However, this analysis was limited to a small set of allotypes. In this study, we have expanded this observation to a much larger set of MHC molecules. For this purpose, we extracted T cell epitopes and MHC-I ligands from the IEDB (27) and SYFPEITHI (28) to use as benchmark data sets for evaluation of the combined predictors. We applied a simple weighted average between affinity- and stability-based predictions to identify the optimal contribution of each to the prediction of T cell epitopes and ligands. In a 5-fold cross-validation, a ratio of affinity-based and stability-based predictions of 0.9:0.1 and 0.8:0.2 was consistently found to be optimal with respect to the prediction of MHC-I ligands and MHC-I epitopes, respectively (Fig. 2). These combinations significantly outperformed each of the methods in

isolation ($p = 0.0013$, binomial test excluding ties for MHC-I ligands, and $p = 0.0005$, binomial test excluding ties for MHC-I epitopes), thus consolidating our earlier findings (23).

Analysis of the predictive performance using balanced networks

The earlier analysis shows that the affinity predictions currently contribute more than the stability predictions to the predictive power of the optimal method. However, the differences in the amounts of data available for the training of the affinity- and stability-based predictors are substantial, with ~140,000 affinity measurements covering >140 MHC-I allotypes compared with a “mere” ~28,000 stability measurements covering 75 HLA alleles. To investigate the impact of this difference, we extracted balanced data sets from the pMHC-I binding affinity and stability data sets containing the same number of “positives” and “negatives” data point (for details, see *Materials and Methods*). Next, predictors of




FIGURE 1. Predictive performance of ANN predictors of pMHC-I complex stability. Graphic shows the average PCC and SCC values for the predictors trained with different t_0 thresholds.

pMHC-I binding affinity and stability were trained as described earlier using a rescaling factor $t_0 = 1$ h for the stability network. Using the aforementioned weighted average scheme and the large set of IEDB and SYFPEITHI T cell epitopes and MHC ligands, we found that the combination of the two methods also in this study significantly improved the prediction beyond each of the individual methods, with an optimal ratio of 0.7:0.3 for MHC-I ligands and 0.6:0.4 for MHC-I epitopes ($p < 0.05$, binomial test excluding ties) (Fig. 3). Thus, using methods trained on size balanced data sets, the relative importance of stability predictions compared with that of affinity predictions is increased. Also, we consistently found in both benchmarks that the relative importance of stability is higher for the prediction of T cell epitopes than for the prediction of MHC ligands.

Source of the performance gain

Having demonstrated that inclusion stability data led to an improved predictive performance, we next asked ourselves what was the source of this gain in performance. The two predictors for binding affinity and stability are highly correlated (Supplemental Fig. 1), and one could ask whether the gain in performance demonstrated the earlier results from the different nature of data (stability versus affinity) or simply because more data were included in the method development. We addressed this question by identifying the list peptides for which both binding affinity and stability measurements existed. This set contained 7600 peptides covering 58 allotypes. Given this data set, we retrained three new pan-specific predictors: a pan-specific affinity predictor trained on the affinity peptide data set excluding the peptide data set with both affinity and stability measurements (termed A), a pan-specific affinity predictor trained on affinity data from the peptides with both affinity and stability measurements (termed B), and a pan-stability predictor trained on stability data from the peptides with both affinity and stability measurements (termed C). Both B and C were trained as described earlier using the same peptide data,

identical data partitioning, and the same number of added artificial random negatives. Predictor A was trained using the *NetMHCpan* protocol described by Hoof et al. (18). Next, linear weighted combinations of A+B and A+C were constructed (the optimal weights were identified separately for each method), and the performance of the optimal combined methods was compared on the epitope benchmark data set. Doing this, we found that the method combining affinity and stability (A+C) significantly outperformed the method combining affinity with affinity (A+B) ($p < 0.001$, binomial test excluding ties). This result thus demonstrates that the observed increased performance is due to the different nature of the data and not due to the simple fact that more data were included.

The *NetMHCstabpan* server

The final pan-specific MHC-I peptide stability prediction method trained on the complete data set of 28,166 9mer measurements covering 75 HLA alleles was implemented as a Web server that is publicly available at <http://www.cbs.dtu.dk/services/NetMHCstabpan>. The server allows for prediction of pMHC-I binding stability to any HLA molecule of known sequence. Predictions for non-9mer peptides are made using the approximation method described in the work of Lundegaard et al. (30). Submissions are accepted in two formats: as a list of peptides, or as protein sequences in FASTA format. If the submission is in FASTA format, the predictions are made for all overlapping peptides of the user-selected length(s) in the sequence. MHC molecules can be selected from a predefined list, or the user can upload a full-length MHC protein sequence in FASTA format. The output is given as a table where each row contains the peptide, the predicted stability score, converted half-life in hours, and a percentile rank score. The predictor also allows the user to include affinity predictions (calculated by *NetMHCpan*-2.8). Selecting this adds the predicted affinity score, the converted binding affinity (in nM), a combination of affinity and stability scores using the previously described ratio of 0.8:0.2 (this ratio can be altered by the user), and a percentile rank score of the combined prediction score estimated from the harmonic mean of the stability and affinity rank values, respectively, to the output.

Discussion

Binding of peptides to MHC-I molecules is the single most selective step in Ag presentation to CD8⁺ T cells, and substantial experimental and computational work has been dedicated to characterize this event. For a peptide to be selected for Ag presentation and trigger translocation to the cell surface in complex with a restricting MHC-I molecule, it must bind with sufficient affinity and stability. Moreover, once presented, it must remain in the complex with the MHC on the cell surface for sufficient time to allow recognition of a circulating CTL. Given this, it has been argued that both binding stability and affinity are essential correlates to peptide immunogenicity and immune dominance; furthermore, we have in a set of recent publications suggested that of the two, stability is the better predictor of CTL immunogenicity (1, 23).

In this study, we have extended those previous studies and generated a large set of almost 30,000 peptide stability binding measurements covering >75 HLA class I molecules. The data were created using a cost-effective screening strategy, where peptides were first screened for predicted binding affinity using *NetMHCpan* v2.8, and second partitioned into 12 groups with predicted shared binding characteristics, allowing screening against multiple HLA allotypes leading to efficient identification of stable binders. From these data, we have constructed a pan-specific predictor of HLA-I peptide binding stability. To the best

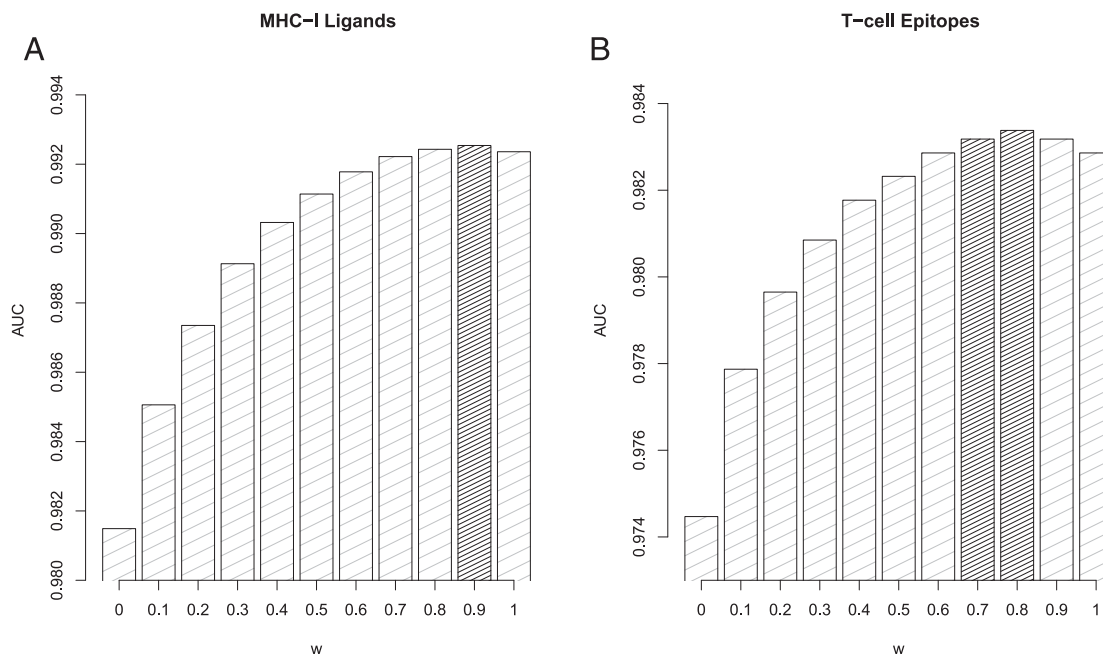


FIGURE 2. Performance of epitope/ligand predictions in terms of AUC. Predictive performance for the stability ($w = 0$) and affinity ($w = 1$) predictors and their combinations for **(A)** MHC-I ligands and **(B)** T cell epitopes. Gray columns show combinations with significantly improved ($p < 0.05$) performance compared with both affinity alone ($w = 1$) and stability alone ($w = 0$) and insignificant difference ($p > 0.05$) compared with the optimal model (the value of w with highest AUC value). Statistical differences were evaluated using binomial test excluding ties.

of our knowledge, this is the first pan-specific predictor of HLA-I stability.

This pan-specific predictor was constructed following the pipeline defined earlier for the pan-specific predictor of HLA-I affinity, *NetMHCpan* (17). The predictive performance was evaluated using data representing T cell epitopes or ligands, which were downloaded from the SYFPEITHI and IEDB databases.

Comparing the predictive performance of the stability predictor with that of the NetMHCpan-2.8 affinity predictor demonstrated that the affinity-based predictor significantly outperformed the stability prediction. However, integrating the two predictors with a relative weight on affinity of 80% showed a significantly improved prediction of CTL epitopes compared with any of the two methods alone.

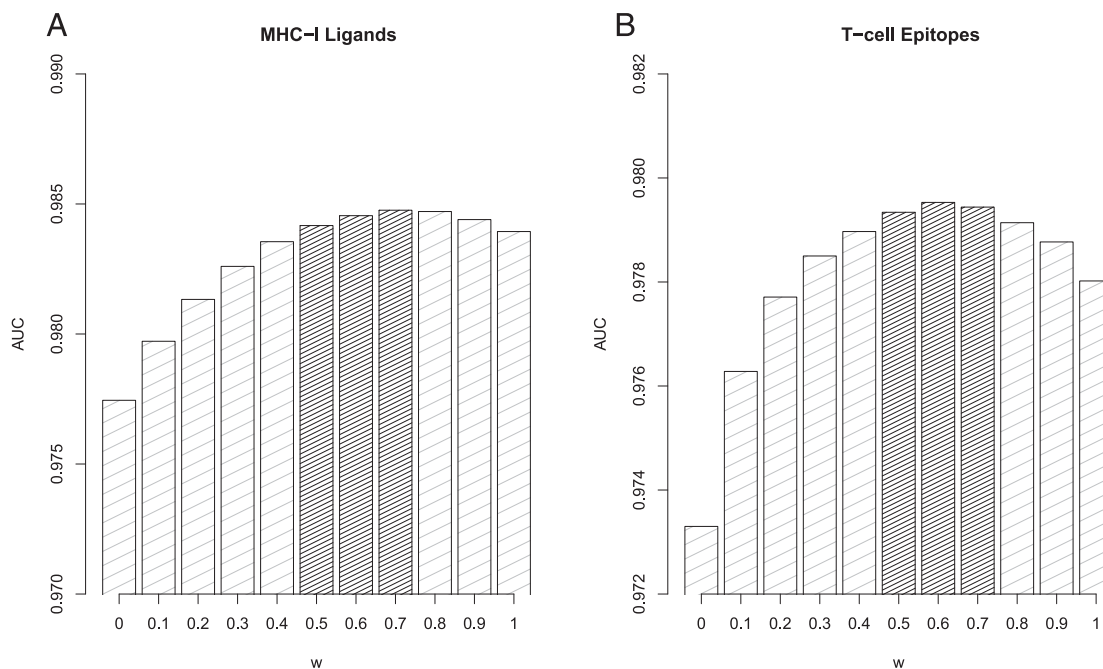


FIGURE 3. Performance of epitope/ligand predictions in terms of AUC for methods trained using balanced data sets. Prediction performance for the stability ($w = 0$) and affinity ($w = 1$) methods training on balanced data sets and their combinations for **(A)** MHC-I ligands and **(B)** T cell epitopes. Gray columns show combinations with significantly improved ($p < 0.05$) performance compared with both affinity alone ($w = 1$) and stability alone ($w = 0$) and insignificant difference ($p > 0.05$) compared with the optimal model (the value of w with highest AUC value). Statistical differences were evaluated using binomial test excluding ties.

At first glance the emphasis on the affinity predictor may appear at odds with the notion that stability is the better correlate of CTL immunogenicity. However, comparing the relative importance of stability and affinity-based predictors is not a straightforward task because there are two main biases that lean in favor of generating an accurate affinity-based method: a greater quantity and diversity of data available for training of the method. The difference in the amount of data available for the training of affinity- and stability-based predictors is substantial, with ~140,000 affinity measurements covering >140 MHC-I allotypes compared with ~28,000 stability measurements covering 70 HLA alleles. To level this difference, we constructed a balanced set of binding data with the same number of allotypes and data points (i.e., binders and non-binders) for the affinity and stability data sets, and used these balanced data to retrain pan-specific binding affinity and stability predictors. When integrating the predictions of these two methods, and evaluating the performance on the large set of HLA ligands and T cell epitopes, we found that the relative weight on binding affinity was decreased to ~60%. Although the balanced data set has leveled the difference in size and allele coverage, one important aspect remains very different between the two data sets. From the ~18,000 data points in the balanced data set, the stability data set has 5581 unique peptides, whereas the affinity data set has 9161 unique peptides. This difference in peptide diversity is a consequence of the strategy used to select the peptides for the stability experiments. In retrospect, we made a suboptimal strategy decision, which unfortunately cannot be resolved with the existing stability data set. As more, and more diverse, stability data will be generated in the future, we speculate that the relative importance between the two methods will be shifted further in favor of emphasizing binding stability.

Even though the present tool is trained on the hitherto largest set of peptide MHC-I stability data, the volume of the data set remains small compared with the data available for peptide MHC-I affinity. Furthermore, the present tool has been trained on nonamer peptide data only limited to the HLA-A and HLA-B isotypes. Given this and experiences from earlier works (18, 21, 31, 32), we strongly expect the performance of the tools will increase as more data covering different peptide length, HLA isotypes/allotypes (potentially also MHC allotypes covering other species) become available.

Notably, even though in this study we have demonstrated a high performance for prediction of T cell epitopes driven by binding affinity and stability, other factors including self- versus non-self-peptides similarity (33, 34) and T cell propensity (35) are essential to characterize and predict peptide immunogenicity.

In conclusion, in this study, we have developed, to our knowledge, the first pan-specific predictor for binding stability of peptide HLA-I complexes and have demonstrated how this tool can significantly enhance the ability to predict T cell peptide immunogenicity. A Web server implementing the method is available publicly at <http://www.cbs.dtu.dk/services/NetMHCstabpan>.

Disclosures

The authors have no financial conflicts of interest.

References

- Harndahl, M., M. Rasmussen, G. Roder, I. Dalgaard Pedersen, M. Sørensen, M. Nielsen, and S. Buus. 2012. Peptide-MHC class I stability is a better predictor than peptide affinity of CTL immunogenicity. *Eur. J. Immunol.* 42: 1405–1416.
- van der Burg, S. H., M. J. Visseren, R. M. Brandt, W. M. Kast, and C. J. Melief. 1996. Immunogenicity of peptides bound to MHC class I molecules depends on the MHC-peptide complex stability. *J. Immunol.* 156: 3308–3314.
- Micheletti, F., R. Guerrini, A. Formentin, A. Canella, M. Marastoni, M. Bazzaro, R. Tomatis, S. Traniello, and R. Gavioli. 1999. Selective amino acid substitutions of a subdominant Epstein-Barr virus LMP2-derived epitope increase HLA/peptide complex stability and immunogenicity: implications for immunotherapy of Epstein-Barr virus-associated malignancies. *Eur. J. Immunol.* 29: 2579–2589.
- Burrows, J. M., K. K. Wynn, F. E. Tynan, J. Archbold, J. J. Miles, M. J. Bell, R. M. Brennan, S. Walker, J. McCluskey, J. Rossjohn, et al. 2007. The impact of HLA-B micropolymorphism outside primary peptide anchor pockets on the CTL response to CMV. *Eur. J. Immunol.* 37: 946–953.
- Nicholls, S., K. P. Piper, F. Mohammed, T. R. Dafforn, S. Tenzer, M. Salim, P. Mahendra, C. Craddock, P. van Ender, H. Schild, et al. 2009. Secondary anchor polymorphism in the HA-1 minor histocompatibility antigen critically affects MHC stability and TCR recognition. *Proc. Natl. Acad. Sci. USA* 106: 3889–3894.
- Spierings, E., S. Gras, J. B. Reiser, B. Mommaas, M. Almekinders, M. G. Kester, A. Chouquet, M. Le Gorrec, J. W. Drijfhout, F. Ossendorp, et al. 2009. Steric hindrance and fast dissociation explain the lack of immunogenicity of the minor histocompatibility HA-1Arg Null allele. *J. Immunol.* 182: 4809–4816.
- Lipford, G. B., S. Bauer, H. Wagner, and K. Heeg. 1995. In vivo CTL induction with point-substituted ovalbumin peptides: immunogenicity correlates with peptide-induced MHC class I stability. *Vaccine* 13: 313–320.
- van Stipdonk, M. J., D. Badia-Martinez, M. Sluiter, R. Offringa, T. van Hall, and A. Achour. 2009. Design of agonistic altered peptides for the robust induction of CTL directed towards H-2Db in complex with the melanoma-associated epitope gp100. *Cancer Res.* 69: 7784–7792.
- Lazarski, C. A., F. A. Chaves, S. A. Jenks, S. Wu, K. A. Richards, J. M. Weaver, and A. J. Sant. 2005. The kinetic stability of MHC class II:peptide complexes is a key parameter that dictates immunodominance. *Immunity* 23: 29–40.
- Hall, F. C., J. D. Rabinowitz, R. Busch, K. C. Visconti, M. Belmares, N. S. Patil, A. P. Cope, S. Patel, H. M. McConnell, E. D. Mellins, and G. Sonderstrup. 2002. Relationship between kinetic stability and immunogenicity of HLA-DR4/peptide complexes. *Eur. J. Immunol.* 32: 662–670.
- Pogue, R. R., J. Eron, J. A. Frelinger, and M. Matsui. 1995. Amino-terminal alteration of the HLA-A*0201-restricted human immunodeficiency virus polypeptide increases complex stability and in vitro immunogenicity. *Proc. Natl. Acad. Sci. USA* 92: 8166–8170.
- Brooks, J. M., R. A. Colbert, J. P. Mear, A. M. Leese, and A. B. Rickinson. 1998. HLA-B27 subtype polymorphism and CTL epitope choice: studies with EBV peptides link immunogenicity with stability of the B27:peptide complex. *J. Immunol.* 161: 5252–5259.
- Abdel-Motal, U. M., R. Friedline, B. Poligone, R. R. Pogue-Caley, J. A. Frelinger, and R. Tisch. 2001. Dendritic cell vaccination induces cross-reactive cytotoxic T lymphocytes specific for wild-type and natural variant human immunodeficiency virus type 1 epitopes in HLA-A*0201/Kb transgenic mice. *Clin. Immunol.* 101: 51–58.
- Vertuani, S., A. Sette, J. Sidney, S. Southwood, J. Fikes, E. Keogh, J. A. Lindencrona, G. Ishioka, J. Levitskaya, and R. Kiessling. 2004. Improved immunogenicity of an immunodominant epitope of the HER-2/neu proto-oncogene by alterations of MHC contact residues. *J. Immunol.* 172: 3501–3508.
- Grohmann, U., M. L. Belladonna, R. Bianchi, C. Orabona, S. Silla, G. Squillacioti, M. C. Fioretti, and P. Puccetti. 1999. Immunogenicity of tumor peptides: importance of peptide length and stability of peptide/MHC class II complex. *Cancer Immunol. Immunother.* 48: 195–203.
- Assarsson, E., J. Sidney, C. Oseroff, V. Paschetto, H. H. Bui, N. Frahm, C. Brander, B. Peters, H. Grey, and A. Sette. 2007. A quantitative analysis of the variables affecting the repertoire of T cell specificities recognized after vaccinia virus infection. *J. Immunol.* 178: 7890–7901.
- Nielsen, M., C. Lundegaard, T. Blicher, K. Lamberth, M. Harndahl, S. Justesen, G. Roder, B. Peters, A. Sette, O. Lund, and S. Buus. 2007. NetMHCpan, a method for quantitative predictions of peptide binding to any HLA-A and -B locus protein of known sequence. *PLoS One* 2: e796.
- Hoof, I., B. Peters, J. Sidney, L. E. Pedersen, A. Sette, O. Lund, S. Buus, and M. Nielsen. 2008. NetMHCpan, a method for MHC class I binding prediction beyond humans. *Immunogenetics* 61: 1–13.
- Karsiene, E., C. Lundegaard, O. Lund, and M. Nielsen. 2012. NetMHCcons: a consensus method for the major histocompatibility complex class I predictions. *Immunogenetics* 64:177–186.
- Peters, B., and A. Sette. 2005. Generating quantitative models describing the sequence specificity of biological processes with the stabilized matrix method. *BMC Bioinformatics* 6: 132.
- Andreatta, M., and M. Nielsen. 2016. Gapped sequence alignment using artificial neural networks: application to the MHC class I system. *Bioinformatics* 32: 511–517.
- Lundegaard, C., K. Lamberth, M. Harndahl, S. Buus, O. Lund, and M. Nielsen. 2008. NetMHC-3.0: accurate web accessible predictions of human, mouse and monkey MHC class I affinities for peptides of length 8–11. *Nucleic Acids Res.* 36:W509–W512.
- Jørgensen, K. W., M. Rasmussen, S. Buus, and M. Nielsen. 2014. NetMHCstab - predicting stability of peptide:MHC-I complexes; impacts for CTL epitope discovery. *Immunology* 141:18–26.
- Harndahl, M., M. Rasmussen, G. Roder, and S. Buus. 2011. Real-time, high-throughput measurements of peptide-MHC-I dissociation using a scintillation proximity assay. *J. Immunol. Methods* 374: 5–12.
- Nielsen, M., C. Lundegaard, P. Worning, S. L. Laue-møller, K. Lamberth, S. Buus, S. Brunak, and O. Lund. 2003. Reliable prediction of T-cell epitopes using neural networks with novel sequence representations. *Protein Sci.* 12: 1007–1017.
- Nielsen, M., C. Lundegaard, and O. Lund. 2007. Prediction of MHC class II binding affinity using SMM-align, a novel stabilization matrix alignment method. *BMC Bioinformatics* 8: 238.

27. Kim, Y., J. Ponomarenko, Z. Zhu, D. Tamang, P. Wang, J. Greenbaum, C. Lundegaard, A. Sette, O. Lund, P. E. Bourne, et al. 2012. Immune epitope database analysis resource. *Nucleic Acids Res.* 40: W525–W530.
28. Rammensee, H. G., T. Friede, and S. Stevanović. 1995. MHC ligands and peptide motifs: first listing. *Immunogenetics* 41: 178–228.
29. Stranzl, T., M. V. Larsen, C. Lundegaard, and M. Nielsen. 2010. NetCTLpan: pan-specific MHC class I pathway epitope predictions. *Immunogenetics* 62: 357–368.
30. Lundegaard, C., O. Lund, and M. Nielsen. 2008. Accurate approximation method for prediction of class I MHC affinities for peptides of length 8, 10 and 11 using prediction tools trained on 9mers. *Bioinformatics* 24: 1397–1398.
31. Rasmussen, M., M. Harndahl, A. Stryhn, R. Boucherma, L. L. Nielsen, F. A. Lemonnier, M. Nielsen, and S. Buus. 2014. Uncovering the peptide-binding specificities of HLA-C: a general strategy to determine the specificity of any MHC class I molecule. *J. Immunol.* 193: 4790–4802.
32. Nielsen, M., and M. Andreatta. 2016. NetMHCpan-3.0; improved prediction of binding to MHC class I molecules integrating information from multiple receptor and peptide length datasets. *Genome Med.* 8: 33.
33. Frankild, S., R. J. de Boer, O. Lund, M. Nielsen, and C. Kesmir. 2008. Amino acid similarity accounts for T cell cross-reactivity and for “holes” in the T cell repertoire. *PLoS One* 3: e1831.
34. Bresciani, A., S. Paul, N. Schommer, M. B. Dillon, T. Bancroft, J. Greenbaum, A. Sette, M. Nielsen, and B. Peters. 2016. T-cell recognition is shaped by epitope sequence conservation in the host proteome and microbiome. *Immunology* 148: 34–39.
35. Calis, J. J., M. Maybeno, J. A. Greenbaum, D. Weiskopf, A. D. De Silva, A. Sette, C. Kesmir, and B. Peters. 2013. Properties of MHC class I presented peptides that enhance immunogenicity. *PLOS Comput. Biol.* 9: e1003266.