Assessing the Generalization Capabilities of TCR Binding Predictors via Peptide Distance Analysis

Leonardo V. Castorina [1,*], **Filippo Grazioli** [2], **Pierre Machart** [2], **Anja Mösch** [2] **and Federico Errica** [2]

[1] *School of Informatics, University of Edinburgh, Edinburgh EH8 9AB, UK*
[2] *NEC Laboratories Europe, 69115 Heidelberg, Germany*

Correspondence*:
Corresponding Author
leonardo.castorina@ed.ac.uk

ABSTRACT

Understanding the interaction between T Cell Receptors (TCRs) and peptide-bound Major Histocompatibility Complexes (pMHCs) is crucial for comprehending immune responses and developing targeted immunotherapies. While recent machine learning (ML) models show remarkable success in predicting TCR-pMHC binding within training data, these models often fail to generalize to peptides outside of their training distributions, raising concerns about their applicability in therapeutic settings. Understanding and improving the generalization of these models is therefore critical to ensure real-world applications.

To address this issue, we evaluate the effect of the distance between training and testing peptide distributions on ML model empirical risk assessments, using sequence-based and 3D structure-based distance metrics. In our analysis we use several state-of-the-art models for TCR-peptide binding prediction: Attentive Variational Information Bottleneck (AVIB), NetTCR-2.0 and -2.2, and ERGO II (pre-trained autoencoder) and ERGO II (LSTM).

In this work, we introduce a novel approach for assessing the generalization capabilities of TCR binding predictors: the Distance Split (DS) algorithm. The DS algorithm controls the distance between training and testing peptides based on both sequence and structure, allowing for a more nuanced evaluation of model performance. We show that lower 3D shape similarity between training and test peptides is associated with a harder out-of-distribution task definition, which is more interesting when measuring the ability to generalize to unseen peptides. However, we observe the opposite effect when splitting using sequence-based similarity. These findings highlight the importance of using a distance-based splitting approach to benchmark models. This could then be used to estimate a confidence score on predictions on novel and unseen peptides, based on how different they are from the training ones. Additionally, our results may hint that employing 3D shape to complement sequence information could improve the accuracy of TCR-pMHC binding predictors.

Keywords: TCR, peptide, MHC, binding prediction, interaction prediction, machine learning, TCR-pMHC, 3D structure

1 INTRODUCTION

The immune system has evolved to recognize different pathogens, such as viruses, which enter cells and exploit their resources for replication. The immune response depends on how well the system can distinguish between healthy and infected/aberrant cells. It does so by leveraging specific molecules on the cell surface called Major Histocompatibility Complexes (MHC). Class I or class II MHCs are presented depending on the cell type (Alberts et al., 2017; Rowen et al., 1996; Glanville et al., 2017). An antigenic peptide is displayed by the MHC forming a peptide-bound MHC (pMHC). These peptides are derived from the proteasome, a complex of protease enzymes in cells which break down proteins into short peptides (Hewitt, 2003).

In infected cells, the presented peptides may be derived from viral proteins, whereas in healthy cells, peptides are derived from housekeeping proteins (Eisenlohr et al., 2007). T cells can recognize viral peptides in the pMHC by binding it with via the T Cell Receptors (TCRs), leading to an immune response (Hewitt, 2003). TCRs are able to recognize many different peptides via diverse sequences at the variable regions of their α and β chains (Smith-Garvin et al., 2009). TCR binding to the pMHC primarily occurs at the Complementarity-Determining Regions (CDRs). The peptide recognition is primarily mediated by the CDR3. The CDR3-α is derived from the alleles of the V and J genes; the D gene, in addition to the V and J, is involved in shaping the sequence and structure of the CDR3-β, (Feng et al., 2007; Rossjohn et al., 2015). These alleles can be recombined extensively, ensuring a high TCR repertoire diversity and allowing for a broad T cell-based immune response (Qi et al., 2014). The exposure of a naive T cell to an antigen leads to its activation and to to the development of a memory T cell population with the same TCR. This allows for a long-lasting immune response (Jameson and Masopust, 2018; Omilusik and Goldrath, 2019). A visual overview of the TCR-pMHC complex is presented in the Supplementary Materials.

Immunotherapies targeting cancer or viral infections use binding specificity to activate T cells, enhancing the immune system response. Binding specificity ensures that T cell binding occurs uniquely on the target, to kill say cancerous or infected cells, and avoid healthy cells. Research in immunotherapy covers two main categories: adoptive T cell therapy and T cell inducing vaccines. In adoptive T cell therapy, cancer patients receive specific T cells targeting and destroying the tumour (He et al., 2019; Waldman et al., 2020). T cell inducing vaccines leverage antigenic peptides or other classes of antigens to trigger specific T cell development against pathogens, such as viruses (Gilbert, 2011; Heslop and Leen, 2013).

Immunotherapy design requires an in-depth understanding of the biochemical interactions between TCRs and pMHCs to accurately stimulate the immune system. The development of computational models that predict whether TCRs and pMHCs interact would allow for a faster *in silico* screening of sequences and consequent design of TCR sequences that bind specific target pMHCs. Recent advancements in machine learning (ML) lead to the development of TCR binding predictors (Springer et al., 2021; Montemurro et al., 2021; Weber et al., 2021; Grazioli et al., 2022a; Springer et al., 2020; Koyama et al., 2023). The data input of these models consists of tuples of short amino acid sequences such as (peptide, CDR3-β) or (peptide, CDR3-β, CDR3-α, MHC). The task is usually formalized as a binary classification between binding or non-binding pairs.

Several computational studies on TCR binding prediction employ data from the Immune Epitope Database (IEDB) (Vita et al., 2019), VDJdb (Bagaev et al., 2020) and McPAS-TCR (Tickotsky et al., 2017). These databases mainly contain CDR3-β sequences, but often lack information on CDR3-α. Additionally, public TCR-pMHC interaction datasets are often limited in diversity and size. They also present sequence bias towards commonly studied viruses or binding/non-binding pairs bias from experiment setups. Furthermore,

laboratory methods that validate the TCR-pMHC interactions, such as surface plasmon resonance, titration calorimetry and fluorescence anisotropy, are usually resource-intensive and time-consuming, hindering the creation of larger datasets (Piepenbrink et al., 2009).

Generally, machine learning (ML) TCR binding predictors achieve high test performance when evaluated on test sets originating from the same source as the training set. However, various studies (Grazioli et al., 2022b; Deng et al., 2023; Mahajan et al., 2019) have shown that these methods exhibit weak cross-dataset generalization, meaning, models performance are significantly lower when training and test samples come from different distributions. Furthermore, Grazioli et al. (2022b) investigated the effect of different splitting techniques on the TChard dataset, which combines samples from IEDB, VDJdb and McPAS-TCR. They also investigated the effect of considering negative samples from wet lab assays and from random shuffling of the positive tuples. They observed that, when using a "vanilla" Random Split (RS) with negative samples from wet lab assays, ML-based models achieve an estimated area under the receiver operator characteristic (AUROC) larger than 95% on the test set. In this setting, the sets of CDR3 sequences in positive and negative samples are disjoint. This allows the models to memorize whether a CDR3 sequence was observed in either the negative or positive samples at training time. Grazioli et al. (2022b) then showed that a simple countermeasure to this problem consists in creating negative samples by randomly shuffling the positive tuples. In this setting, when the RS is employed, models achieve an AUROC score between 70% and 80%. Nevertheless, using the RS, peptides and CDR3 sequences may appear in both the training and test sets, leading to inflated estimates of the real-world model generalization capabilities.

Hence, to approximately test models performance for unseen peptides, they propose an alternative splitting method, named Hard Split (HS). The HS ensures that test peptides are never observed at training time. Therefore, the peptide sequences in the training, validation and test sets are unique to their respective sets. When using the HS and negative samples are obtained via random shuffling, the model performs barely better than random guesses, with an AUROC smaller than 55%.

To ensure real-world applicability, the HS proposed by Grazioli et al. (2022b) aims at evaluating models' generalization abilities under the toughest setup possible. In fact, new peptides arising from new pathogens may differ significantly from those used at training time. ML models can only be safely employed for broad real-world applications if they show *sufficient* generalization to unseen sequences.

In practice, however, RS and HS represent two extremes of a wider spectrum of data splitting options. Assessing other splitting options can provide further insights on the robustness of ML models.

Our work is driven by the following research question: can we estimate the performance of ML models given an unseen test peptide and its distance from the training ones? We hypothesize that, as the difference in distributions between training and test sets increases, the prediction task becomes gradually harder. Knowing how well ML predictors generalize to an increasing shift in distributions between training and test sets would provide a tangible metric to measure when it is appropriate and to which extent we can employ these methods in the real world.

For this reason, in this work we introduce a new dataset splitting algorithm called *Distance Split (DS)*. Analogously to the HS, peptides placed in the training set are absent from the test set and vice versa. However, given a distance metric, we control the distances between sequences in the training and in the test sets. We increase such distance to test performance on peptides that are "further apart" from the ones seen by the model at training time, making the task increasingly harder.

Previous efforts have predominantly relied on sequence-based metrics, such as Levenshtein distance and BLOSUM substitution matrices, to measure similarity between peptides (Korpela et al., 2023; Moris

et al., 2020; Croce et al., 2024). However, sequence metrics may not fully capture the structural aspects of TCR-peptide interactions, which are essential for accurate binding predictions. The integration of 3D structural information, such as Root Mean Square Deviation (RMSD) between peptides, could offer a more comprehensive approach to evaluating model performance on unseen peptides. Therefore, we explore different distance metrics, i.e., sequence metrics such as Levenshtein and BLOSUM (Cock et al., 2009), as well as shape metrics such as RMSD (Schrödinger, LLC, 2015), using the predicted structures of peptides. As opposed to binning the HS based on distance, the DS allows for control over the median distance of the peptides in the test and validation.

We show that the increased peptide distance between the train and test split directly correlates with the models performances. In particular, increased RMSD (shape) distance between training and test peptides leads to decreased performance. Surprisingly, for BLOSUM (sequence) distance we see the inverse relationship. In real-world scenarios, when predictions on new unseen viral epitopes are required, we believe that leveraging RMSD between the 3D structures of the training and test peptides could serve as a valuable indicator of model reliability, with higher RMSD implying increased prediction uncertainty. Conversely, the inverse relationship observed with BLOSUM distance suggests that incorporating sequence-diverse training data may actually improve model generalization to novel peptides. Thus, a combination of structural and sequence-based metrics could provide a balanced approach, enhancing both the robustness and reliability of TCR-pMHC binding predictions.

2 METHOD

2.1 Dataset Creation

With the goal of vaccine development against viruses, we select a viral subset of the VDJdb dataset, focusing on human host and MHC class I (Shugay et al., 2018). We omit MHC class II due to the small number of available samples. The dataset includes peptides from SARS-CoV-2, Influenza A, Human Immunodeficiency Virus (HIV), Hepatitis C Virus (HCV) and Cytomegalovirus (CMV). We discard data points which do not include both the CDR3-β and the peptide. The dataset contains 52 unique MHC A and 1 MHC B alleles, 16,504 unique CDR3-α sequences, 28,831 unique CDR3-β sequences and 757 unique peptides, for a total of 34,415 binding samples. To generate non-binding (i.e., negative) samples, we randomly shuffle the available (peptide, CDR3-β) pairs. This process is commonly employed in the literature (Jurtz et al., 2018; Grazioli et al., 2022a) and leverages the hypothesis that random pairs will most likely not bind. To create a balanced dataset, we randomly generate 36,641 samples of non-binding combinations of CDR3-β and peptide sequences, increasing the total number of data points to 65,946. In this study, the data consists of pairs of (peptide, CDR3-β) and a binding/non-binding label.

2.1.1 Obtaining 3D Structures of Peptides

The VDJdb dataset contains information on the primary structure of proteins, i.e., the sequence of amino acids. In reality, these sequences exist as 3D shapes and as part of a complex (Shugay et al., 2018). However, public datasets like VDJDB lack the 3D shapes of peptides. We use ESMFold, a light-weight sequence-to-shape Language Model (Lin et al., 2022) and its online API [1] to generate 3D structures for the peptide sequences. For structures that gave errors, we use OmegaFold (Wu et al., 2022), a similar sequence-to-shape Language Model. We use the shapes to calculate the 3D distance between peptides, as described in the following section.

[1] https://esmatlas.com/resources?action=fold

2.2 The Distance Split Algorithm

Given a distance metric, the goal of the Distance Split (DS) algorithm is to create data splits of the available (peptide, CDR3-β) samples, enforcing a specified median peptide distance between training and test (and validation).

As distance metrics, we use Levenshtein, BLOSUM and RMSD. The Levenshtein distance is a string edit distance measuring the minimum number of edits required to change one peptide sequence into another (Wagner and Fischer, 1974). The BLOSUM distance is a substitution matrix-based metric that quantifies the evolutionary similarity between two peptide sequences by comparing their amino acid global alignments (Henikoff and Henikoff, 1992). The RMSD distance is a 3D structural metric that measures the average deviation between the atomic positions of two superimposed peptide structures, providing an estimate of their conformational similarity. We calculated the RMSD shape distance using PyMol (Schrödinger, LLC, 2015). PyMol aligns any pairs of peptide 3D structures computes the RMSD between Cα of the two structures. The Cα atoms are the backbone carbon atoms in each amino acid, which are commonly used in RMSD calculations to provide a simplified yet accurate representation of the overall peptide structure. We focus on Cα atoms to avoid over-penalizing minor deviations in side chains, as protein folding models can generate physically impossible conformations that would otherwise lead to inflated RMSD values (Moore et al., 2022).

For each metric (Levenshtein, BLOSUM, RMSD), we calculate the pairwise distances between all 757 peptides, resulting in three distance matrices: $\mathbf{M}_{Levenshtein}$, $\mathbf{M}_{BLOSUM}$ $\mathbf{M}_{RMSD} \in \mathbb{R}^{757\times757}$. For each distance matrix, we then calculate the row-wise median $\mathbf{m}_{med} \in \mathbb{R}^{757}$. Each value of $\mathbf{m}_{med}$ corresponds to the median distance between a given peptide and all other peptides. We then select three bin ranges over the cumulative distribution of the median distances of $\mathbf{m}_{med}$, defined by a lower and an upper bound (b_l, b_u) distance, i.e., $(0, 33)$, $(33, 66)$ and $(66, 100)$. Given the median vector $\mathbf{m}_{med}$, for each interval, we compute the distances corresponding to the lower and upper percentiles, i.e., d_l and d_u, for b_l and b_u, respectively. We then filter out all the peptides whose median distance falls outside the interval of interest. The sets of peptides that fall inside this percentile interval are then sampled for test and validation. For each training-validation-test split, we use a 90-5-5 ratio. The DS algorithm iteratively samples the validation and test sets based on this ratio and the total available data. To guide this process, we calculate the expected number of data points for each split N_{train}, N_{test}, $N_{validation}$, which we refer to as the test and validation budget.

We select the test samples by iteratively sampling peptides from the specified percentile interval. In each iteration, we randomly choose a peptide from this set, and move all corresponding (peptide, CDR3-β) pairs to the test set. This process continues until the target number of test samples (test budget) is reached. Peptides that were not selected are used to form the validation set, following the same iterative procedure. Then, the remaining peptides are assigned to the training set.

Additionally, we enforce a minimum and maximum count for peptides to be included in the validation and test sets. This is meant to constrain the peptide diversity in the set. In this work, we use a minimum count of 5 maximum count of 5,000. The complete algorithm for the DS is available in 1.

Using the DS, as well as the RS and HS, we generate training, test and validation splits of the available data points for all distance metrics, $\mathbf{M}_{Levenshtein}$, $\mathbf{M}_{BLOSUM}$ and $\mathbf{M}_{RMSD}$. As shown in Supplementary Figure S2, the average number of unique peptides is roughly the same across splits.

2.3 TCR-Peptide Interaction Prediction Models

We select 5 state-of-the-art models for TCR-peptide interaction prediction: Attentive Variational Information Bottleneck (AVIB) (Grazioli et al., 2022b), NetTCR-2.0 (Jurtz et al., 2018), NetTCR-2.2 (Jensen and Nielsen, 2024), ERGO II (with pre-trained TCR autoencoder), and ERGO II-LSTM (Springer et al., 2021) (see details in Supplementary Section 2). With the exception of ERGO, all models encode the TCRβ and peptide sequences using the BLOSUM substitution matrix (Henikoff and Henikoff, 1992).

Algorithm 1: Distance Split (DS) algorithm for splitting data into training, validation, and test sets based on peptide distances within specified percentile bounds.

Input: Dataset of peptide-TCR pairs $\mathcal{D} = \{(pep_i, TCR_i)\}_{i=1}^{N}$;
Peptide distance metric (e.g., BLOSUM or RMSD);
Training split ratio s (fraction of data for training);
Lower and upper percentile bounds (b_l, b_u) over the distance distribution;
Leeway factor l: flexibility when the exact sample budget cannot be met. If the number of available samples falls below the target by up to $l \times N$, the algorithm will accept the smaller set .
Output: Training, validation, and test sets.

1) Compute unique peptides and counts:
- Extract the set of unique peptides $\mathcal{P}$ from $\mathcal{D}$.
- For each peptide $pep \in \mathcal{P}$, count the number of (pep, TCR) pairs and store in $CountMap(pep)$.

2) Calculate peptide distances:
- Compute the distance matrix $\mathbf{M}_{\text{dist}}$ between peptides using the chosen method.
- Calculate the median distance $\mathbf{m}_{\text{med}}$ for each peptide from $\mathbf{M}_{\text{dist}}$.

3) Select peptides within percentile bounds:
- Determine distance thresholds d_l and d_u corresponding to percentiles b_l and b_u of $\mathbf{m}_{\text{med}}$.
- Select peptides $\mathcal{P}^* = \{pep \in \mathcal{P} \mid \mathbf{m}_{\text{med}}(pep) \in [d_l, d_u]\}$.

4) Determine split sizes:
- Calculate total sample count: $N_{\text{total}} = \sum_{pep \in \mathcal{P}} CountMap(pep)$.
- Compute training sample count: $N_{\text{train}} = s \times N_{\text{total}}$.
- Remaining samples: $N_{\text{remaining}} = N_{\text{total}} - N_{\text{train}}$.
- Set test and validation counts: $N_{\text{test}} = N_{\text{validation}} = N_{\text{remaining}}/2$.

5) Initialize sets:
- $\mathcal{P}_{\text{train}} \leftarrow \mathcal{P}$, $\mathcal{P}_{\text{test}} \leftarrow \{\}$, $\mathcal{P}_{\text{val}} \leftarrow \{\}$.
- Counters: $t_{\text{test}} \leftarrow 0$, $t_{\text{val}} \leftarrow 0$.

6) Sample peptides for test and validation sets:

```
for each set in {test, validation} do
    while counter t < N × (1 − l) do
        Randomly sample pep from P* not already selected.
        if t + CountMap(pep) ≤ (1 + l) × N then
            Add pep to the current set.
            Remove pep from P_train and P*.
            Update counter: t ← t + CountMap(pep).
        end
    end
end
```

7) Construct final datasets:
- Training set: all (pep, TCR) pairs where $pep \in \mathcal{P}_{\text{train}}$.
- Validation set: all (pep, TCR) pairs where $pep \in \mathcal{P}_{\text{val}}$.
- Test set: all (pep, TCR) pairs where $pep \in \mathcal{P}_{\text{test}}$.

In AVIB, the TCRβ and Peptide encoders estimate Gaussian posteriors over a latent space. The various posteriors are then combined using an attention mechanism to estimate a single multi-sequence posterior distribution, which is then used to estimate the binding probability (Grazioli et al., 2022b). In NetTCR-2.0, the TCRβ and peptide encoders are encoded using several convolutional and global max pooling layers. The encoding is then concatenated and passed through two fully connected layers (Jurtz et al., 2018). NetTCR-2.2 introduces a dropout layer before the fully connected layers, ReLU over Sigmoid, and 64 units rather than 32 for the fully connected layer (Jensen and Nielsen, 2024). In the ERGO II model, the TCRβ is encoded using a pre-trained autoencoder. ERGO II-LSTM encodes the peptide using an LSTM encoder. The encodings are then passed to two fully connected layers (Springer et al., 2021).

3 RESULTS

We compute the pairwise distance between all peptides using Levenshtein, BLOSUM and RMSD distance for the viral VDJdb dataset. Descriptive statistics of the metrics are summarized in Table 1 and the median visualised in Figure 1. The peptides count distribution is heavily skewewd towards a count of less than 10 instances, with several outliers, as shown by the high standard deviation. The median RMSD has a bimodal distribution with peaks at ≈ 0.0 and ≈ 0.5 Å. The median BLOSUM is normally distributed with most values ranging from 50 to 100. Finally, the median Levenshtein distance is ≈ 8, and shows very little distribution of values.

Then, we calculate the correlation between these distances in Figure 2. Levenshtein and BLOSUM distances are very significantly correlated (Spearman ρ: 0.32, p-value: 0), as expected since BLOSUM distance uses a sequence alignment. However, additional physico-chemical information is implicitly accounted for in BLOSUM, which considers evolutionary substitution patterns of amino acids (Henikoff and Henikoff, 1992). This inclusion captures the likelihood of biologically relevant substitutions based on factors like size, charge, and hydrophobicity, which are not reflected in simple sequence-based metrics like Levenshtein distance. We observe no correlation between RMSD and sequence-based metrics.

Metric	**Average**	**Median**	**Stdev**
Peptide (counts)	45	2	576
Levenshtein	8.29	8.25	0.60
BLOSUM	85.04	71.10	28.97
RMSD (Å)	0.46	0.36	0.27

Table 1. Descriptive statistics of peptides counts, as well as Levenshtein, BLOSUM, and RMSD distances for the viral VDJdb dataset. Levenshtein and BLOSUM are sequence distance metrics. RMSD refers to the average 3D distance between peptides.

For our experiments we use five state-of-the-art deep learning models for TCR-peptide interaction prediction: NetTCR-2.0 (Jurtz et al., 2018), NetTCR-2.2 (Jensen and Nielsen, 2024), AVIB (Grazioli et al., 2022a), ERGO II and ERGO II-LSTM (Springer et al., 2021). We evaluate how increasing the distance between the train and test set affects their performance. We train and test all models with RS, HS as baselines and and DS. For the DS, we use lower and upper bound pairs over the cumulative distribution of the median distances: $(0, 33)$, $(33, 66)$ and $(66, 100)$. We repeat each split with 5 different random seeds and the average results and standard deviation are shown in Figure 3. More detailed results with Levenshtein DS, overall and per-peptide metrics are available in Supplementary Section 2.

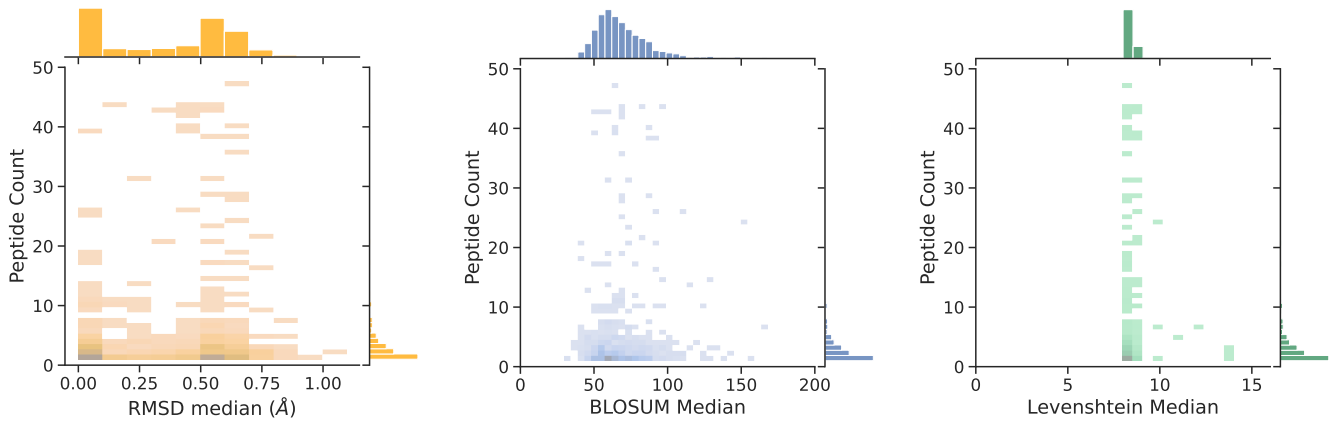

Figure 1. Distribution of peptide count against median distance of the peptide against all other peptides. Distances are RMSD (3D shape), BLOSUM (sequence), and Levenshtein (sequence), respectively.

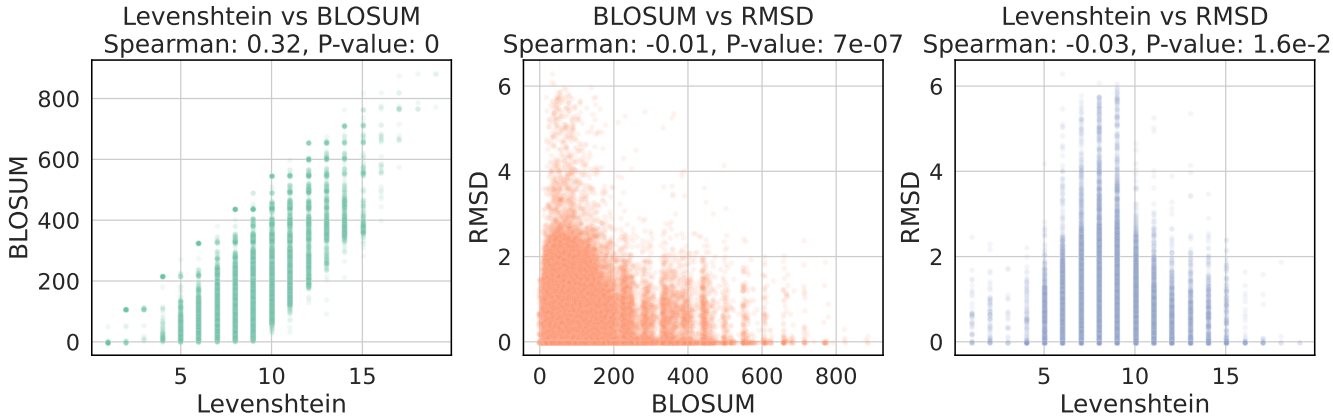

Figure 2. Correlations between peptides distance metrics in the Viral VDJDB dataset. Levenshtein and BLOSUM are sequence-based distances, while RMSD is a shape-based distance. There is a positive correlation between sequence distances and no correlation between shape and sequence distances.

As shown in Figure 3, AVIB and NetTCR-2.2 perform similarly in the sequence and shape DS splits, followed by NetTCR-2.0 and both ERGO II models. As expected, all models perform best in the RS, as peptides in the train set also appear in the test set. On the other hand, performance on the HS is significantly lower compared to the RS.

Regarding the DS splits, as the median distance between the training and test sets increases, model performance worsens in RMSD-based DS splits, as shown by significant Spearman correlations ($p < 0.05$) for NetTCR-2.2, NetTCR-2.0, and AVIB, all within a 95% confidence interval. In contrast, we observe the opposite trend for BLOSUM-based DS splits, where models perform better as the median BLOSUM distance between the training and test sets increases, with all p-values below 0.05. In general, the highest performance in the DS splits is slightly higher in the RMSD $(0, 33)$ bin, than in the BLOSUM $(66, 100)$ bin, except for ERGO II (LSTM). There is generally no significant effect on the performance trend when using Levenshtein-based distance split, except for ERGO II (LSTM) which shows a similar trend to BLOSUM (see Supplementary Figure S13).

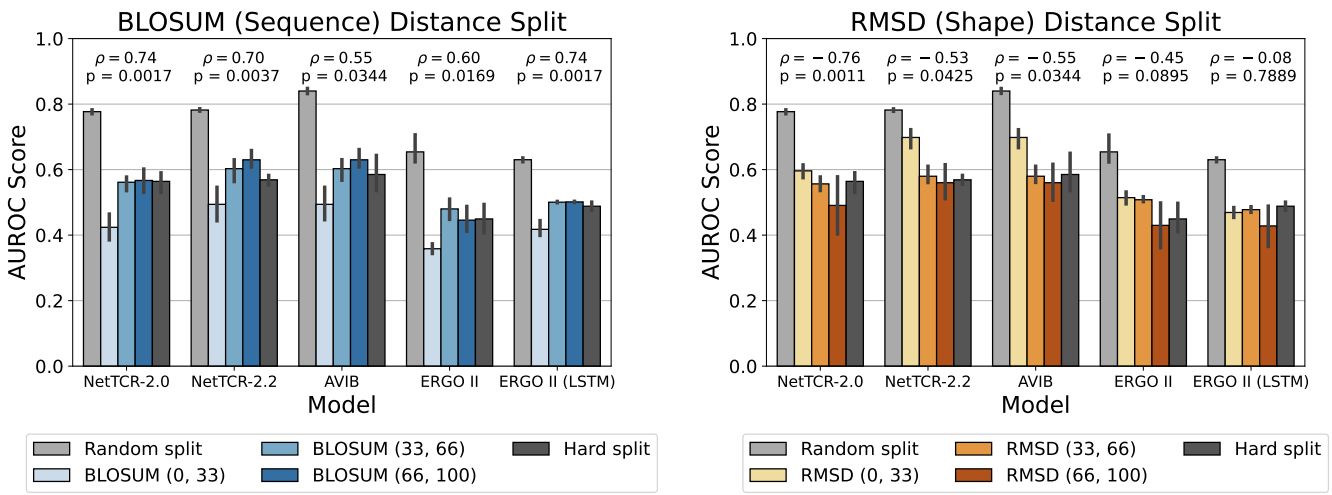

Figure 3. AUROC Scores for TCR's CDR3β-peptide binding prediction for models trained and tested using various splitting techniques. The commonly used *Random Split* allows test peptides to be observed also during training. In the *Hard Split*, peptides are exclusively allocated to either the training or test set. In the *Distance Split*, we enforce specific median distances between the training and test peptides, using either sequence distance (BLOSUM) or shape distance (RMSD). The training-test distance between peptides is controlled by selecting three percentile intervals over the cumulative median peptide-peptide distance distribution: $(0, 33)$, $(33, 66)$ and $(66, 100)$. Mean and standard deviation are calculated over 5 repeated experiments with different random seeds. More detailed results with Levenshtein DS, overall and per-peptide metrics are available in Supplementary Section 2.

4 DISCUSSION

The TCR-peptide/pMHC interaction prediction problem presents evaluation challenges (Grazioli et al., 2022b). When data is randomly split, models can achieve over-optimistic test performance, as test sequences can be observed at training time.

As a solution, Grazioli et al. (2022b) proposed the Hard Split (HS), which consists in randomly sampling test peptides and moving all instances of that peptide to the test set. The HS ensures that test peptides are never seen at training time, thus simulating what happens in the real world, for example, when new peptides arise from new pathogens.

This raises the question of whether it is possible to infer how well a model will generalize to an unseen peptide, by knowing how different the novel peptide is from those seen at training time.

We divide the peptides into different training and test sets using a distance-based algorithm using a sequence metric (BLOSUM) and a shape metric (RMSD). To obtain the RMSD, we generate the 3D structures of the antigenic peptides using ESMFold and OmegaFold (Wu et al., 2022; Lin et al., 2022). We then use PyMol to align and calculate the Root Mean Squared Distance (RMSD) between the Cα atoms in the protein backbone (Schrödinger, LLC, 2015). We use Cα RMSD, as opposed to all-atoms RMSD, as it is computationally less intensive and produces similar results (Mechelke and Habeck, 2010). Additionally, predicted 3D structures can (and often will) have artifacts such as physically-unlikely positioning of the amino acid side chains, which could unfairly increase the all-atom RMSD (Moore et al., 2022).

To simulate real-world scenarios involving unseen peptides in a controlled manner, we propose the *Distance Split* (DS) approach (see Algorithm 1). Using the DS, we chose each test split to guarantee a given median distance between peptides in the training and test peptides. By using the median, some

peptides with properties similar to those in the training set may still be included in the test set, thereby allowing for a degree of overlap while maintaining sufficient variability. Nevertheless, when employing the DS, analogously to the HS, test and validation peptides cannot appear at training time. For more stringent conditions, the minimum distance could be used with the DS to ensure that all test peptides are distinctly different from those in the training set.

Specifically, the test peptides are chosen by sampling from a subset of peptides, whose median distance from the other peptides is between a lower and an upper bound. In our experiments, we consider the following intervals over the cumulative median distance distribution: $(0, 33)$, $(33, 66)$ and $(66, 100)$.

As shown in Figure 2, we observe no correlation between RMSD and BLOSUM. Several factors could explain this. The 3D structure, as opposed to the amino acid sequence, is more relevant to describe protein-protein interactions. For example, certain angles in structures could make specific areas interact more than others. Related to this, some small differences in sequence metrics may result in larger structural differences, for instance, if two oppositely charged amino acids are placed next to each other. The 3D structure and RMSD metric can therefore provide a stronger signal for binding prediction models and may be used for testing models performance in out-of-distribution settings, or for the development of new models. Interestingly, we observe that models' performance is slightly lower for DS in the ranges $(33, 66)$ and $(66, 100)$ compared to that of the HS. The HS ensures that both test and train peptides are *unique*. However, even with this uniqueness, test and train peptides might still possess a certain degree of similarity. This similarity provides the model with additional chemical information about the peptide. In contrast, DS enforces a threshold distance between the train and test sets, thereby controlling the difficulty.

Obtaining a 3D structure of a protein often requires lengthy experiments to crystallize it, which may take days, months, or years (Ballone et al., 2020). With the advent of sequence-to-shape models like AlphaFold (Jumper et al., 2021), OmegaFold (Wu et al., 2022) and ESMFold (Lin et al., 2022), this lengthy process can often be reduced to minutes or hours and makes it possible to obtain a fairly accurate 3D shapes of proteins. Structural validation in the lab of these models is currently underway, with some researchers reporting partial success, even when solving for previously unknown structures (Van Breugel et al., 2022). Successful applications of sequence-to-shape models in protein research include vaccine design (Goswami et al., 2022), binding affinity ranking (Chang and Perez, 2023), protein sequence design (Castorina et al., 2024; Watson et al., 2022) and benchmarking (Castorina et al., 2023).

The structures produced by sequence-to-shape models like the ones mentioned above, may be biased towards the training sequences in the Protein Data Bank (PDB), which is composed of larger structures compared to peptides, which are made up of just a few amino acids in length (Berman et al., 2003). Nevertheless, in this study, we use the shape as a representation of the real shape and assume that differences between two shapes would be consistent. This analysis could be extended to the full VDJdb dataset to test how the findings transfer to non-viral peptides.

In the structure-based DS, as the median RMSD distance between training and test peptide is increased, the model performance generally worsens. On the other hand, we observe the opposite effect in sequence-based DS, where increasing the BLOSUM sequence between training-test leads to better generalization performance (See Figure 3).

The immune system has evolved to recognize multitudes of pathogens, some of which may have never existed. The TCR-pMHC interaction is therefore degenerate, meaning that many TCRs can recognize the same peptide and that a TCR can recognize millions of peptides, even if the sequences are not identical (Lee et al., 2020). Additionally, TCRs may bind different peptides using different binding modes

(Coles et al., 2020). However, as shown in Figure 2, we find no correlation between sequence and shape distance metrics *for the peptide*. Coupled with the results of Figure 3, this may imply that TCRs bind conformationally similar viral peptides even if the sequences are different. This is supported by the fact that the peptide binding pocket of MHCs have anchor positions, like the B pocket, which can accommodate different secondary anchor amino acids without altering the overall peptide conformation (Kersh et al., 2001; Nguyen et al., 2021). This flexibility allows TCRs to recognize structurally similar peptides despite sequence variations, suggesting that structural rather than purely sequence-based features drive TCR recognition efficiency in some cases (Kersh et al., 2001; Nguyen et al., 2021). Additionally, viruses, like SARS-CoV-2, evolve by changing their sequences to escape immune detection while keeping the structural features needed to interact with host cells. For example, mutations in the spike protein allow SARS-CoV-2 to evade the immune system but still bind effectively to the ACE2 receptor, ensuring the virus remains functional (Barton et al., 2021; Xue et al., 2024). It is therefore possible that in the BLOSUM $(0, 33)$ split, models learn how similar peptides are bound by TCRs and is able to generalise better.

Additionally, in the field of protein design, benchmarking has revealed that even with a sequence similarity below 50%, the predicted 3D structures look very similar (Castorina et al., 2023). These results suggest that BLOSUM-based sequence augmentation, aimed to generate similar binding peptides for a given TCR, may reduce generalization. In contrast, a shape-based augmentation, for example by using protein sequence design models to generate new peptides, may enhance generalization to unseen peptides (Qi and Zhang, 2020; Castorina et al., 2024). Additionally, given that the peptide sequence and shape are not correlated, including the 3D shape may increase the overall performance, as shown by recent papers (Ji et al., 2024; Li et al., 2024; Bradley, 2023; Pham et al., 2024) However, as TCR and MHC structures are generally similar, it may be worthwhile to limit the 3D information to the regions that vary in shape, like the peptide, the CDRs, and the binding pocket to reduce the model complexity and maintain performance (Bradley, 2023). Future work can explore the differences in structural distance between models such as AlphaFold, ESMFold and OmegaFold in the context of peptides, to further evaluate biases and divergences of structures generated by different sequence-to-shape models.

A short-coming of the BLOSUM distance for the peptide is that changes in the anchor amino acids should be weighted more than changes outside the anchors. Future work could explore a more peptide-centric metric of distance based on anchor positions. Similarly, a short-coming of the RMSD is that due to differences in peptide lengths, it is not possible to normalise the distance by the length of the peptide. In our pairwise comparison, however, most (about 87%) [2] of the peptide distances were same-length comparison. Future work could explore the RMSD of the core amino acids of the peptide and using length-normalised RMSD like RMSD_{100} (Carugo, 2007).

Sequence-to-shape models could also be used to predict full TCR complexes for binding prediction. In their current state, however, tools for complex modelling are computationally expensive and, although they may be comparable in speed with molecular dynamics simulations, it is currently unfeasible to test all possible combinations of TCR, peptides and pMHC for the binding prediction task.

Our results suggest that the use of 3D shapes in the context of TCR-peptide/pMHC interaction prediction could help reduce the uncertainty about the generalization capabilities of ML models to unseen sequences. Given enough computational resources, 3D shapes could be predicted for the whole TCR structure, as well as for the MHCs presenting the peptide. This would enable the design of models that, given the individual structures of each input sequence, will predict more accurate binding interactions.

[2] 248,630 matching length pairs comparisons out of 285,090 using 757 peptides

CONFLICT OF INTEREST STATEMENT

Authors AM, PM and FE are employed by NEC Laboratories Europe GmbH. LVC and FG were employed by NEC Laboratories Europe GmbH during the project.

AUTHOR CONTRIBUTIONS

All authors were involved in discussions, planning, brainstorming, and writing of the paper. FG and LVC both wrote and analysed the distance split algorithm and trained the AVIB, NetTCR-2.0, NetTCR-2.2, ERGO II (AutoEncoder and LSTM) on the dataset. LVC produced the 3D structures of epitopes.

FUNDING

Funding for this project was provided by NEC Labs Europe GmbH.

ACKNOWLEDGMENTS

We would like to thank Timo Sztyler (NEC Labs Europe) and Martin Renqiang Min (NEC Labs America) for helpful discussions, especially at the inception of the project.

DATA AVAILABILITY STATEMENT

The VDJdb dataset used for this study can be found in https://vdjdb.cdr3.net. The AVIB model is available at: https://github.com/nec-research/vibtcr. The NetTCR-2.0 is available at: https://github.com/mnielLab/NetTCR-2.0. The Distance Split code is available at: https://github.com/nec-research/peptide-distance-analysis

REFERENCES

Alberts, B., Johnson, A., Lewis, J., Morgan, D., Raff, M., Roberts, K., et al. (2017). *Molecular biology of the cell* (WW Norton & Company)

Bagaev, D. V., Vroomans, R. M., Samir, J., Stervbo, U., Rius, C., Dolton, G., et al. (2020). Vdjdb in 2019: database extension, new analysis infrastructure and a t-cell receptor motif compendium. *Nucleic Acids Research* 48, D1057–D1062

Ballone, A., Lau, R. A., Zweipfenning, F. P. A., and Ottmann, C. (2020). A new soaking procedure for X-ray crystallographic structural determination of protein–peptide complexes. *Acta Crystallogr F Struct Biol Commun* 76, 501–507. doi:10.1107/S2053230X2001122X

Barton, M. I., MacGowan, S. A., Kutuzov, M. A., Dushek, O., Barton, G. J., and van der Merwe, P. A. (2021). Effects of common mutations in the sars-cov-2 spike rbd and its ligand, the human ace2 receptor on binding affinity and kinetics. *eLife* 10. doi:10.7554/elife.70658

Berman, H., Henrick, K., and Nakamura, H. (2003). Announcing the worldwide Protein Data Bank. *Nat Struct Mol Biol* 10, 980–980. doi:10.1038/nsb1203-980

Bradley, P. (2023). Structure-based prediction of t cell receptor: peptide-mhc interactions. *Elife* 12, e82813

Carugo, O. (2007). Statistical validation of the root-mean-square-distance, a measure of protein structural proximity. *Protein Engineering, Design and Selection* 20, 33–37. doi:10.1093/protein/gzl051

Castorina, L. V., Petrenas, R., Subr, K., and Wood, C. W. (2023). PDBench: Evaluating computational methods for protein-sequence design. *Bioinformatics* 39, btad027. doi:10.1093/bioinformatics/btad027

Castorina, L. V., Ünal, S. M., Subr, K., and Wood, C. W. (2024). Timed-design: flexible and accessible protein sequence design with convolutional neural networks. *Protein Engineering, Design and Selection* 37, gzae002

Chang, L. and Perez, A. (2023). Ranking Peptide Binders by Affinity with AlphaFold**. *Angewandte Chemie* 135. doi:10.1002/ange.202213362

Cock, P. J. A., Antao, T., Chang, J. T., Chapman, B. A., Cox, C. J., Dalke, A., et al. (2009). Biopython: Freely available Python tools for computational molecular biology and bioinformatics. *Bioinformatics* 25, 1422–1423. doi:10.1093/bioinformatics/btp163

Coles, C. H., Mulvaney, R. M., Malla, S., Walker, A., Smith, K. J., Lloyd, A., et al. (2020). Tcrs with distinct specificity profiles use different binding modes to engage an identical peptide–hla complex. *The Journal of Immunology* 204, 1943–1953

Croce, G., Bobisse, S., Moreno, D. L., Schmidt, J., Guillame, P., Harari, A., et al. (2024). Deep learning predictions of tcr-epitope interactions reveal epitope-specific chains in dual alpha t cells. *Nature Communications* 15. doi:10.1038/s41467-024-47461-8

Deng, L., Ly, C., Abdollahi, S., Zhao, Y., Prinz, I., and Bonn, S. (2023). Performance comparison of TCR-pMHC prediction tools reveals a strong data dependency. *Frontiers in Immunology* 14. doi:10.3389/fimmu.2023.1128326

Eisenlohr, L. C., Huang, L., and Golovina, T. N. (2007). Rethinking peptide supply to MHC class i molecules. *Nature Reviews Immunology* 7, 403–410. doi:10.1038/nri2077

Feng, D., Bond, C. J., Ely, L. K., Maynard, J., and Garcia, K. C. (2007). Structural evidence for a germline-encoded t cell receptor–major histocompatibility complex interaction'codon'. *Nature immunology* 8, 975–983

Gilbert, S. C. (2011). T-cell-inducing vaccines - what's the future. *Immunology* 135, 19–26. doi:10.1111/j.1365-2567.2011.03517.x

Glanville, J., Huang, H., Nau, A., Hatton, O., Wagar, L. E., Rubelt, F., et al. (2017). Identifying specificity groups in the t cell receptor repertoire. *Nature* 547, 94–98

Goswami, A., Kumar, S. M., and Gore, M. M. (2022). De Novo *Design of Anti-Variant COVID Vaccine with T-cell Memory*. Preprint, Bioinformatics. doi:10.1101/2022.10.20.513049

Grazioli, F., Machart, P., Mösch, A., Li, K., Castorina, L., Pfeifer, N., et al. (2022a). Attentive Variational Information Bottleneck for TCR-peptide Interaction Prediction. *Bioinformatics* , btac820doi:10.1093/bioinformatics/btac820

Grazioli, F., Mösch, A., Machart, P., Li, K., Alqassem, I., O'Donnell, T. J., et al. (2022b). On TCR binding predictors failing to generalize to unseen peptides. *Front. Immunol.* 13, 1014256. doi:10.3389/fimmu.2022.1014256

He, Q., Jiang, X., Zhou, X., and Weng, J. (2019). Targeting cancers through TCR-peptide/MHC interactions. *Journal of Hematology and Oncology* 12. doi:10.1186/s13045-019-0812-8

Henikoff, S. and Henikoff, J. G. (1992). Amino acid substitution matrices from protein blocks. *Proceedings of the National Academy of Sciences* 89, 10915–10919. doi:10.1073/pnas.89.22.10915

Heslop, H. E. and Leen, A. M. (2013). T-cell therapy for viral infections. *Hematology* 2013, 342–347. doi:10.1182/asheducation-2013.1.342

Hewitt, E. W. (2003). The MHC class i antigen presentation pathway: strategies for viral immune evasion. *Immunology* 110, 163–169. doi:10.1046/j.1365-2567.2003.01738.x

Jameson, S. C. and Masopust, D. (2018). Understanding Subset Diversity in T Cell Memory. *Immunity* 48, 214–226. doi:10.1016/j.immuni.2018.02.010

Jensen, M. F. and Nielsen, M. (2024). Nettcr 2.2 - improved tcr specificity predictions by combining pan- and peptide-specific training strategies, loss-scaling and integration of sequence similarity doi:10.7554/elife.93934.2

Ji, H., Wang, X.-X., Zhang, Q., Zhang, C., and Zhang, H.-M. (2024). Predicting TCR sequences for unseen antigen epitopes using structural and sequence features. *Briefings in Bioinformatics* 25, bbae210. doi:10.1093/bib/bbae210

Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., et al. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature* 596, 583–589. doi:10.1038/s41586-021-03819-2

Jurtz, V. I., Jessen, L. E., Bentzen, A. K., Jespersen, M. C., Mahajan, S., Vita, R., et al. (2018). Nettcr: sequence-based prediction of tcr binding to peptide-mhc complexes using convolutional neural networks. *BioRxiv* , 433706

Kersh, G. J., Miley, M. J., Nelson, C. A., Grakoui, A., Horvath, S., Donermeyer, D. L., et al. (2001). Structural and functional consequences of altering a peptide mhc anchor residue. *The Journal of Immunology* 166, 3345–3354. doi:10.4049/jimmunol.166.5.3345

Korpela, D., Jokinen, E., Dumitrescu, A., Huuhtanen, J., Mustjoki, S., and Lähdesmäki, H. (2023). Epic-trace: predicting tcr binding to unseen epitopes using attention and contextualized embeddings. *Bioinformatics* 39. doi:10.1093/bioinformatics/btad743

Koyama, K., Hashimoto, K., Nagao, C., and Mizuguchi, K. (2023). Attention network for predicting t-cell receptor–peptide binding can associate attention with interpretable protein structural properties. *Frontiers in Bioinformatics* 3. doi:10.3389/fbinf.2023.1274599

Lee, C. H., Salio, M., Napolitani, G., Ogg, G., Simmons, A., and Koohy, H. (2020). Predicting cross-reactivity and antigen specificity of t cell receptors. *Frontiers in Immunology* 11. doi:10.3389/fimmu.2020.565096

Li, F., Qian, X., Liu, Y., Lai, X., Zhang, X., Zhu, X., et al. (2024). Tcrcost: A deep learning model utilizing tcr 3d structure for enhanced tcr-peptide binding prediction. *Frontiers in Genetics* 15, 1346784

Lin, Z., Akin, H., Rao, R., Hie, B., Zhu, Z., Lu, W., et al. (2022). *Language Models of Protein Sequences at the Scale of Evolution Enable Accurate Structure Prediction*. Preprint, Synthetic Biology. doi:10.1101/2022.07.20.500902

Mahajan, S., Yan, Z., Jespersen, M. C., Jensen, K. K., Marcatili, P., Nielsen, M., et al. (2019). Benchmark datasets of immune receptor-epitope structural complexes. *BMC Bioinformatics* 20. doi:10.1186/s12859-019-3109-6

Mechelke, M. and Habeck, M. (2010). Robust probabilistic superposition and comparison of protein structures. *BMC Bioinformatics* 11. doi:10.1186/1471-2105-11-363

Montemurro, A., Schuster, V., Povlsen, H. R., Bentzen, A. K., Jurtz, V., Chronister, W. D., et al. (2021). NetTCR-2.0 enables accurate prediction of TCR-peptide binding by using paired TCRα and β sequence data. *Communications Biology* 4. doi:10.1038/s42003-021-02610-3

Moore, P. B., Hendrickson, W. A., Henderson, R., and Brunger, A. T. (2022). The protein-folding problem: Not yet solved. *Science* 375, 507–507. doi:10.1126/science.abn9422

Moris, P., De Pauw, J., Postovskaya, A., Gielis, S., De Neuter, N., Bittremieux, W., et al. (2020). Current challenges for unseen-epitope tcr interaction prediction and a new perspective derived from image classification. *Briefings in Bioinformatics* 22. doi:10.1093/bib/bbaa318

Nguyen, A. T., Szeto, C., and Gras, S. (2021). The pockets guide to hla class i molecules. *Biochemical Society Transactions* 49, 2319–2331. doi:10.1042/bst20210410

Omilusik, K. D. and Goldrath, A. W. (2019). Remembering to remember: T cell memory maintenance and plasticity. *Current Opinion in Immunology* 58, 89–97. doi:10.1016/j.coi.2019.04.009

Pham, M.-D. N., Su, C. T.-T., Nguyen, T.-N., Nguyen, H.-N., An Nguyen, D. D., Giang, H., et al. (2024). epitcr-kda: Knowledge distillation model on dihedral angles for tcr-peptide prediction. *bioRxiv* doi:10.1101/2024.05.18.594806

Piepenbrink, K. H., Gloor, B. E., Armstrong, K. M., and Baker, B. M. (2009). Methods for Quantifying T cell Receptor Binding Affinities and Thermodynamics. In *Methods in Enzymology* (Elsevier), vol. 466. 359–381. doi:10.1016/S0076-6879(09)66015-8

Qi, Q., Liu, Y., Cheng, Y., Glanville, J., Zhang, D., Lee, J.-Y., et al. (2014). Diversity and clonal selection in the human T-cell repertoire. *Proceedings of the National Academy of Sciences* 111, 13139–13144. doi:10.1073/pnas.1409155111. Number: 36

Qi, Y. and Zhang, J. Z. H. (2020). Densecpd: Improving the accuracy of neural-network-based computational protein sequence design with densenet. *Journal of Chemical Information and Modeling* 60, 1245–1252. doi:10.1021/acs.jcim.0c00043

Rossjohn, J., Gras, S., Miles, J. J., Turner, S. J., Godfrey, D. I., and McCluskey, J. (2015). T cell antigen receptor recognition of antigen-presenting molecules. *Annual review of immunology* 33, 169–200

Rowen, L., Koop, B. F., and Hood, L. (1996). The complete 685-kilobase dna sequence of the human β t cell receptor locus. *Science* 272, 1755–1762

Schrödinger, LLC (2015). The PyMOL molecular graphics system, version 1.8. PyMOL The PyMOL Molecular Graphics System, Version 1.8, Schrödinger, LLC.

Shugay, M., Bagaev, D. V., Zvyagin, I. V., Vroomans, R. M., Crawford, J. C., Dolton, G., et al. (2018). VDJdb: A curated database of T-cell receptor sequences with known antigen specificity. *Nucleic Acids Research* 46, D419–D427. doi:10.1093/nar/gkx760

Smith-Garvin, J. E., Koretzky, G. A., and Jordan, M. S. (2009). T cell activation. *Annual Review of Immunology* 27, 591–619. doi:10.1146/annurev.immunol.021908.132706. PMID: 19132916

Springer, I., Besser, H., Tickotsky-Moskovitz, N., Dvorkin, S., and Louzoun, Y. (2020). Prediction of specific tcr-peptide binding from large dictionaries of tcr-peptide pairs. *Frontiers in Immunology* 11. doi:10.3389/fimmu.2020.01803

Springer, I., Tickotsky, N., and Louzoun, Y. (2021). Contribution of t cell receptor alpha and beta cdr3, mhc typing, v and j genes to peptide binding prediction. *Frontiers in immunology* 12, 664514

Tickotsky, N., Sagiv, T., Prilusky, J., Shifrut, E., and Friedman, N. (2017). Mcpas-tcr: a manually curated catalogue of pathology-associated t cell receptor sequences. *Bioinformatics* 33, 2924–2929

Van Breugel, M., Rosa E Silva, I., and Andreeva, A. (2022). Structural validation and assessment of AlphaFold2 predictions for centrosomal and centriolar proteins and their complexes. *Commun Biol* 5, 312. doi:10.1038/s42003-022-03269-0

Vita, R., Mahajan, S., Overton, J. A., Dhanda, S. K., Martini, S., Cantrell, J. R., et al. (2019). The immune epitope database (iedb): 2018 update. *Nucleic acids research* 47, D339–D343

Wagner, R. A. and Fischer, M. J. (1974). The string-to-string correction problem. *Journal of the ACM (JACM)* 21, 168–173

Waldman, A. D., Fritz, J. M., and Lenardo, M. J. (2020). A guide to cancer immunotherapy: From T cell basic science to clinical practice. *Nat Rev Immunol* 20, 651–668. doi:10.1038/s41577-020-0306-5

Watson, J. L., Juergens, D., Bennett, N. R., Trippe, B. L., Yim, J., Eisenach, H. E., et al. (2022). Broadly applicable and accurate protein design by integrating structure prediction networks and diffusion generative models. *bioRxiv* doi:10.1101/2022.12.09.519842

Weber, A., Born, J., and Martínez, M. R. (2021). TITAN: T-cell receptor specificity prediction with bimodal attention networks. *Bioinformatics* 37, i237–i244. doi:10.1093/bioinformatics/btab294

Wu, R., Ding, F., Wang, R., Shen, R., Zhang, X., Luo, S., et al. (2022). *High-Resolution de Novo Structure Prediction from Primary Sequence*. Preprint, Bioinformatics. doi:10.1101/2022.07.21.500999

Xue, S., Han, Y., Wu, F., and Wang, Q. (2024). Mutations in the sars-cov-2 spike receptor binding domain and their delicate balance between ace2 affinity and antibody evasion. *Protein & Cell* 15, 403–418. doi:10.1093/procel/pwae007

Supplementary Tables and Figures

TCR-pMHC Complex Overview

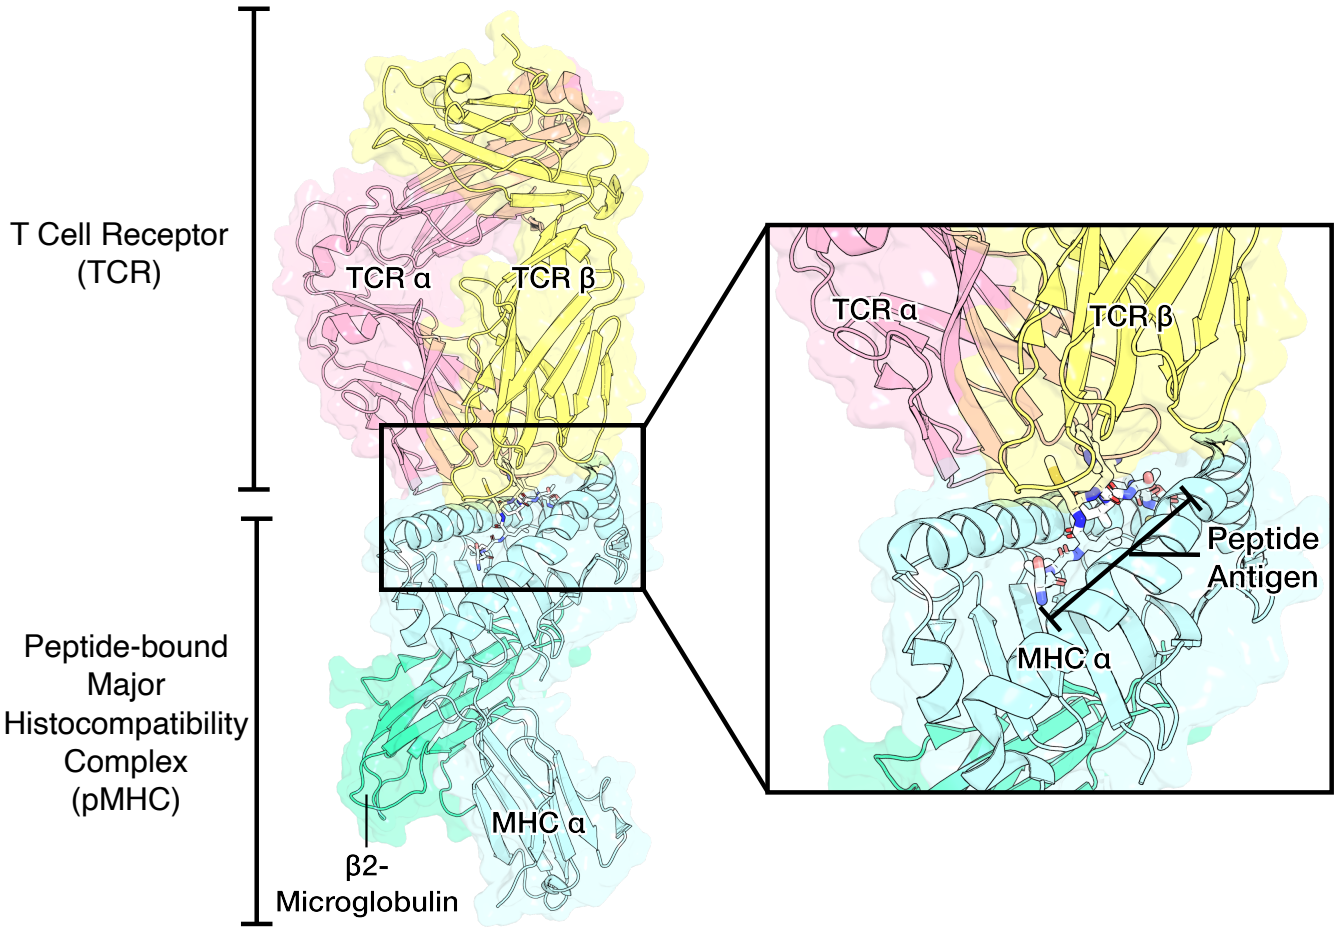

Figure S1. Overview of the TCR-pMHC Complex. The TCR is composed of an α and β chain which directly interact with the peptides. The MHC binds the peptide. Class 1 MHCs (pictured above) are composed of an α chain and a β2-Microglobulin. Class 2 MHCs, on the other hand, are composed of an α and β chain.

1 TCR-PMHC MODELS IMPLEMENTATION

We used the TCRModels library available on GitHub at this link to run Attentive Variational Information Bottleneck (AVIB) [Grazioli et al., 2022b], NetTCR-2.0 Jurtz et al. [2018], NetTCR-2.2 Jensen and Nielsen [2024], ERGO II (with pre-trained TCR autoencoder), and ERGO II-LSTM Springer et al. [2021b].

All models were trained for 5 seeds and 300 epochs with early stopping patience of 50 epochs. We used optimal parameters for each of these models provided in the corresponding publications.

1.1 Splits Peptide Numbers

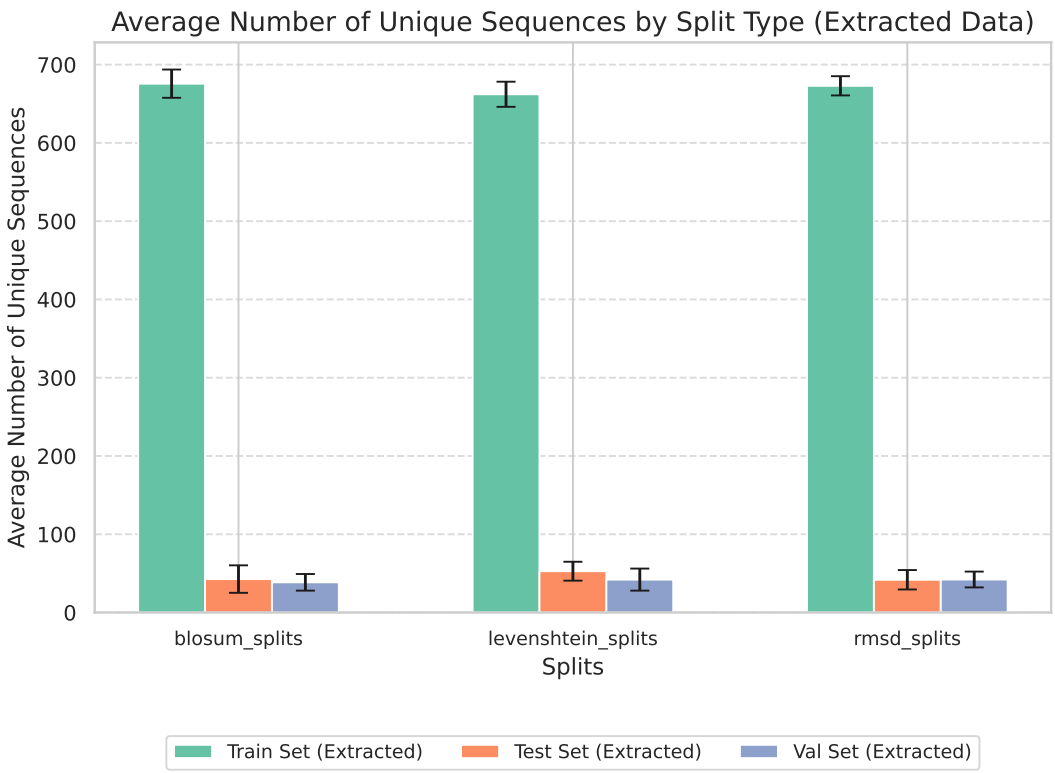

Figure S2. Number of peptides per split type in the Distance Split.

2 PERFORMANCE SCORES TCR-PMHC MODELS

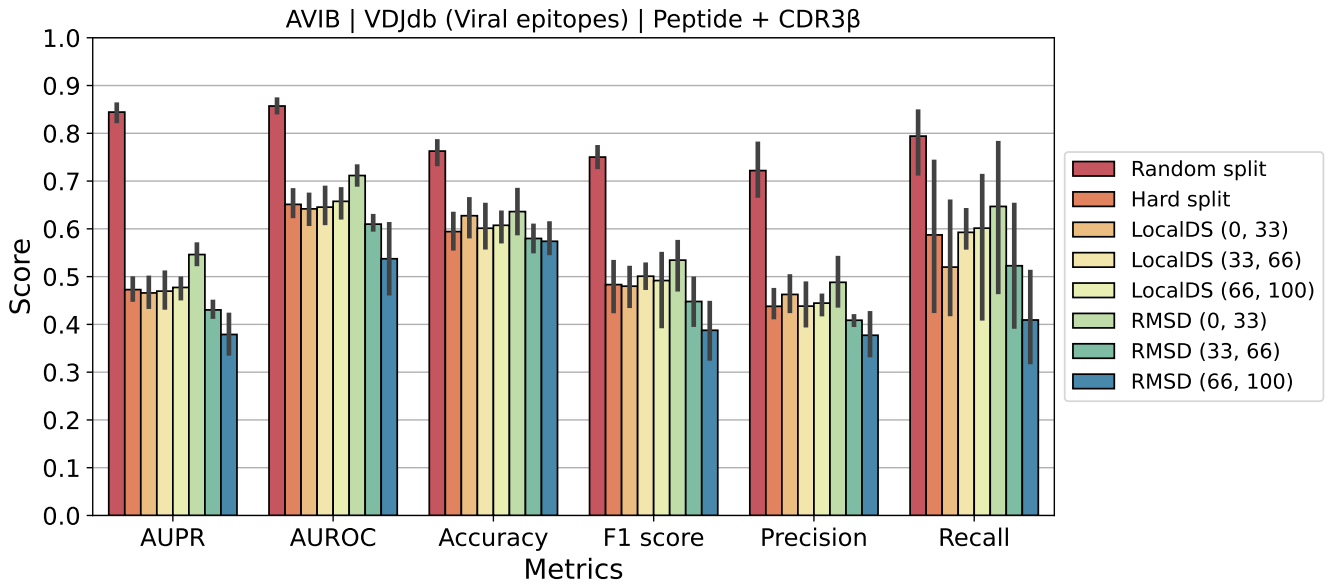

Figure S3. Performance metrics for the AVIB model [Grazioli et al., 2022a].

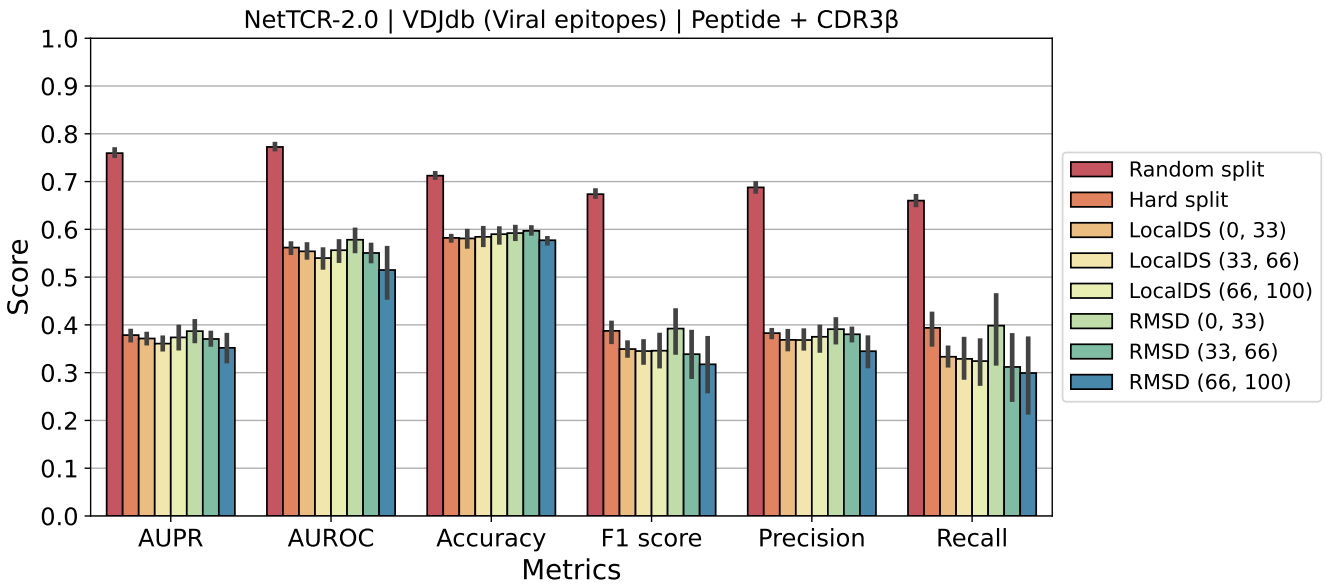

Figure S4. Performance metrics for the NetTCR-2.0 model [Jurtz et al., 2018].

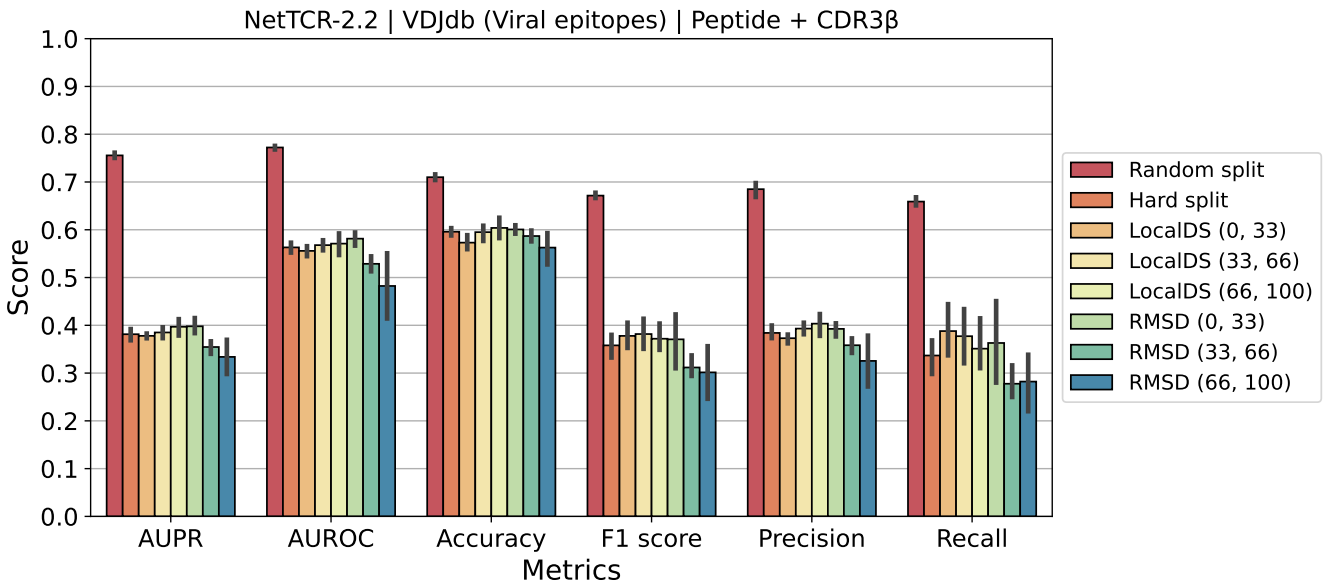

Figure S5. Performance metrics for the NetTCR-2.2 model [Jensen and Nielsen, 2024].

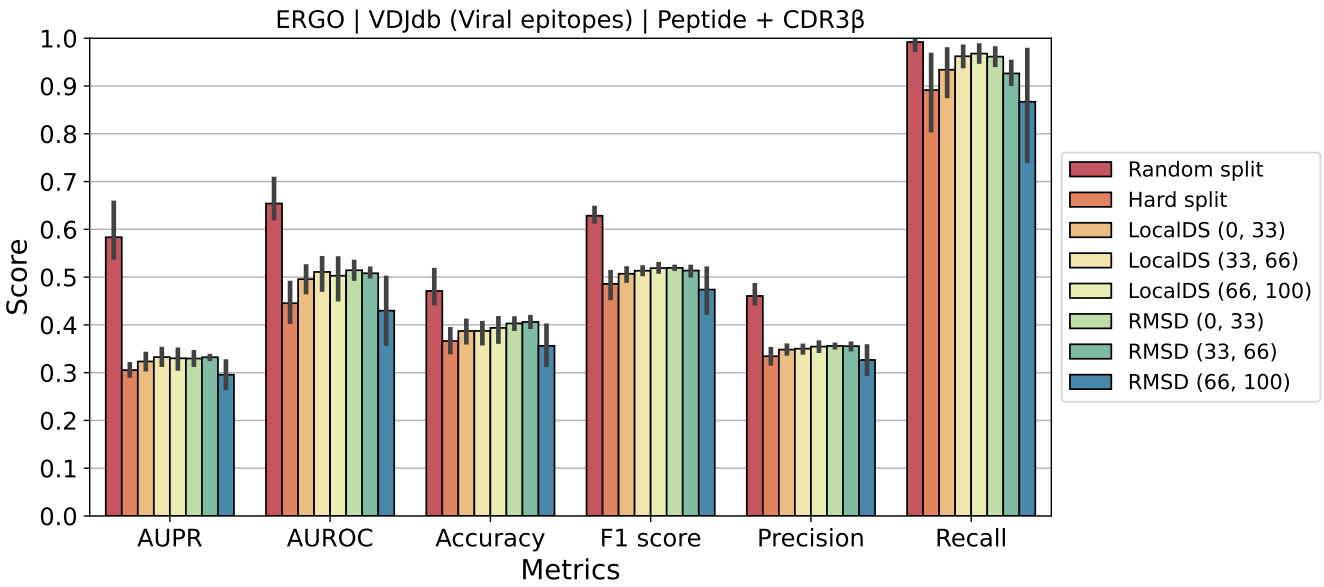

Figure S6. Performance metrics for the ERGO II model with autoencoder [Springer et al., 2021a].

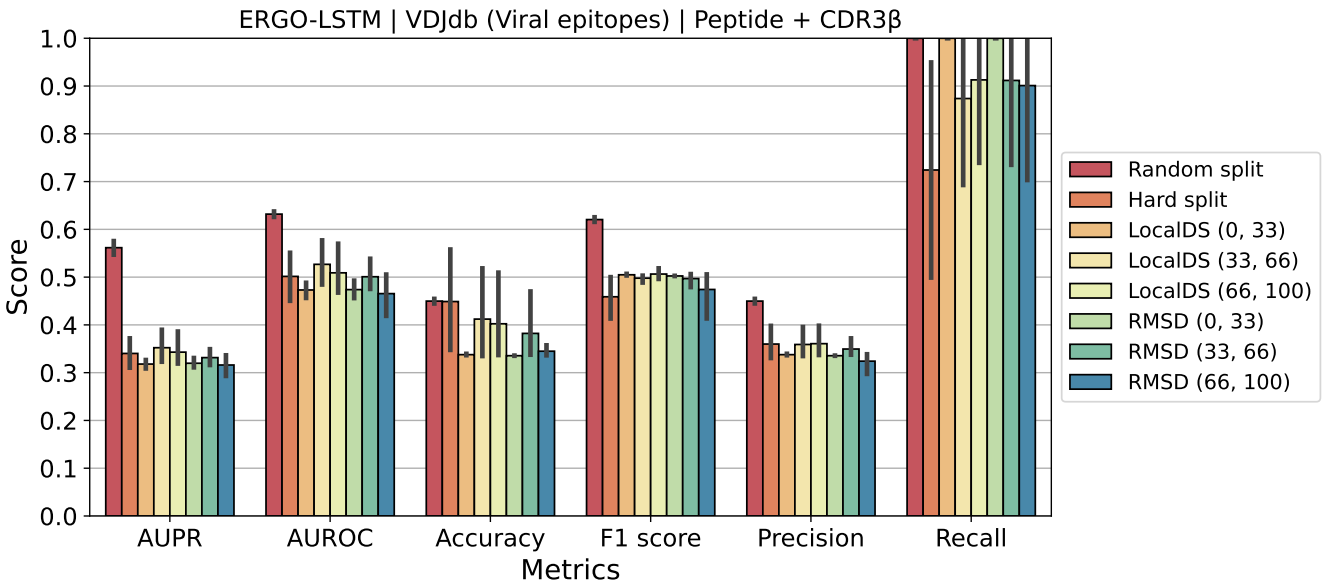

Figure S7. Performance metrics for the ERGO II model with the LSTM encoder [Springer et al., 2021a].

2.1 Macro (per-petide) Performance Scores for TCR-pMHC models

Presented below are per-peptide metrics. However, as the dataset contains many peptides with very few instances (e.g. 1-2) these metrics may not accurately represent the performance of the models.

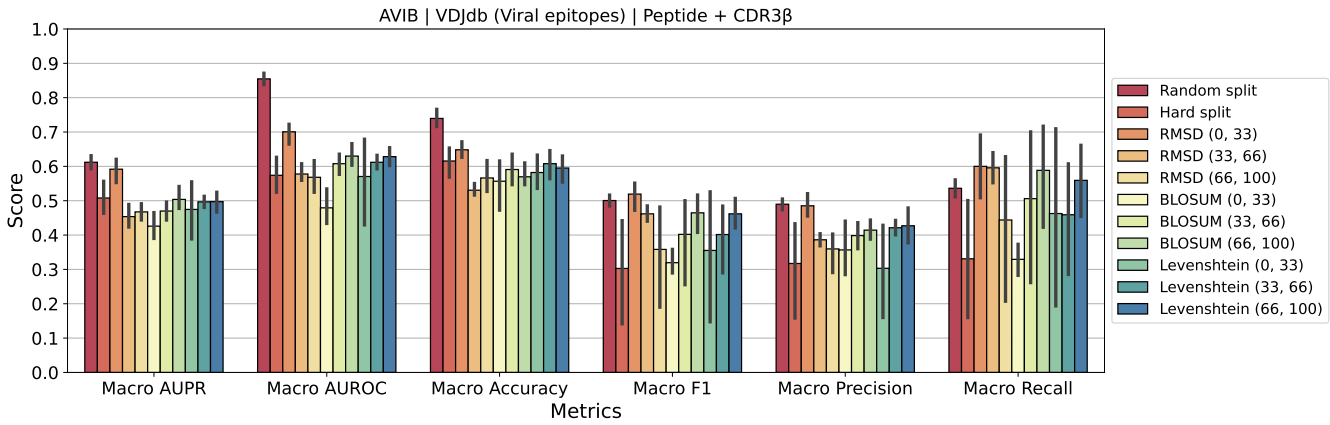

Figure S8. Macro performance metrics for the AVIB model [Grazioli et al., 2022a].

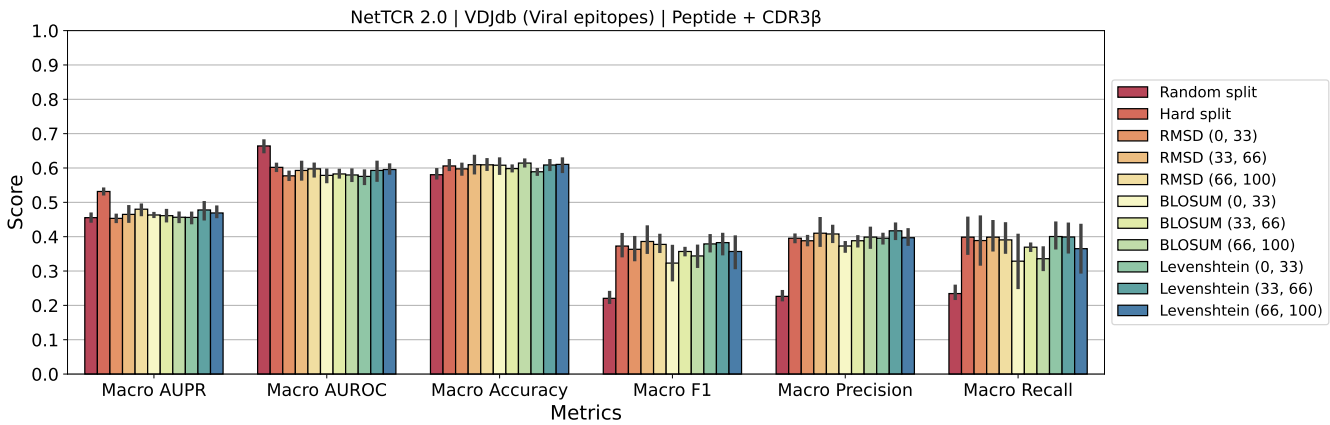

Figure S9. Macro performance metrics for the NetTCR-2.0 model [Jurtz et al., 2018].

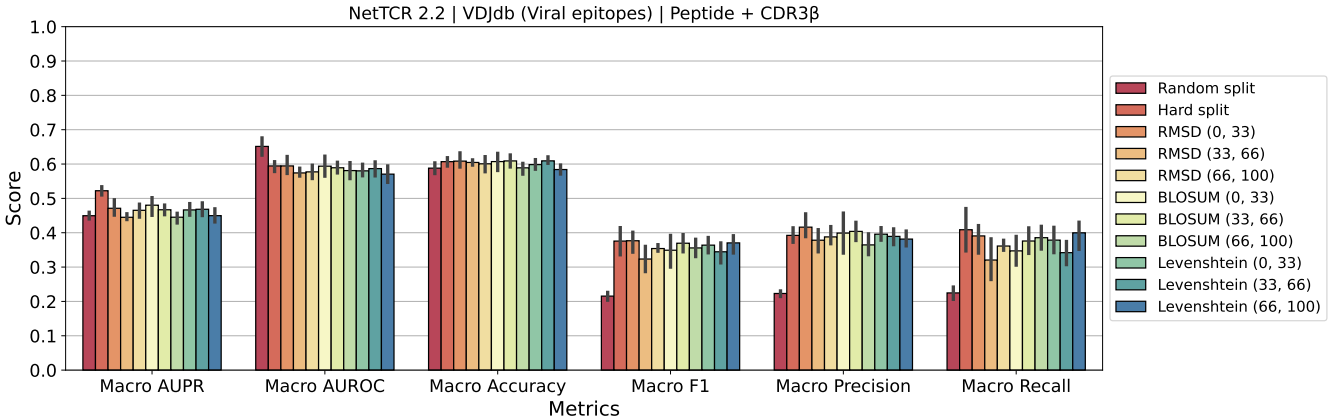

Figure S10. Macro performance metrics for the NetTCR-2.2 model [Jensen and Nielsen, 2024].

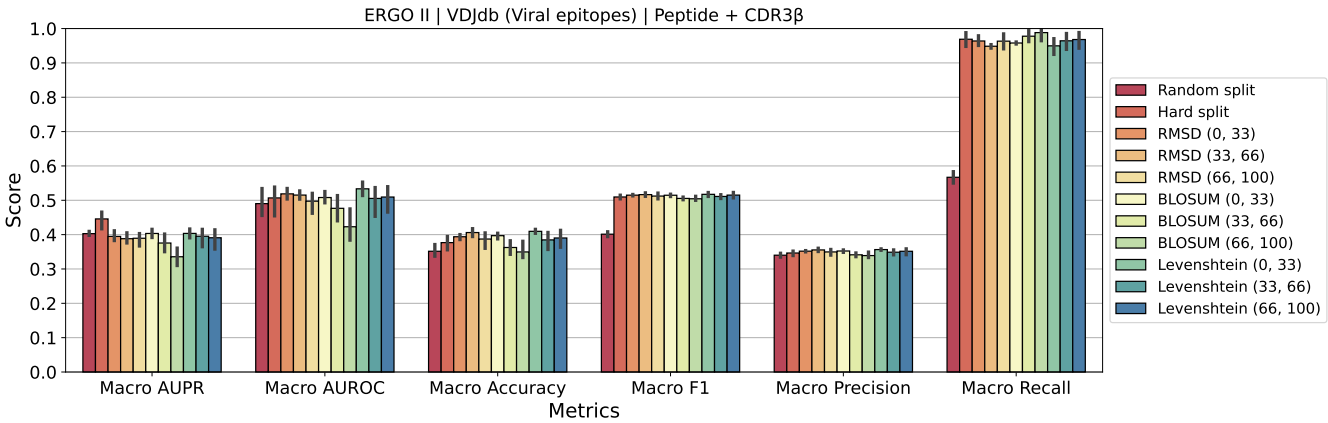

Figure S11. Macro performance metrics for the ERGO II model with autoencoder [Springer et al., 2021a].

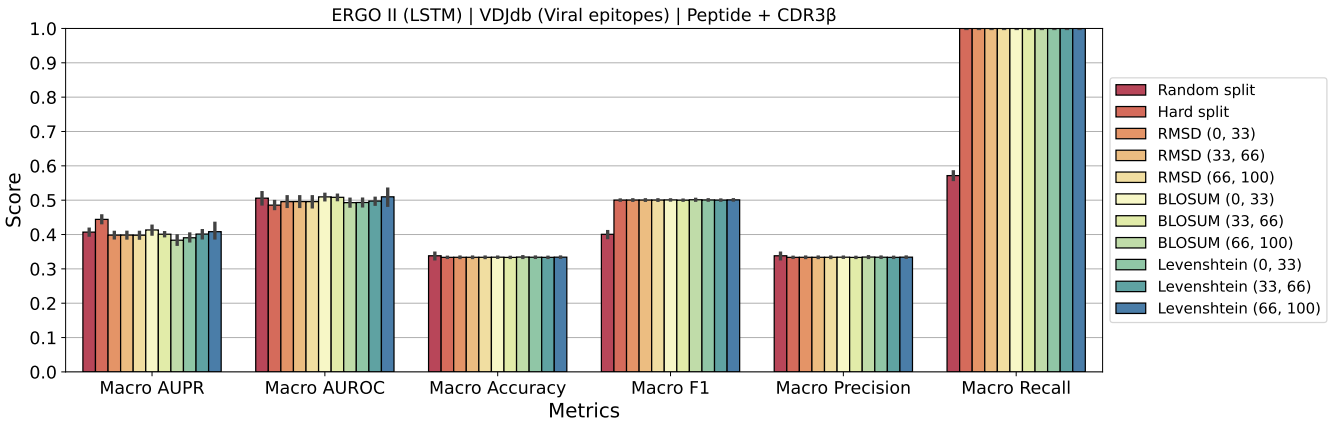

Figure S12. Macro performance metrics for the ERGO II model with the LSTM encoder [Springer et al., 2021a].

2.2 AUROC vs Difficulty Correlation

2.2.1 Sequence-based Performance

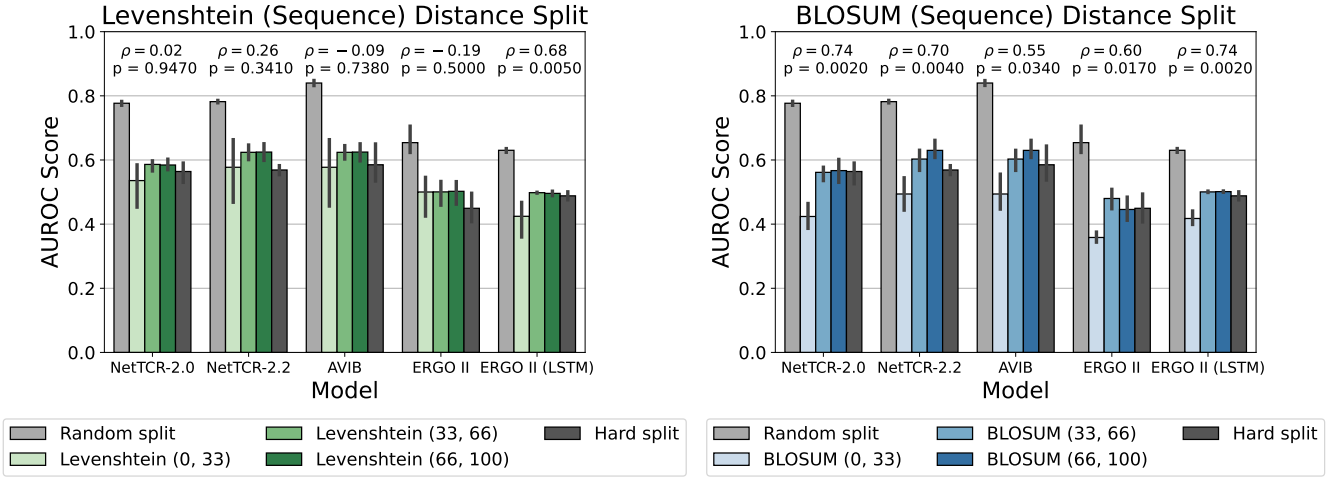

Figure S13. AUROC scores for TCR's CDR3β-peptide binding prediction for models trained and tested using various splitting techniques. The commonly used *Random Split* allows test peptides to be observed also during training. In the *Hard Split*, peptides are exclusively allocated to either the training or test set. In the *Distance Split*, we enforce specific median distances between the training and test peptides, using sequence-based distances (Levenshtein and BLOSUM). The training-test distance between peptides is controlled by selecting three percentile intervals over the cumulative median peptide-peptide distance distribution: $(0, 33)$, $(33, 66)$ and $(66, 100)$.

2.2.2 Correlation Overall (Micro)

Model	RMSD	BLOSUM	Levenshtein
AVIB	**-0.76 (0.001)**	**0.74 (0.002)**	0.02 (0.947)
NetTCR2.0	**-0.53 (0.043)**	**0.70 (0.004)**	0.26 (0.341)
NetTCR2.2	**-0.55 (0.034)**	**0.55 (0.034)**	-0.09 (0.738)
ERGO II	-0.45 (0.089)	**0.60 (0.017)**	-0.19 (0.500)
ERGO II (LSTM)	-0.08 (0.789)	**0.74 (0.002)**	**0.68 (0.005)**

Table S1. Spearman Correlation Coefficients for Micro AUROC (with p-values in parentheses)

2.2.3 Correlation Per Peptide (Macro)

Model	RMSD	BLOSUM	Levenshtein
AVIB	**-0.70 (0.004)**	**0.76 (0.001)**	0.08 (0.789)
NetTCR2.0	0.40 (0.143)	0.08 (0.789)	0.36 (0.189)
NetTCR2.2	-0.26 (0.341)	-0.21 (0.457)	-0.17 (0.544)
ERGO II	-0.25 (0.377)	**-0.57 (0.028)**	-0.13 (0.638)
ERGO II (LSTM)	0.51 (0.052)	**-0.68 (0.005)**	**0.81 (0.000)**

Table S2. Spearman Correlation Coefficients for Macro AUROC (with p-values in parentheses)

2.3 Entropy vs Difficulty Correlation

2.3.1 Overall (Micro)

Model	RMSD	BLOSUM	Levenshtein
AVIB	-0.06 (0.841)	0.43 (0.105)	0.17 (0.544)
NetTCR2.0	-0.02 (0.947)	-0.02 (0.947)	-0.04 (0.894)
NetTCR2.2	-0.28 (0.306)	0.23 (0.416)	0.51 (0.052)
ERGO II	0.32 (0.243)	**-0.76 (0.001)**	-0.34 (0.215)
ERGO II (LSTM)	0.25 (0.377)	0.23 (0.416)	-0.36 (0.189)

Table S3. Spearman Correlation Coefficients for Micro Entropy (with p-values in parentheses)

2.3.2 Per Peptide (Macro)

Model	RMSD	BLOSUM	Levenshtein
AVIB	-0.03 (0.502)	**0.39 (9.08e-25)**	0.03 (0.466)
NetTCR2.0	0.06 (0.107)	0.04 (0.279)	**-0.09 (0.009)**
NetTCR2.2	-0.02 (0.667)	0.07 (0.092)	**0.09 (0.010)**
ERGO II	0.02 (0.700)	**-0.27 (2.72e-12)**	-0.03 (0.388)
ERGO II (LSTM)	**0.13 (0.001)**	0.07 (0.072)	-0.04 (0.211)

Table S4. Spearman Correlation Coefficients for Macro Entropy (with p-values in parentheses)

REFERENCES

Filippo Grazioli, Pierre Machart, Anja Mösch, Kai Li, Leonardo Castorina, Nico Pfeifer, and Martin Renqiang Min. Attentive Variational Information Bottleneck for TCR-peptide Interaction Prediction. *Bioinformatics*, page btac820, December 2022a. ISSN 1367-4803, 1367-4811. doi: 10.1093/bioinformatics/btac820. URL https://academic.oup.com/bioinformatics/advance-article/doi/10.1093/bioinformatics/btac820/6960920.

Filippo Grazioli, Anja Mösch, Pierre Machart, Kai Li, Israa Alqassem, Timothy J. O'Donnell, and Martin Renqiang Min. On TCR binding predictors failing to generalize to unseen peptides. *Front. Immunol.*, 13:1014256, October 2022b. ISSN 1664-3224. doi: 10.3389/fimmu.2022.1014256. URL https://www.frontiersin.org/articles/10.3389/fimmu.2022.1014256/full.

Mathias Fynbo Jensen and Morten Nielsen. Nettcr 2.2 - improved tcr specificity predictions by combining pan- and peptide-specific training strategies, loss-scaling and integration of sequence similarity. February 2024. doi: 10.7554/elife.93934.2. URL http://dx.doi.org/10.7554/eLife.93934.2.

Vanessa Isabell Jurtz, Leon Eyrich Jessen, Amalie Kai Bentzen, Martin Closter Jespersen, Swapnil Mahajan, Randi Vita, Kamilla Kjærgaard Jensen, Paolo Marcatili, Sine Reker Hadrup, Bjoern Peters, et al. Nettcr: sequence-based prediction of tcr binding to peptide-mhc complexes using convolutional neural networks. *BioRxiv*, page 433706, 2018.

Ido Springer, Nili Tickotsky, and Yoram Louzoun. Contribution of t cell receptor alpha and beta cdr3, mhc typing, v and j genes to peptide binding prediction. *Frontiers in Immunology*, 12, 2021a. ISSN 1664-3224. doi: 10.3389/fimmu.2021.664514. URL https://www.frontiersin.org/journals/immunology/articles/10.3389/fimmu.2021.664514.

Ido Springer, Nili Tickotsky, and Yoram Louzoun. Contribution of t cell receptor alpha and beta cdr3, mhc typing, v and j genes to peptide binding prediction. *Frontiers in immunology*, 12:664514, 2021b.