



IMMUNOLOGY

RACER-m leverages structural features for sparse T cell specificity prediction

Ailun Wang^{1,2}, Xingcheng Lin^{3,4*}, Kevin Ng Chau^{1,2}, José N. Onuchic^{5,6}, Herbert Levine^{1,2,7}, Jason T. George^{6,8*}

Reliable prediction of T cell specificity against antigenic signatures is a formidable task, complicated by the immense diversity of T cell receptor and antigen sequence space and the resulting limited availability of training sets for inferential models. Recent modeling efforts have demonstrated the advantage of incorporating structural information to overcome the need for extensive training sequence data, yet disentangling the heterogeneous TCR-antigen interface to accurately predict MHC-allele-restricted TCR-peptide interactions has remained challenging. Here, we present RACER-m, a coarse-grained structural model leveraging key biophysical information from the diversity of publicly available TCR-antigen crystal structures. Explicit inclusion of structural content substantially reduces the required number of training examples and maintains reliable predictions of TCR-recognition specificity and sensitivity across diverse biological contexts. Our model capably identifies biophysically meaningful point-mutant peptides that affect binding affinity, distinguishing its ability in predicting TCR specificity of point-mutants from alternative sequence-based methods. Its application is broadly applicable to studies involving both closely related and structurally diverse TCR-peptide pairs.

INTRODUCTION

T cell immunity is determined by the interaction of a T cell receptor (TCR) with antigenic peptide (p) presented on the cell surface via major histocompatibility molecules (MHCs) (1). T cell activation occurs when there is a favorable TCR-pMHC interaction and, for the case of CD8⁺ effector cells, ultimately results in T cell killing of the pMHC-presenting cell (2). T cell-mediated antigen recognition confers broad immunity against intracellular pathogens as well as tumor-associated antigenic signatures (3). Thus, a detailed understanding of the specificity of individual T cells in a repertoire composed of many (~10⁸) unique T cell clones is required for understanding and accurately predicting many important clinical phenomena, including infection, cancer immunogenicity, and autoimmunity.

Because of the immense combinatorial complexity of antigen (~10¹³) and T cell (~10¹⁸) sequence space, initial conceptual progress in the field was made by studying simple forms of amino acid interactions, motivated by either protein folding ideas (4, 5) or random energy approaches (6, 7). Recent advances in high-throughput studies interrogating T cell specificity (8–10) together with the development of statistical learning approaches have at last enabled data-driven modeling as a tractable approach to this problem. Consequently, a number of approaches have been developed to predict TCR-antigen specificity (11–15). A majority of developed approaches input only TCR and pMHC primary sequence information. The persistent challenge with this lies in limited training data given that any reasonable sampling of antigens and T cells—or even an entire

human T cell repertoire—represents a very small fraction of sequence space. One persistent and notorious challenge of virtually all current models involves an inability to make reasonable specificity predictions on unseen epitopes that are excluded from training. As a result, many models underperform on sequences that are moderately dissimilar from their nearest neighbor in the training set, an issue that we refer to as “global sparsity.”

While global sparsity complicates inference extension to moderately dissimilar antigens, another distinct challenge exists for reliably predicting the behavior of closely related TCR-pMHC pairs that differ by a single-amino acid substitution, which we refer to as “local resolvability.” These “point-mutated” TCR-pMHC pairs require predictive methods capable of quantifying the effects of single-amino acid changes on the entire TCR-peptide interaction, a task often limited by lack of sufficient training examples required for reliable estimation of the necessary pairwise residues. Instead, a modeling framework aiming to discern such subtle differences between point mutants may benefit from learning the general rules of amino acid interactions at the TCR-peptide interface and their varied contributions to binding affinity. Resolving this very particular problem, discerning relevant point mutations in self-peptide and viral antigens, promises to deliver enhanced therapeutic utility in targeting cancer neoantigens, optimally selecting hematopoietic stem cell transplant donors, and predicting the immunological consequences of viral variants. Thus, local resolvability represents a distinct learning task wherein detailed reliable predictions need to be made on many small variations around a very specific TCR-pMHC pair.

Several structure-based approaches have also been used to better understand TCR-pMHC specificity. Detailed structural models that focus on a comprehensive description of TCR-pMHC interaction, including all-atom simulation and structural relaxation, are computationally limited to describing a few realized structures of interest (16, 17). Another strategy develops an AlphaFold-based pipeline to generate accurate three-dimensional (3D) structures from primary sequence information to improve the accuracy of TCR-pMHC binding predictions for hundreds of TCR-pMHC pairs (18). A previous hybrid

¹Center for Theoretical Biological Physics, Northeastern University, Boston, MA, USA. ²Department of Physics, Northeastern University, Boston, MA, USA. ³Department of Physics, North Carolina State University, Raleigh, NC, USA. ⁴Bioinformatics Research Center, North Carolina State University, Raleigh, NC, USA. ⁵Departments of Physics and Astronomy, Chemistry, and Biosciences, Rice University, Houston, TX, USA. ⁶Center for Theoretical Biological Physics, Rice University, Houston, TX, USA. ⁷Department of Bioengineering, Northeastern University, Boston, MA, USA. ⁸Department of Biomedical Engineering, Texas A&M University, Houston, TX, USA. *Corresponding author. Email: xingcheng_lin@ncsu.edu (X.L.); jason.george@tamuedu (J.T.G.)

approach (14) used crystal structural data together with known binding sequences to train an optimized binding energy model for describing TCR-pMHC interactions. This approach offered several advantages, including the ability to perform repertoire-level predictions within a reasonable time, along with a reduced demand for extensive training data. However, this model largely focused on a restricted set of peptide or TCR pairs using a single MHC class II (MHC-II) structural template and performed best at explaining mouse I-E^k-restricted systems. Thus, its ability to make reliable predictions for a structurally diverse collection of TCR and peptide pairs with a conserved human leukocyte antigen (HLA) allele restriction remains unknown.

Here, we leverage all available protein crystal structures of the most common human MHC-I allele variant, HLA-A*02:01, to develop a combined sequence-structural model of TCR-pMHC specificity that features biophysical information from a diversity of known structural templates. The general strategy of our approach is outlined in Fig. 1. We quantify the structural diversity in available crystal structures of TCR-pMHC complexes (19–21) and demonstrate that incorporating a small subset of available structural information is sufficient to enable reliable predictions of favorable interactions across a diverse set of TCR-antigen pairs. We show that, by using structural templates from closely related amino acid sequences, RACER-m generates reasonable predictions for previously unseen epitopes. Our results further suggest that the availability of structural information having close proximity to the true structure of a TCR-pMHC pair can ameliorate both global sparsity and local resolvability in discerning the immunogenicity of diverse and point-mutated antigenic variants.

RESULTS

Model development and identification of TCR-peptide pairs with structural templates

We build on our previous RACER framework developed primarily on the mouse MHC-II I-E^k system (14). Our approach, termed RACER multi-template (RACER-m), represents a comprehensive pipeline that leverages published crystal structures of known human TCR-pMHC pairs.

All 66 HLA-A*02:01-restricted systems with a TCR-pMHC published structure [Protein Data Bank (PDB)/Immune Epitope Database (IEDB)] available through www.rcsb.org were used as the structures of strong binders for training (22–24). Their 66 corresponding peptide and TCR variable CDR3 α and CDR3 β sequences were also used, and this list of TCR-pMHC pairs was further augmented by identification of all reported TCR-pMHC pairs in the publications that referenced the above structures, as part of the “ATLAS dataset.” In addition, the ATLAS database containing affinity information (K_d) for related TCR-peptide pairs (19) was used for cases where either a TCR or an epitope had substantial overlap with that of the sequences having structures. A threshold of 200 nM was used to define strong binders to be included in the ATLAS dataset, based on the reported K_d . Last, grouping by template was performed using hierarchical clustering based on structural similarity using an approach previously developed in the protein folding community (25, 26) followed by hierarchical clustering. In total, 163 unique TCR-peptide pairs and 66 structural templates were identified for training and validation (see the Supplementary Materials).

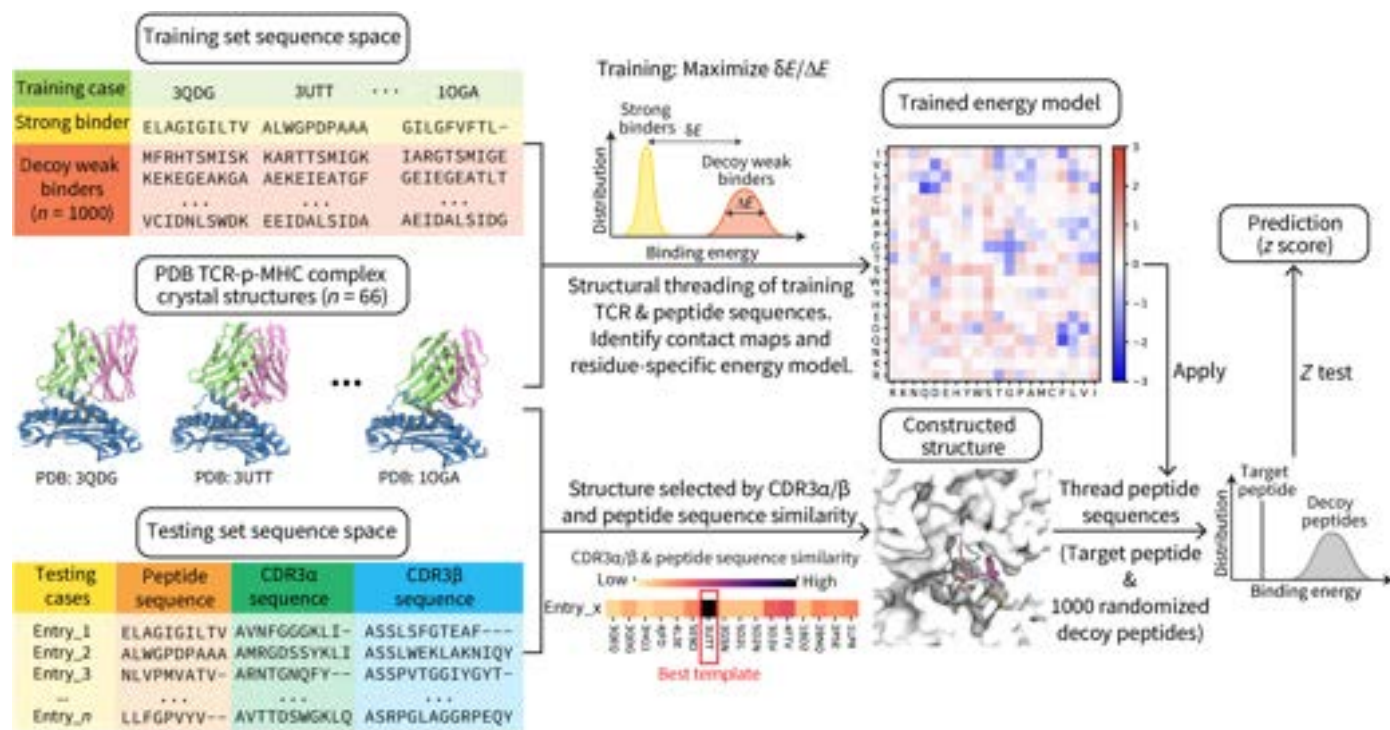


Fig. 1. RACER-m model architecture. Schematic representation of the training (top row) and testing (bottom row) processes in RACER-m. Sixty-six crystal structures of known strong binders were used as both training set and template structures for the testing processes, which cover several major clusters of TCR repertoires (MART-1, TAX, 1E6, NLV, and FLU) and other clusters with smaller size.

We next assessed the structural diversity of training templates by pairwise evaluation of structural similarity using a previously developed method referred to as mutual Q (25, 26). Mutual Q similarity defines a structural metric consisting of a sum of transformed pairwise distances between each residue in two structures normalized within the range of 0 to 1, which was then used to perform hierarchical clustering. We found that the identified structural clusters largely partition TCR-pMHC pairs according to immunological function (for example, TCR-pMHC pairs sharing a conserved antigen) with a few exceptions (Fig. 2A). Despite our focus only on a specified

HLA-restricted repertoire, the analysis, nonetheless, revealed clustering heterogeneity across all included structures: In some cases [e.g., Melanoma-associated antigen recognized by T cells-1 (MART-1) and TAX], substantial heterogeneity was observed and associated with enriched pairwise dissimilarity of TCR and peptide sequences. This, together with cross-cluster structural diversity, is a consequence of global sparsity given limited observed structures. On the other hand, we also identified structurally homogeneous clusters composed of TCR-pMHC pairs having near-identical pairwise sequence similarity (e.g., 1E6), yet these pairs have substantial differences in

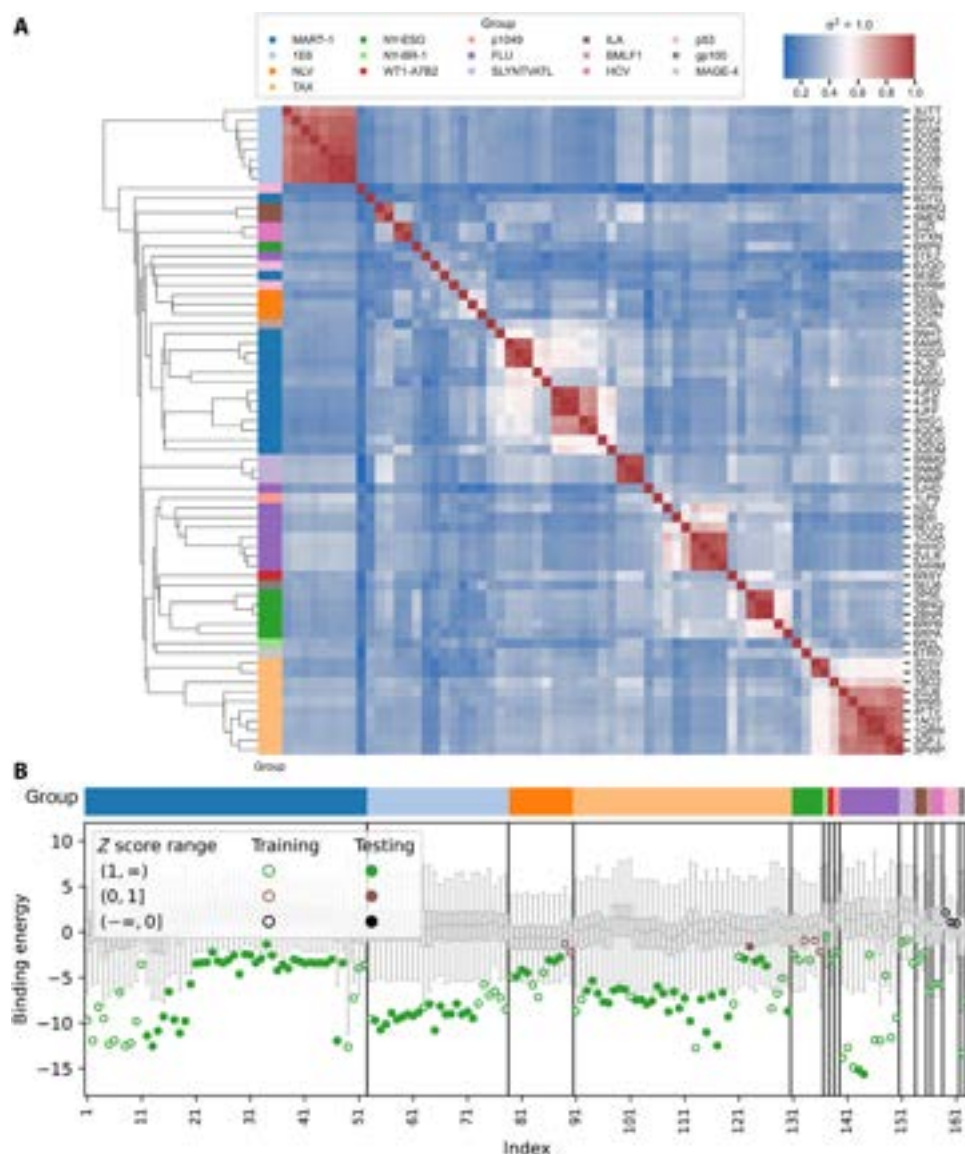


Fig. 2. Performance on ATLAS dataset. (A) Mutual Q calculation results between all crystal structures in training set of RACER-m, which measures the structural similarity between every pair of structures from the training set. The linkage map shows the hierarchical clustering result based on the pairwise mutual Q values. Color blocks next to the linkage map indicates the corresponding cluster of the crystal structure in the row. (B) Predicted binding energies for ATLAS dataset (open circles and closed dots) in comparison with the binding energies for corresponding weak binders (box plots). Each open circle represents the predicted binding energy for a structure in the training set, while each closed dot represents the predicted binding energy for a testing case from ATLAS dataset. Each training/testing case is associated with 1000 decoy weak binders generated by randomizing the peptide sequence and pairing with the TCR in the corresponding training/testing structure. Box plots represent the distribution of the predicted energies of the decoy weak binders with the box representing the lower (Q1) to upper (Q3) quartiles and a horizontal line representing the median. The whiskers extended from the box by 1.5 IQR, where IQR indicates the interquartile range.

binding affinity, consistent with earlier predictions (6, 7). This simultaneous manifestation of global sparsity and local resolvability among TCR-peptide pairs with identical HLA restriction represents a dual challenge for the development of robust predictive models of TCR-peptide specificity.

Given the inter-cluster structural diversity for TCR-pMHC complexes as well as the intra-cluster variability, it is necessary to suitably select a list of structures with sufficient coverage of the identified structural clusters as training data for the model and structural templates for test cases. In particular, we hypothesized that our hybrid structural and sequence-based methodology could benefit from the inclusion of multiple template structures, and the modeling approach presented here was developed with this motivation in mind.

The flow chart in Fig. 1 illustrates the training (top row) and testing (bottom row) algorithm in RACER-m. For training, contact interactions between peptide and TCR were calculated for each of the strong binding pairs with available TCR-pMHC crystal structures. Here, contact interactions were defined by a switching function based on the distance between structural residues and a characteristic interaction length (see Methods). For each strong binder, 1000 decoy (weak binder) sequences were generated by pairing the original TCR with a randomized version of the peptide. Contact interactions derived from the topology of known TCR-pMHC structures, together with a pairwise 20-by-20 symmetric amino acid energy matrix, determine total binding energy. Each value of the energy matrix corresponds to a particular contribution by an amino acid combination, with negative numbers corresponding to attractive contacts. The training objective aims to select the energy matrix that maximizes separability between the binding energy distributions of strong and weak binders.

In the testing phase, a sequence threading method is used to construct 3D structures for testing cases that lack a solved crystal structure. Here, constructed structures are based on using a chosen known template with shortest (CDR3 α/β and peptide) sequence distance to the specific testing case. Using the constructed 3D structure, a contact interface can be similarly calculated for each testing case, and 1000 decoy weak binders can be generated by randomizing the peptide sequence. The optimal energy model is then applied to assign energies to the testing TCR-pMHC pair and decoy binders, and the testing pair is identified as a strong binder if its predicted binding energy is substantially lower than the decoy energy distributions based on a standardized z score. Here, z score calculation was adopted from the statistical z test applied to the predicted binding energy of test TCR-pMHC pairs and decoy weak binders, the latter of which were used as a null distribution to compare against a given test binder. The z score of binding energies is defined as $z = (\bar{E}_{\text{decoy}} - E_{\text{test}}) / \sigma_{\text{decoy}}$, where \bar{E}_{decoy} is the average predicted binding energy of decoy weak binders, E_{test} is the predicted binding energy of the testing TCR-pMHC pair, and σ_{decoy} is the standard deviation (SD) of the binding energies of decoy weak binders. While model output is composed of continuous values of energy (or normalized z score), we consider test TCR-pMHC pairs with z scores exceeding 1 to be strong binding for categorization purposes.

Structural information enhances recognition specificity of pMHC-TCR complexes

RACER-m was developed to explicitly leverage the available structural information obtained from experimentally determined TCR-pMHC complexes for test predictions. While a prior modeling effort

(14) relied on a single structural template for both training and testing and achieved reasonable results given reduced training data, structural differences became prominent as the testing data expanded to include additional TCR and peptide diversity, which resulted in reduced predictive utility. Structural variation has been previously observed and quantified in high molecular detail (22, 27) using docking angles (28) and interface parameters.

For HLA-A*02:01 TCR-pMHC systems, the docking angles (between the peptide binding groove on the MHC and the vector between the TCR domains, which corresponds to the twist of the TCR over the pMHC) ranged from 29° to 73.1°, while the incident angle varied from 0.3° to 39.5° (22, 27, 29). The observed structural differences among different TCR-pMHC complexes suggest that a single TCR-pMHC complex structure may not accurately represent the contact interfaces of other TCR-pMHC complexes, particularly those with substantially different docking orientations. These distinct docking orientations lead to large variations in the contact interfaces between peptide and CDR3 α/β loops, which can be observed from the diversity in contact maps as shown in fig. S1. RACER-m overcomes this limitation by the inclusion of 66 TCR-pMHC crystal structures, which are distributed over distinct structural groups, including MART-1, 1E6, TAX, native Cytomegalovirus (NLV), and influenza (FLU) and serve as both the training dataset and reference template structures for testing cases.

In testing TCR-peptide pairs, all corresponding crystal structures were omitted from predictions. Thus, selecting an appropriate template from available structures became crucial for accurately reconstructing the TCR-pMHC interface and estimating the binding energy. To accomplish this, RACER-m assumed that high sequence similarity corresponds to high similarities in the structure space, which is supported by the correlation between mutual Q score and sequence similarity measured from the 66 solved crystal structures of TCR-pMHC complexes (fig. S2). This assumption was implemented by calculating sequence similarity scores of the testing peptide and TCR CDR3 α/β sequences with those of all 66 reference templates. In each case, a position-wise uniform hamming distance on amino acid sequences was calculated to quantify the similarity. The sum of CDR3 α and CDR3 β similarities generated the TCR similarity score, and a composite score was created by taking the product of peptide and TCR scores (see Methods). The template structure having the highest sequence similarity was then selected as the template for threading the sequences of the testing TCR-peptide pair.

To evaluate the extent to which the RACER-m approach can address global sparsity by accurately recapitulating observed specificity in the setting of limited training data, we trained a model using 42.3% of the total experimentally confirmed strong binders [in addition to the 66 HLA-A*02:01 TCR-pMHC crystal structures plus structures with PDB ID 3GSR, 3GSU, and 3GSV for NLV peptide strong binders (30)] which sparsely cover all the structural groups involved in the mutual Q analysis shown in Fig. 2A. The remaining 57.7% of TCR-peptide sequences that lack solved structures were used as testing cases to validate the sensitivity of the trained energy model. RACER-m effectively recognizes strong binding peptide-TCR pairs and correctly predicts 98.9% of the testing TCR-pMHC pairs using the criteria that z score is greater than 1. Among the 94 testing pairs, only one TCR-peptide pair in the TAX structural group was mis-predicted as a weak binder with a binding energy deviating from the average binding energies of decoy weak binders by 0.64σ , where σ is the SD of the decoy energies. These initial results

(Fig. 2) confirm that the model is effectively able to learn the specificity rules from TCR-pMHC pairs exhibiting distinct structural representations. Moreover, RACER-m computes a continuous value capable of illustrating differences in the relative binding affinities within functional TCR-peptide clusters (fig. S3).

While the reliable identification of strong-binding TCR-pMHC pairs is clinically useful and one important measure of model performance, simultaneous evaluation of model specificity is equally crucial for generating useful predictions on the level of a TCR repertoire. To evaluate the specificity of a global sparsity task, we next tested RACER-m's ability to discern experimentally confirmed weak-binding TCR-pMHC pairs. We selected peptides or TCRs from the most abundant structural groups (MART-1 and TAX) in the training set to create "scrambled" TCR-pMHC pairs by cross-cluster mismatching of either TCRs or peptides (see Methods for full details). Proceeding in this manner enables a specificity test on biologically realized sequences instead of randomly generated ones. Specifically, every peptide selected from a given structural group (e.g., peptide EAAGIGILTV in the MART-1 group) was mismatched with a list of TCRs specific for peptides belonging to other groups (e.g., TAX, 1E6, and FLU) to form a set of scrambled weak binders.

Following our aforementioned testing protocols, we next calculated z scores for these mismatched interactions, which were then compared to correctly matched TCR-pMHC pairs with the same peptide sequence (e.g., EAAGIGILTV). We also conducted the complementary test on TCRs using scrambled peptides. The primary advantages of this approach include (i) the ability to match amino acid empirical distributions in binding and nonbinding pairs and (ii) utilization of realized TCR sequences for specificity assessment instead of random sequences that have minimal, if any, overlap with physiological sequences.

A representative example of these tests using the MART-1 epitope and MART-1-specific TCRs is given in Fig. 3. First, seven sets of weak binders were constructed by mismatching 36 MART-1-specific TCRs each with seven non-MART-1 peptides sampled from distinct clusters. We applied RACER-m on each weak binder to predict its binding energy and then compared this value to the distribution of decoy binding energies to obtain a binding Z score. Z scores of mismatched weak binders, together with those of

correctly matched MART-1-TCR strong binders, were used to derive the receiver operating characteristic (ROC) curve (Fig. 3A and fig. S4). The area under the curve (AUC) was greater than or equal to 0.98 for five of the seven test sets, while the others had AUCs of 0.80 and 0.75, illustrating RACER-m's ability to successfully distinguish strong binding peptides from mismatched ones in the available MART-1-specific TCR cases.

An analogous test was performed on the five available peptide variants from the MART-1 structural group by mismatching them with 35 TCR sequences contained in the NLV, FLU, 1E6, or TAX clusters. Relative to the binding energies of correctly matched MART-1-specific TCRs, RACER-m performs well in discerning matched versus mismatched TCRs for four of the five tested MART-1 peptides (Fig. 3B and fig. S5), the one initial exception being peptide ELAGIGILTV. Further inspection of the TCRs in this group revealed that the TAX-specific TCR A6 (triangle sign in Fig. 3C) together with several closely associated point mutants had a z score distribution resembling that of the RD1-MART1^{High} TCR and its associated point mutants (fig. S5E). This could be explained by the fact that the RD1-MART1^{High} TCR was engineered from the A6 TCR to achieve MART-1 specificity (31), wherein A6 was selected because of its similarity with MART-1-specific TCRs in the V α region and similar docking mode (16, 31). However, the engineered (RD1-MART1^{High}) TCR is no longer specific to the TAX peptide (LLFGYPVYV), which is consistent with the z scores predicted from RACER-m. When the A6-specific TAX peptide is paired with RD1-MART1^{High} TCR, a relatively lower z score (cross sign in Fig. 3C) is predicted in comparison with the z scores from strong binders (violin shape in Fig. 3C) of the same peptide.

Evaluation on extended datasets highlights the added value of structural information

Given RACER-m's performance on the ATLAS data, we then applied the model to additional datasets to further validate its ability in the setting of global sparsity. The 10x Genomics (32) dataset details many TCR-peptide binders collected from five healthy donors. HLA-A*02:01-restricted samples in this dataset include 23 unique peptides, and the number of TCRs specific for each peptide varied from 8365 (e.g., GILGFVFTL) to 1 (e.g., ILKEPVHGV). We remark

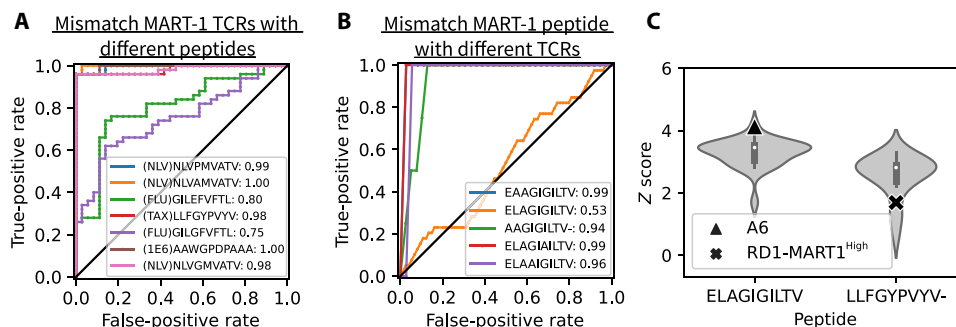


Fig. 3. Prediction performance on weak binders generated by mismatching peptides with TCRs. (A) Receiver operating characteristic (ROC) curves for RACER-m classification performance on differentiating weak binders generated by mismatching peptides from NLV, TAX, FLU, and 1E6 clusters with MART-1 TCRs from MART-1 strong binders with the same set of TCRs. (B) ROC curves for RACER-m classification performance on distinguishing MART-1 strong binders from mismatched weak binders generated by pairing MART-1-specific peptides with TCRs from NLV, TAX, FLU, and 1E6 clusters. (C) When TAX A6 TCR is paired with MART-1 peptide ELAGIGILTV, the z score of the mismatched TCR-pMHC pair (triangle) resembles the values from the strong binders (violin shape) formed by the same peptide and TCR RD1-MART1^{High} and its point mutants, which was engineered from A6. In the reverse scenario, TCR RD1-MART1^{High} shows lower z score (cross) than TAX strong binders (violin shape) when paired with TAX specific peptide LLFGYPVYV.

that the diversity of HLA-A*02:01 samples was substantially reduced to 1741 TCR-pMHC pairs having unique CDR3 α/β and peptide sequences after removing redundancies. We selected this large dataset as a reasonable test because 89.26% of the 1741 testing pairs did not share either the same CDR3 α or CDR3 β sequence in common with the list of available TCR-pMHC pairs used in the training set, and 99.89% of the testing TCR-pMHC pairs did not have the same CDR3 α -CDR3 β combination with the training set, although 7 of the 23 peptides were shared with the training set.

Given this relative lack of overlap with our training data, we applied RACER-m to all unique HLA-A*02:01 pairs. In a majority (88.9%) of these cases across a large immunological diversity of peptides, RACER-m successfully identifies enriched z scores in the distribution of binding TCRs (Fig. 4A). The distinction of TCRs belonging to testing versus training sets, together with the notable difference in the size of training and testing TCR-pMHC pairs, suggests that shared structural features were able to augment RACER-m's predictive power on distinct tests. Thus, the inclusion of structural information in model training enhances RACER-m's predictive ability across distinct TCR-pMHC tests.

There were several cases where RACER-m's predicted distributions overlapped substantially with low z scores, indicating a failed

prediction; in these cases, we investigated whether this could be explained by the lack of an appropriate structural template. A positive correlation was observed between a testing case's optimal structural template similarity and the RACER-m-predicted z scores, consistent with a decline in model applicability whenever the closest available template is inadequate for representing the TCR-pMHC pair in question (fig. S6). Despite this, the RACER-m approach, trained on 69 cases, was able to predict roughly 90% of strong binders contained in over 1700 distinct testing cases in the 10x Genomics dataset. A similar trend was also seen when applying RACER-m to the "global true" test set curated from the VDJdb (33) that were not included in training. RACER-m again exhibited optimal predictive performance when a reasonable structural template was available (figs. S7 and S8). Overall, RACER-m was able to successfully predict 56.7% of the strong binders in this set. For groups with high sequence similarities with our template structures, such as the cases of peptide "GILGFVFTL," RACER-m yields a higher success rate of strong binder prediction (91.1% for cases with peptide "GILGFVFTL").

We then compared RACER-m's performance to NetTCR-2.0 (11), a well-established convolutional neural network model for predictions of TCR-peptide binding that is trained on over 16,000 combinations of peptide/CDR3 α/β sequences. This comparison was

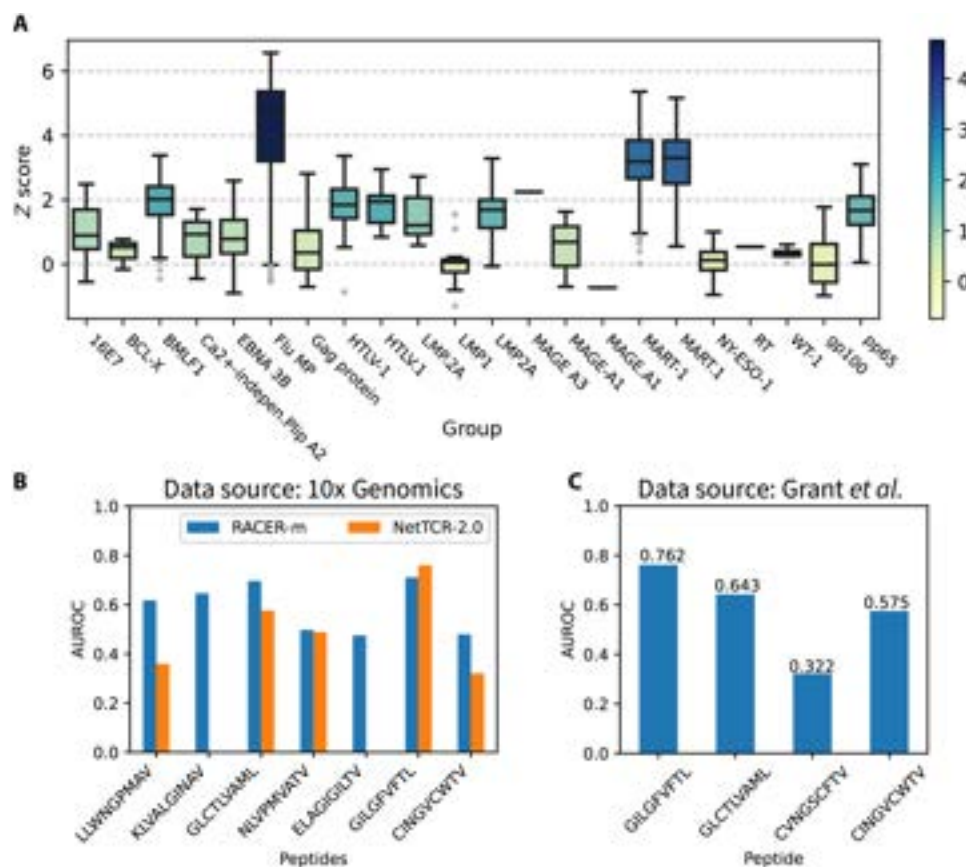


Fig. 4. Validate the predictive power of RACER-m with external datasets. (A) Prediction results of RACER-m on the HLA-A*02:01-restricted systems from 10x Genomics dataset collected from five healthy donors. A total of 1741 unique pairs of TCR-peptide sequences were tested, and the prediction results of z score were grouped by the immunological profile of the test TCR-pMHC pairs and depicted as box plots. (B) Comparison of classification performance between RACER-m and NetTCR-2.0 (11) on a curated list of public TCR-pMHC repertoires (12, 42) composed of both strong binders and mismatched weak binder. Because of the restriction of NetTCR-2.0 on the peptide length (9-mer), there are no data from NetTCR-2.0 for the two 10-mer peptides (KLVALGINAV and ELAGIGILTV). (C) The classification performance of RACER-m on another set of TCR-pMHC test TCR-pMHC pairs (34). AUROC, area under the ROC curve.

performed on a publicly available list of TCR-pMHC repertoires curated by Zhang *et al.* (12), which were mutually independent of RACER-m or NetTCR-2.0 training data, wherein we included known strong binders and mismatched weak binders for eight unique peptides of HLA-A*02:01. Because NetTCR-2.0 has a restricted length for antigen peptide (no longer than 9-mer), it cannot be applied on testing TCR-pMHC pairs with 10-mer peptides such as KLVALGINAV and ELAGIGILTV, which are absent from the NetTCR-2.0 evaluation in Fig. 4B. The area under the ROC curve was used as a standard measure of classification success. In the majority of cases, RACER-m outperformed NetTCR-2.0 in diagnostic accuracy with higher ROC values (Fig. 4B). Last, RACER-m was further evaluated using an unrelated set of TCR-pMHC data composed of 400 samples made up of the strong binders and mismatched weak binders with four peptides and 100 TCRs (34), which also gives us good distributional performance (Fig. 4C). In one of the four peptides included in this dataset, RACER-m seems to have difficulty providing correct classification about strong and weak binders for peptide CVNGSCFTV, which could again be explained by the lack of appropriate structure templates for this pMHC and related strong binding TCRs (fig. S9).

RACER-m specificity of point-mutated variants and preservation of local resolvability

Encouraged by model handling of global sparsity in tests of disparate binding TCR-pMHC pairs having high sequence diversity, we next evaluated RACER-m's ability in maintaining local resolvability of point-mutated peptides with near-identical sequence similarity to a known strong binder, which represents a distinct and usually more difficult computational problem. Understanding in detail which available point mutants enhance or break immunogenicity is directly relevant for assessing the efficacy of tumor neoantigens and T cell responses to viral evolution. In addition, the performance of structural models in accomplishing this task are a direct readout on their utility over sequence-based methods because the latter case will struggle to accurately cluster and, therefore, resolve TCR-pMHC pairs having single-amino acid differences. To evaluate RACER-m's ability to recognize point mutants, we performed an additional test on an independent comprehensive dataset of TCR 1E6 containing a point mutagenic screening of the peptide displayed on MHC. This testing set includes 20 strong binders and 73 weak binders (21), wherein strong binding to the 1E6 TCR was confirmed by tumor necrosis factor- α activity. RACER-m demonstrates enrichment of

the distribution of binding energies for strong binders versus confirmed weak cases (Fig. 5A). ROC analysis of the RACER-m's ability to resolve these groups gives an AUC of 0.78. Note that only two strong binders of this group were included in the training of RACER-m's energy model.

Inspired by these initial results on the 1E6 mutagenic screen, we extended this analysis to all point-mutated weak binding TCR-pMHC pairs in the ATLAS dataset, specifically those with K_d values greater than 200 μ M. Our results, presented template-wise for each structure in the point-mutant data, demonstrate that RACER-m improves in this recognition task when compared to NetTCR-2.0 (Fig. 5C). Last, to explicitly explore the value of structural modeling for predicting the impact of immunologically important single-amino acid differences, we quantified the predicted z scores for both strong and weak binders based on a measure of total sequence similarity (fig. S10). This measure was obtained by taking the maximum product of CDR3 α , CDR3 β , and peptide Hamming similarity between a test TCR-peptide pair and each of the training TCR-peptide pairs with an available structure. The results demonstrate that the inclusion of information from correctly identified structural templates enhances RACER-m's predictive power. Collectively, our results suggest that RACER-m offers a unique computational advantage over traditional, sequence-only methods of prediction by leveraging substantially fewer training sequences with key structural information to efficiently identify the contribution of each amino acid change.

DISCUSSION

Reliable and efficient estimation of TCR-pMHC interactions is of central importance in understanding and thus optimizing the adaptive immune response. The field has experienced considerable recent research activity in the development of inference-based computational methods to predict TCR-pMHC specificity (35). Decoding the predictive rules of TCR-pMHC specificity is a formidable challenge, largely owing to the extreme sparsity of available training data relative to the diversity of sequences that need to be interrogated in meaningful investigation. A majority of approaches (11, 36, 37) take a complementary approach to RACER-m by training on TCR and/or peptide primary sequence data alone. One recent method achieves training by relaxing a common requirement of having paired CDR3 α/β sequences (36). We developed RACER-m to augment the predictive power of a relatively small number of

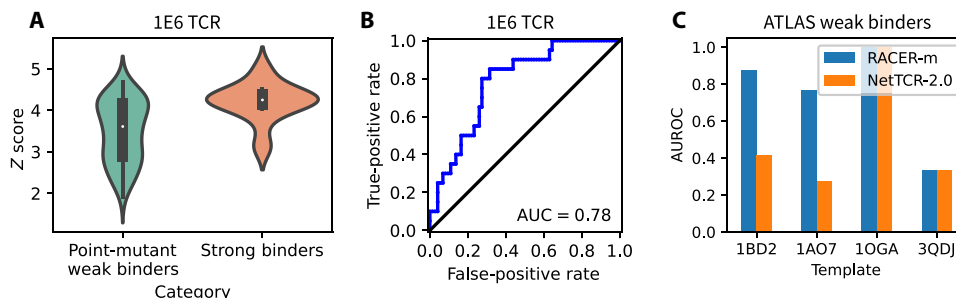


Fig. 5. RACER-m's performance on differentiating strong binders from point-mutant weak binders. (A) Distribution of z scores from strong binders of 1E6 TCR and weak binders from point mutagenic screen. (B) ROC curve for RACER-m classification performance using the strong and point-mutant weak binders for 1E6 TCR. (C) Comparison of RACER-m and NetTCR-2.0 in classification of strong and point-mutant weak binders from ATLAS dataset. Here, RACER-m predictions used the known crystal structure selected by the sequence similarity calculation results as a representative template for threading each test case.

TCR and epitope sequences by leveraging the structural information contained in solved TCR-pMHC crystal structures. Our analysis focused on the most common human MHC allele variant, due to the abundance of sequence and structural data. Despite this restriction, we observed structural heterogeneity underpinning the specificity of various TCR-pMHC pairs in distinct immunological contexts. Enhancement in predictive accuracy was largely driven by the availability of a small list of structural templates, which included 66 crystal structures of TCR-pMHC complexes from the PDB.

Using our minimal list, together with mutually independent testing TCR-pMHC pairs for RACER-m and NetTCR-2.0, we find that our model is able to outperform NetTCR-2.0 on both detection of strong binders as well as avoidance of weak binders, both representing distinct but equally important tasks. We advocate for the inclusion of such mixed performative tests for rigorous validation as a necessary and standardized component in model evaluation, in addition to model comparisons using testing data that are equally dissimilar from the training data included in competing models.

Intriguingly, incorporation of structural information into the training approach enables the development of a model that maintains predictive accuracy while dealing with both global sparsity and local resolvability, all while requiring substantially reduced training sequence data. Because of RACER-m's ability to deal with both global sparsity and local resolvability, we anticipate that this approach may be applicable to future applications that require reliable predictions on TCR responses against disparate and closely related collections of antigens. Such an approach may provide a useful theoretical tool to design, for example, tumor antigen vaccines. Our results suggest that a wealth of information is contained in the structural templates pertaining to key contributors of a favorable TCR-peptide interaction, wherein conserved features across distinct TCR-pMHC pairs can be learned to mitigate global sparsity. Conversely, structural encoding of information pertinent to residues whose amino acid substitutions either preserve or break immunogenicity also assists RACER-m trained on only a small subset of all possible point mutations by identifying key contributing positions and residues, thereby preserving local resolvability.

Our current approach has been successfully applied to resolve unknown strong and weak binding TCR-pMHC pairs given those identified as such in the previously published test datasets under consideration. We note that perfect resolvability in the setting of repertoire-level studies that assess large numbers of randomly sampled TCR and peptide pairs would require larger z scores for distinguishing strong binders. In several test cases, our model does assign strong binders a larger score ($z = 4$; Fig. 4A; FLU and MART-1), especially when sufficient positive training data exist. We also note that some tasks (for example, picking out single-amino acid mutants that retain strong binding) do not require competing against a large number of possible choices, and so the needed z score should be much lower.

Moreover, model accuracy correlated directly with the availability of a template having sufficient proximity to the sequences of testing TCR-pMHC pairs. As a result, we anticipate that RACER-m will improve as more structures become readily available for inclusion. Existing computational methods for identifying structural models from primary sequence data (18) may provide an efficient method of adding highly informative structures into the candidate pool for testing. This task, together with identifying the minimal sufficient number of distinct structural classes within a given MHC allele

restriction, remains for subsequent investigation. Our current results suggest that this is doable given the small number of structures available for explaining the diverse TCR-pMHC pairs studied herein. Notably, the inclusion of only 66 template structure augmented RACER-m's ability to accurately differentiate strong and weak binders when evaluated with hundreds and even thousands of testing TCR-pMHC pairs. This structural advantage was enhanced both by the approach of hybridizing sequence and structural information into the training and testing protocols and the availability of templates that shared sufficient sequence-based similarity to testing cases so that an adequate threading template was available.

METHODS

RACER-m model

To predict the binding affinity between a given TCR-peptide pair, we used a pairwise energy model to assess the TCR-peptide binding energy (14). The CDR3 α and CDR3 β regions were used to differentiate between different TCRs because CDR3 loops primarily interact with the antigen peptides, while CDR1 and CDR2 interact with MHC (38). However, the binding energy was evaluated on the basis of the entire binding interface between TCR and peptide. As illustrated in Fig. 1, we included 66 experimentally determined TCR-pMHC complex structures and three additional TCR-pMHC complex structures composed of experimentally determined pMHC complexes with corresponding TCR structures as strong binders for training an energy model (details in Supporting Methods), which was subsequently used to evaluate binding energies of other TCR-peptide pairs based on their CDR3 and peptide sequences. In addition, for each strong binder, we generated 1000 decoy binders by randomizing the peptide sequence. These 69,000 decoys constitute an ensemble of weak binders within our training set.

To parameterize this energy model, we optimized the parameters by maximizing the gap of binding energies between the strong and weak TCR-peptide binders, represented by δE in Fig. 1. The resulting optimized energy model will be used for predicting the binding specificity of a peptide toward a given TCR based on their sequences. Further details regarding the calculation of binding energy are provided below.

Detailed calculation of TCR-peptide binding energies

To evaluate the binding affinity between a TCR and a peptide, RACER-m used the framework of the AWSEM force field (39), which is a residue-resolution protein force field widely used for studying protein folding and binding (39, 40). To adapt the AWSEM force field for predicting TCR-peptide binding energy, we used its direct protein-protein interaction component to calculate the inter-residue contacting interactions at the TCR-peptide interface. Specifically, we used the C β atoms (except for glycine, where C α atom was used instead) of each residue to calculate the contacting energy using the following expression

$$V_{\text{direct}} = \sum_{i \in \text{TCR}, j \in \text{peptide}} \gamma_{ij}(a_i, a_j) \Theta_{ij}^I \quad (1)$$

In Eq. 1, Θ_{ij}^I represents a switching function that defines the effective range of interactions between each amino acid from the peptide and the TCR

$$\Theta_{ij}^I = \frac{1}{4} \{1 + \tanh[5.0 \times (r_{ij} - r_{\min}^I)]\} \{1 + \tanh[5.0 \times (r_{\max}^I - r_{ij})]\} \quad (2)$$

where $r_{\min}^I = 6.5 \text{ \AA}$ and $r_{\max}^I = 8.5 \text{ \AA}$. The coefficients $\gamma_{ij}(a_i, a_j)$ define the strength of interactions based on the types of amino acids (a_i, a_j). The $\gamma_{ij}(a_i, a_j)$ coefficients are also the parameters that are trained in the optimization protocols described as follows.

Optimization of energy model for predicting the TCR-peptide binding specificity

To predict the binding specificity between a given TCR and peptide, the energy model is trained using interactions gathered from the known strong binders and their corresponding randomly generated decoy binders.

Following the protocol specified in our previous paper (14), the energy model of RACER-m was trained to maximize the gap between the binding energies of strong and weak binders. In addition, a larger training set was used to achieve a more comprehensive coverage of the structural and sequence space. Specifically, the binding energies were calculated from individual strong binders (E_{strong}) and their corresponding decoy weak binders (E_{decoy}) as described in Eq. 1. We then calculated the average binding energy of the strong ($\langle E_{\text{strong}} \rangle$), the average binding energy of the decoy weak binders ($\langle E_{\text{decoy}} \rangle$), and the SD of the energies of the decoy weak binders (ΔE).

To train the model, the parameters $\gamma_{ij}(a_i, a_j)$ were optimized to maximize $\delta E / \Delta E$, where $\delta E = \langle E_{\text{decoy}} \rangle - \langle E_{\text{strong}} \rangle$, resulting in the maximal separation between strong and weak binders. Mathematically, δE can be represented as $\mathbf{A}^T \boldsymbol{\gamma}$, where

$$\mathbf{A} = \langle \phi_{\text{decoy}} \rangle - \langle \phi_{\text{strong}} \rangle \quad (3)$$

Furthermore, the SD of the decoy binding energies ΔE can be calculated as $\Delta E^2 = \boldsymbol{\gamma}^T \mathbf{B} \boldsymbol{\gamma}$, where

$$\mathbf{B} = \langle \phi_{\text{decoy}} \phi_{\text{decoy}}^T \rangle - \langle \phi_{\text{decoy}} \rangle \langle \phi_{\text{decoy}} \rangle^T \quad (4)$$

here, ϕ takes the functional form of V_{direct} and summarizes interactions between different types of amino acids. Therefore, the vector \mathbf{A} specifies the difference in interaction strengths for each pair of amino acid types between the strong and decoy binders, with a dimension of (1, 210), while the matrix \mathbf{B} is a covariance matrix with a dimension of (210, 210).

With the definition above, maximizing the objective function of $\delta E / \Delta E$ can be reformulated as maximization of $\mathbf{A}^T \boldsymbol{\gamma} / \sqrt{\boldsymbol{\gamma}^T \mathbf{B} \boldsymbol{\gamma}}$. This maximization can be effectively achieved through maximizing the functional objective $R(\boldsymbol{\gamma}) = \mathbf{A}^T \boldsymbol{\gamma} - \lambda_1 \sqrt{\boldsymbol{\gamma}^T \mathbf{B} \boldsymbol{\gamma}}$. By setting $\partial R(\boldsymbol{\gamma}) / \partial \boldsymbol{\gamma}^T$ to 0, the optimization process leads to $\boldsymbol{\gamma} \propto \mathbf{B}^{-1} \mathbf{A}$, where $\boldsymbol{\gamma}$ is a (210, 1) vector encoding the trained strength of each type of amino acid-amino acid interactions. For visualization purposes, the vector $\boldsymbol{\gamma}$ is reshaped into a symmetric 20-by-20 matrix, as shown in Fig. 1. In addition, a filter is applied to reduce the noise caused by the finite sampling of decoy binders. In this filter, the first 50 eigenvalues of the \mathbf{B} matrix are retained, and the remaining eigenvalues are replaced with the 50th eigenvalue.

Construction of target TCR-pMHC complex structures from sequences

Because RACER-m calculates the binding energy based on the interaction contacts between a given peptide and a TCR, it relies on the 3D structure of the TCR-pMHC complex for contact calculation. Although the training data include a 3D structure for each of the TCR-peptide strong binders, we usually lack 3D structures for

most of the testing cases. To address this limitation, we used the software MODELLER (41) to construct a structure based on the target peptide/CDR3 sequences in the test TCR-pMHC pair and a template crystal structure selected from the training set.

Specifically, for each testing TCR-pMHC pair, a position-wise uniform Hamming distance was computed between the target sequence and each of the sequences from the 66 training strong binders with complete TCR-pMHC complex structures, separately for peptide, CDR3 α , and CDR3 β regions. Then, sequence similarity scores were assigned to peptide, CDR3 α , and CDR3 β , respectively, with the number of amino acids that remain the same between target and template sequences. To calculate a composite similarity score for the target TCR-peptide complex, we summed the similarity scores of the CDR3 α and CDR3 β regions and multiplied this sum by the peptide similarity score. The template structure with the highest similarity score was selected as the template for the subsequent sequence replacement using MODELLER (Fig. 1, bottom).

To perform the sequence replacement, the peptide, CDR3 α , and CDR3 β sequences in the template structure were replaced with the corresponding target sequences in the testing TCR-peptide pair. The location of the target sequence in the template structure was determined by aligning the first amino acid of the target sequence with the original template sequence. If the two sequences had different lengths, then the remaining locations were patched with gaps. This sequence alignment and the selected template structure were then used as input for MODELLER to generate a new structure. The constructed structure was then used for the estimation of the binding energy of the testing TCR-pMHC pair.

Generation of weak binders by mismatching sequences of known TCR-peptide pairs

To test the performance of RACER-m in distinguishing strongly bound TCR-peptide pairs from weak binders, we generated a set of weak binders by introducing sequence mismatches between the peptides and TCRs from the known strongly bound TCR-peptide pairs. As shown in Fig. 2, the strong binders were grouped on the basis of their immunological systems, such as MART-1 and TAX. Note that pairs within the same group also share similar TCR-peptide structural interfaces.

To generate the weak binders, we mismatched the sequences of peptides and the CDR3 α/β pairs from different groups. For example, 36 pairs of MART-1-specific CDR3 α/β sequences were mismatched with seven non-MART-1 peptides to form weak binders for Fig. 3A, while five MART-1-specific peptides were mismatched with 35 pairs of non-MART-1 CDR3 α/β sequences to form weak binders in Fig. 3B. The newly generated combinations of sequences were then used to create 3D structures of the TCR-pMHC complexes, following the protocol specified in the ‘‘Construction of target TCR-pMHC complex structures from sequences’’ section.

Mutual Q calculation

To quantify the structural distances between the 66 crystal structures of TCR-pMHC complexes, a pairwise mutual Q score was used to calculate the structural similarity between every pair of the 66 structures. Because our focus is on the contact interface between the peptide and the CDR3 α /CDR3 β loops of the TCR, the mutual Q score was computed between these regions. We adopted a similar protocol used in (25) and calculated the mutual Q score between structures A and B with the following expression

$$Q^{A,B} = c \sum_{i \in \text{peptide}, j \in \text{CDR3}} \exp \left[- \frac{(r_{ij}^A - r_{ij}^B)^2}{2\sigma^2} \right] \quad (5)$$

where i and j are indices of atoms from the peptide and CDR3 loops, respectively. r_{ij}^A and r_{ij}^B denote the contact distances between atom i and j in structure A and B, respectively. For simplicity, σ was set as 1 Å instead of using the sequence distance between i and j as done in (25). The coefficient c normalizes the value of Q to fall within the range of 0 and 1. This definition ensures that a larger value of Q indicates a greater structural similarity between the two TCR-pMHC pairs.

Prediction protocols with NetTCR-2.0

To test the predictive performance of RACER-m, we compared the prediction accuracy of RACER-m with NetTCR-2.0, another widely used computational tool trained with a convolutional neural network model, as described by Montemurro *et al.* (11). To ensure a fair comparison, we retrained the NetTCR-2.0 model with the dataset with paired α/β TCR CDR3 regions and a 95% partitioning threshold (file train_ab_95_alpha_beta.csv, provided in <https://github.com/mmielab/NetTCR-2.0>). The trained model was then used to classify the strong and weak binders, as shown in Fig. 5C. Because of the peptide length restriction in the application of NetTCR-2.0, we excluded peptides longer than nine residues from our testing prediction.

Supplementary Materials

This PDF file includes:

Supporting Methods
Figs. S1 to S10
References

REFERENCES AND NOTES

- L. Klein, B. Kyewski, P. M. Allen, K. A. Hogquist, Positive and negative selection of the T cell repertoire: What thymocytes see (and don't see). *Nat. Rev. Immunol.* **14**, 377–391 (2014).
- A. Grakoui, S. K. Bromley, C. Sumen, M. M. Davis, A. S. Shaw, P. M. Allen, M. L. Dustin, The immunological synapse: A molecular machine controlling T cell activation. *Science* **285**, 221–227 (1999).
- S. Ilyas, J. C. Yang, Landscape of tumor antigens in T cell immunotherapy. *J. Immunol.* **195**, 5117–5122 (2015).
- A. Košmrlj, A. K. Jha, E. S. Huseby, M. Kardar, A. K. Chakraborty, How the thymus designs antigen-specific and self-tolerant T cell receptor sequences. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 16671–16676 (2008).
- A. K. Chakraborty, A. Košmrlj, Statistical mechanical concepts in immunology. *Annu. Rev. Phys. Chem.* **61**, 283–303 (2010).
- J. T. George, D. A. Kessler, H. Levine, Effects of thymic selection on T cell recognition of foreign and tumor antigenic peptides. *Proc. Natl. Acad. Sci. U.S.A.* **114**, E7875–E7881 (2017).
- K. N. Chau, J. T. George, J. N. Onuchic, X. Lin, H. Levine, Contact map dependence of a T-cell receptor binding repertoire. *Phys. Rev. E* **106**, 014406 (2022).
- M. E. Birnbaum, J. L. Mendoza, D. K. Sethi, S. Dong, J. Glanville, J. Dobbins, E. Özkan, M. M. Davis, K. W. Wucherpfennig, K. Christopher Garcia, Deconstructing the peptide-MHC specificity of T cell recognition. *Cell* **157**, 1073–1087 (2014).
- P. Dash, A. J. Fiore-Gartland, T. Hertz, G. C. Wang, S. Sharma, A. Souquet, J. C. Crawford, E. Bridie Clemens, T. H. Nguyen, K. Kedzierska, Quantifiable predictive features define epitope-specific T cell receptor repertoires. *Nature* **547**, 89–93 (2017).
- S. C. Boutet, D. Walter, M. J. Stubbington, K. A. Pfeiffer, J. Y. Lee, S. E. B. Taylor, L. Montesclaros, J. K. Lau, D. P. Riordan, A. M. Barrio, L. Brix, K. Jacobsen, B. Yeung, X. Zhao, T. S. Mikkelsen, Scalable and comprehensive characterization of antigen-specific CD8 T cells using multi-omics single cell analysis. *J. Immunol.* **202**, 131.4 (2019).
- A. Montemurro, V. Schuster, H. R. Povlsen, A. K. Bentzen, V. Jurtz, W. D. Chronister, A. Crinklaw, S. R. Hadrup, O. Winther, B. Peters, NetTCR-2.0 enables accurate prediction of tcr-peptide binding by using paired tcr α and β sequence data. *Commun. Biol.* **4**, 1060 (2021).
- W. Zhang, P. G. Hawkins, J. He, N. T. Gupta, J. Liu, G. Choonoo, S. W. Jeong, C. R. Chen, A. Dhanik, M. Dillon, A framework for highly multiplexed dextramer mapping and prediction of T cell receptor sequences to antigen specificity. *Sci. Adv.* **7**, eabf5835 (2021).
- J.-W. Sidhom, H. B. Larman, D. M. Pardoll, A. S. Baras, DeepTCR is a deep learning framework for revealing sequence concepts within T-cell repertoires. *Nat. Commun.* **12**, 1605 (2021).
- X. Lin, J. T. George, N. P. Schafer, K. N. Chau, M. E. Birnbaum, C. Clementi, J. N. Onuchic, H. Levine, Rapid assessment of T-cell receptor specificity of the immune repertoire. *Nat. Comput. Sci.* **1**, 362–373 (2021).
- I. Springer, N. Tickotsky, Y. Louzoun, Contribution of T cell receptor alpha and beta CDR3, MHC typing, V and J genes to peptide binding prediction. *Front. Immunol.* **12**, 664514 (2021).
- B. G. Pierce, L. M. Hellman, M. Hossain, N. K. Singh, C. W. Vander Kooi, Z. Weng, B. M. Baker, Computational design of the affinity and specificity of a therapeutic T cell receptor. *PLoS Comput. Biol.* **10**, e1003478 (2014).
- C. Yanover, P. Bradley, Large-scale characterization of peptide-MHC binding landscapes with structural simulations. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 6981–6986 (2011).
- P. Bradley, Structure-based prediction of T cell receptor: Peptide-MHC interactions. *eLife* **12**, e82813 (2023).
- T. Borrmann, J. Cimons, M. Cosiano, M. Purcaro, B. G. Pierce, B. M. Baker, Z. Weng, ATLAS: A database linking binding affinities with structures for wild-type and mutant TCR-pMHC complexes. *Proteins* **85**, 908–916 (2017).
- C. Szeto, C. A. Lobos, A. T. Nguyen, S. Gras, TCR recognition of peptide-MHC-I: Rule makers and breakers. *Int. J. Mol. Sci.* **22**, 68 (2020).
- A. M. Bulek, D. K. Cole, A. Skowera, G. Dolton, S. Gras, F. Madura, A. Fuller, J. J. Miles, E. Gostick, D. A. Price, J. W. Drijfhout, R. R. Knight, G. C. Huang, N. Lissin, P. E. Molloy, L. Wooldridge, B. K. Jakobsen, J. Rossjohn, M. Peakman, P. J. Rizkallah, A. K. Sewell, Structural basis for the killing of human beta cells by CD8⁺ T cells in type 1 diabetes. *Nat. Immunol.* **13**, 283–289 (2012).
- R. Gowthaman, B. G. Pierce, TCR3d: The T cell receptor structural repertoire database. *Bioinformatics* **35**, 5323–5325 (2019).
- H. M. Berman, The Protein Data Bank. *Nucleic Acids Res.* **28**, 235–242 (2000).
- R. Vita, S. Mahajan, J. A. Overton, S. K. Dhanda, S. Martini, J. R. Cantrell, D. K. Wheeler, A. Sette, B. Peters, The Immune Epitope Database (IEDB): 2018 update. *Nucleic Acids Res.* **47**, D339–D343 (2018).
- M. Chen, X. Lin, W. Lu, J. N. Onuchic, P. G. Wolynes, Protein folding and structure prediction from the ground up II: AAWSEM for α/β Proteins. *J. Phys. Chem. B* **121**, 3473–3482 (2017).
- S. S. Cho, Y. Levy, P. G. Wolynes, P versus Q: Structural reaction coordinates capture protein folding on smooth landscapes. *Proc. Natl. Acad. Sci. U.S.A.* **103**, 586–591 (2006).
- R. Gowthaman, B. G. Pierce, Modeling and viewing T cell receptors using TCRmodel and TCR3d. *Methods Mol. Biol.* **2120**, 197–212 (2020).
- M. G. Rudolph, R. L. Stanfield, I. A. Wilson, How TCRs bind MHCs, peptides, and coreceptors. *Annu. Rev. Immunol.* **24**, 419–466 (2006).
- B. G. Pierce, Z. Weng, A flexible docking approach for prediction of T cell receptor-peptide-MHC complexes. *Protein Sci.* **22**, 35–46 (2013).
- S. Gras, X. Saulquin, J.-B. Reiser, E. Debeaupuis, K. Echasserieu, A. Kissenpfennig, F. Legoux, A. Chouquet, M. L. Gorrec, P. Machillot, B. Neveu, N. Thielens, B. Malissen, M. Bonneville, D. Housset, Structural bases for the affinity-driven selection of a public TCR against a dominant human cytomegalovirus epitope. *J. Immunol.* **183**, 430–437 (2009).
- S. N. Smith, Y. Wang, J. L. Baylon, N. K. Singh, B. M. Baker, E. Tajkhorshid, D. M. Kranz, Changing the peptide specificity of a human T-cell receptor by directed evolution. *Nat. Commun.* **5**, 5223 (2014).
- 10x Genomics, “A new way of exploring immunity—linking highly multiplexed antigen recognition to immune repertoire and phenotype” (Tech. Rep., 10x Genomics, 2019).
- P. Meysman, J. Barton, B. Bravi, L. Cohen-Lavi, V. Karnaukhov, E. Lilleskov, A. Montemurro, M. Nielsen, T. Mora, P. Pereira, A. Postovskaya, M. R. Martinez, J. Fernandez-de-Cossio-Diaz, A. Vujkovic, A. M. Walczak, A. Weber, R. Yin, A. Eugster, V. Sharma, Benchmarking solutions to the t-cell receptor epitope prediction problem: Immrep22 workshop report. *Immunoinformatics* **9**, 100024 (2023).
- E. J. Grant, T. M. Josephs, S. A. Valkenburg, L. Wooldridge, M. Hellard, J. Rossjohn, M. Bharadwaj, K. Kedzierska, S. Gras, Lack of heterologous cross-reactivity toward HLA-A*02:01 restricted viral epitopes is underpinned by distinct α/β T cell receptor signatures. *J. Biol. Chem.* **291**, 24335–24351 (2016).
- Z. S. Ghoreyshi, J. T. George, Quantitative approaches for decoding the specificity of the human T cell repertoire. *Front. Immunol.* **14**, (2023).
- B. Meynard-Piganeau, C. Feinauer, M. Weigt, A. M. Walczak, T. Mora, TULIP-A transformer based unsupervised language model for interacting peptides and T-cell receptors that generalizes to unseen epitopes. *bioRxiv* 549669 [Preprint]. 2023. <https://doi.org/10.1101/2023.07.19.549669>.

37. B. P. Kwee, M. Messemaker, E. Marcus, G. Oliveira, W. Scheper, C. Wu, J. Teuwen, T. Schumacher, STAPLER: Efficient learning of TCR-peptide specificity prediction from full-length TCR-peptide data. *bioRxiv* 538237 [Preprint]. 2023. <https://doi.org/10.1101/2023.04.25.538237>.
38. N. L. La Gruta, S. Gras, S. R. Daley, P. G. Thomas, J. Rossjohn, Understanding the drivers of MHC restriction of T cell receptors. *Nat. Rev. Immunol.* **18**, 467–478 (2018).
39. A. Davtyan, N. P. Schafer, W. Zheng, C. Clementi, P. G. Wolynes, G. A. Papoian, AWSEM-MD: Protein structure prediction using coarse-grained physical potentials and bioinformatically based local structure biasing. *J. Phys. Chem. B* **116**, 8494–8503 (2012).
40. W. Zheng, N. P. Schafer, A. Davtyan, G. A. Papoian, P. G. Wolynes, Predictive energy landscapes for protein–protein association. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 19244–19249 (2012).
41. B. Webb, A. Sali, Comparative protein structure modeling using MODELLER. *Curr. Protoc. Bioinformatics* **54**, 5–6 (2016).
42. 10x Genomics. *Tech. rep* 2019.

Acknowledgments

Funding: Work by the Center for Theoretical Biological Physics was supported by the NSF

(grant PHY-2019745). J.T.G. was supported by CPRIT grant RR210080. J.N.O. was also supported by the NSF (grant PHY-2210291) and by the Welch Foundation (grant C-1792). J.N.O. and J.T.G. are CPRIT Scholars in Cancer Research. **Author contributions:** X.L., H.L., J.N.O., and J.T.G. conceived of the research. A.W., X.L., H.L., J.N.O., and J.T.G. designed the research. A.W., X.L., and K.N.C. performed the research. A.W., X.L., K.N.C., H.L., J.N.O., and J.T.G. analyzed the results. A.W., X.L., K.N.C., H.L., J.N.O., and J.T.G. wrote the paper. J.T.G. supervised the research. All authors approve of the final manuscript. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials. All scripts and input files needed to reproduce the results are available in a publicly accessible repository (<https://zenodo.org/records/8374294>).

Submitted 25 September 2023

Accepted 10 April 2024

Published 15 May 2024

10.1126/sciadv.adl0161