

<https://doi.org/10.1038/s42003-025-07708-6>

T-cell receptor structures and predictive models reveal comparable alpha and beta chain structural diversity despite differing genetic complexity

Check for updates

Nele P. Quast ¹, Brennan Abanades ¹, Bora Guloglu ¹, Vijaykumar Karupiah ², Stephen Harper ², Matthew I. J. Raybould ¹ & Charlotte M. Deane ¹ ✉

T-cell receptor (TCR) structures are currently under-utilised in early-stage drug discovery and repertoire-scale informatics. Here, we leverage a large dataset of solved TCR structures from Immunocore to evaluate the current state-of-the-art for TCR structure prediction, and identify which regions of the TCR remain challenging to model. Through clustering analyses and the training of a TCR-specific model capable of large-scale structure prediction, we find that the alpha chain VJ-recombined loop (CDR3 α) is as structurally diverse and correspondingly difficult to predict as the beta chain VDJ-recombined loop (CDR3 β). This differentiates TCR variable domain loops from the genetically analogous antibody loops and supports the conjecture that both TCR alpha and beta chains are deterministic of antigen specificity. We hypothesise that the larger number of alpha chain joining genes compared to beta chain joining genes compensates for the lack of a diversity gene segment. We also provide over 1.5M predicted TCR structures to enable repertoire structural analysis and elucidate strategies towards improving the accuracy of future TCR structure predictors. Our observations reinforce the importance of paired TCR sequence information and capture the current state-of-the-art for TCR structure prediction, while our model and 1.5M structure predictions enable the use of structural TCR information at an unprecedented scale.

T-cell receptors (TCRs) are highly sequence diverse, somatically recombined receptor proteins. They govern the T-cell-mediated adaptive immune response by interacting with antigens such as peptides presented by the Major Histocompatibility Complex (pMHC)^{1–3}. The potential of TCRs as therapeutics for cancer and other diseases has led to substantial research effort dedicated to understanding and predicting their pMHC interactions and specificity^{4–20}. Despite the importance of this task²¹, TCR complementarity remains extremely difficult to predict beyond well-explored case studies, with existing computational models failing to generalise to unseen data^{22,23}. The development of more generalisable approaches was recently labelled a ‘Cancer Grand Challenge’ by the US National Institutes of Health²⁴.

Currently, most TCR:pMHC specificity predictors use sequence based features only^{4–18}. The incorporation of structural information offers a compelling strategy to encourage models to learn more generalisable physicochemical principles of complementarity^{25–27}. However, accessing and

generating structural information for TCR:pMHC specificity prediction poses new challenges, which is perhaps why only a handful of structure-aware methods exist to date^{20,25,28}. One bottleneck is the low quantity of available TCR structure data and the high experimental cost of obtaining more: only 605 publicly available crystal structures deposited in the PDB²⁹ at the time this study was conducted contained TCR structures³⁰, compared to the over 20,000 epitope-labelled paired TCR sequences available in VDJdb³¹. However, advances in computational TCR structure prediction provide a cheaper and higher throughput alternative to experimentally solving TCR structures. While historically the accuracy of adaptive immune receptor structure prediction algorithms has been poor^{32–34}, recent progress applying deep learning methods has made faster and higher-accuracy *in-silico* predictions of TCR structures viable^{35,36}. Simultaneously, federated learning initiatives are increasing accessibility to highly valuable datasets from the Pharmaceutical industry, which could potentially significantly boost the amount of available data for model training^{37,38}.

¹Department of Statistics, University of Oxford, Oxford, UK. ²Immunocore Ltd., Milton Park, Abingdon, UK. ✉e-mail: deane@stats.ox.ac.uk

A second limitation of many TCR:pMHC specificity predictors is their inclusion of only beta chain information of the TCR^{1,7,8,11,12,22,39,40}. TCRs function as heterodimers, typically consisting of an alpha and a beta chain whose recombination and sequence diversification process relies on the same genetic mechanisms underpinning the synthesis of the heavy and light chain of naïve B-cell receptors (BCRs). Briefly, sequence diversity results from the recombination of multiple genes (VDJ for beta/heavy chains, VJ for alpha/light chains) leading to six structurally proximal complementarity determining region (CDR) loops, three per chain³. While the CDR1 and CDR2 loops are relatively less sequence variable because they are templated by the V-genes, the CDR3 loops lie across the gene junctions giving rise to combinatorial diversity in their sequence. Further N- and P- nucleotide insertions and deletions at the junctions in the CDR3 lead to substantially greater length and sequence diversity⁴¹. Overall, it is understood that the VDJ-recombined beta/heavy chain CDR3 (CDR3 β /CDR3H) loop exhibits the greatest sequence diversity and thus contributes more to determining target specificity than the corresponding VJ-recombined CDR3 of the alpha/light chain (CDR3 α /CDR3L)^{3,42,43}. However, while there is strong evidence that this holds for BCRs/antibodies, where the CDR3H loop appears to contribute disproportionately to antigen specificity relative to the CDR3L loop^{44,45}, it remains unclear whether the TCR CDR3 β is more deterministic of antigen recognition than the CDR3 α ⁴⁶, implying that the exclusion of the alpha chain could lead to the loss of important information for specificity prediction.

Here, leveraging a proprietary dataset of TCR crystal structures solved by Immunocore to supplement the publicly available data, we evaluated the current headline performance of computational TCR modelling and identified which CDRs remain challenging to predict, representing barriers to the utility of structural models for predicting properties of TCRs such as pMHC specificity. We first clustered the complete structural dataset to identify relevant biases, which unexpectedly revealed that the CDR3 α and CDR3 β loops exhibit similar structural diversity, despite their differing recombination mechanisms.

We then built TCRBuilder2+ by retraining TCRBuilder2³⁶, a best-in-class TCR specific structure predictor, on the expanded dataset of structures. The performance of TCRBuilder2+ improved for genes better sampled in the new training set, and TCRBuilder2+ achieves comparable overall accuracy to AlphaFold Multimer at a fraction of the computational cost. Consistent with the clustering analysis, both TCRBuilder2+ and AlphaFold Multimer found the VJ-recombined CDR3 α loop at least as challenging to model as the corresponding VDJ-recombined CDR3 β loop, differentiating TCR structure prediction from BCR/antibody structure prediction, where the CDR3L is substantially easier to predict than the CDR3H^{36,47}. In order to further explore the structural diversity of TCRs we used TCRBuilder2+ to predict the structures of over 1.5 million TCR sequences from the Observed T-cell receptor Space (OTS) database⁴⁸.

Finally, we investigated potential sources of the high structural diversity in the CDR3 α . The amino acid composition and genetic coherence of distinct structures suggest that the greater number of TRAJ genes relative to both TRBJ and IGKJ/LJ genes, which leads to greater combinatorial diversity in the alpha chain, may also lead to enhanced structural variability.

Our study is the first to elucidate the surprisingly high structural complexity of the TCR CDR3 α and the challenge of predicting its structure, and reinforces the value of paired-chain information for capturing the function of TCRs.

Results

Curating and analysing a supplemented dataset of solved TCR structures

The amount of experimentally determined TCR structure data in the public domain is significantly less than TCR sequence data³⁰, therefore harnessing structural features in general TCR property prediction pipelines requires the ability to make accurate, and ideally fast, predictions of TCR structures. In recent years, deep learning methods have emerged as state of the art for this task³⁶. However, all machine learning models are limited by the data

available to train them on and there are countless examples of deep learning models that generalise poorly or that learn dataset biases rather than the underlying function governing the phenomenon^{23,49–51}.

In order to explore how increased structural data could improve TCR structure prediction, we supplemented the publicly available data with an internal repository of TCR structures solved by Immunocore, retraining our latest high-throughput TCR-specific deep learning model, TCRBuilder2³⁶ (Fig. 1a). As an initial step we analysed the novel dataset of TCR structures to identify any residual or newly introduced data biases and to enable fairer downstream model evaluation.

The original training data for TCRBuilder2³⁶ contained 704 experimentally resolved TCR structures sourced from STCRDab^{30,52}. These structures were obtained from 463 distinct PDB files, with 349 of the 704 structures having a unique sequence. Our novel data contained 204 Immunocore TCR structures, of which 201 sequences are unique and only 6 overlap with the existing training set (Fig. 1b). This expansion contributed to a combined set of 908 structures, of which 544 are unique in sequence. Filtering for alpha-beta TCR pairs resulted in 881 TCR structures to train TCRBuilder2+, of which 524 are sequence non-redundant. Furthermore, since the original TCRBuilder2 was published more TCR structures have become publicly available, allowing for the curation of a larger and more robust test set. We generated a new, independent, and non-redundant test set of 45 alpha-beta TCR structures for downstream benchmarking (SI Table 1).

Figure 1c shows the distribution of TRAV and TRBV gene subgroups of the 881 TCR structures contained in our new training set, characterising the biases present in the data. We report the distribution at the level of individual genes in SI Fig. 1. There are strong biases towards a small number of genes and overall the joint distribution over pairs is sparsely sampled, with the majority of alpha/beta pairs unsampled or only appearing in the dataset once. The joint distribution is disproportionately weighted towards a small subset of gene pairs; for example, TRAV12-2 paired with TRBV6-5 and TRAV1-1 with TRBV6-1 account for 79 and 59 structures respectively, 15.7% of the training data.

Observations of naturally occurring alpha/beta chain pairings imply variable gene pairings are relatively uniform after accounting for gene abundance⁴⁸, indicating that it is the selection of TCR structures that have been experimentally resolved that is giving rise to the TRAV and TRBV gene bias, not the natural underlying TCR distribution. To disentangle the contributions of the STCRDab and the Immunocore datasets to the bias we also considered the distributions independently (SI Fig. 2). While we observed similar biases within each set, our new dataset improves the sampling of human TRAV-TRBV gene subgroups and gene pairs. We saw a similar trend of bias and improved coverage in our new dataset for TRAJ-TRBJ pairs, as well as VJ pairings, shown in SI Figs. 3–5.

Computationally predicting TCR structures from their sequences using data-driven methods such as deep learning or homology relies on there being a functional mapping from sequence to structure that can be extracted from the existing data. Therefore, we used the supplemented structural data to investigate the relationship between sequence and structure identity of most variable TCR loops, the CDR3 α and CDR3 β (Fig. 1d–e).

There is a disparity in trend between sequence identities of CDR3 loops greater than 65% and between 0% and 65%. At lower sequence identities, structural distances are weakly correlated (PCC -0.17 to -0.29). Higher sequence identities are moderately correlated with smaller structural distances (PCC -0.52 to -0.66), though, even between loops of identical sequence, substantial structural deviation can be observed: CDR3 β and CDR3 α loops with very high sequence identities ($\geq 90\%$) can display structural dissimilarities of up to 4 Å (Fig. 1d, e, SI Fig. 6). We observed a similar trend when considering the pairwise LDDT distribution against sequence identity (SI Fig. 7). These results demonstrate the known limitations of homology approaches to TCR structure prediction, as there is often more complexity in what determines CDR3 structure than sequence identity of the loop alone. Deep learning approaches, such as the

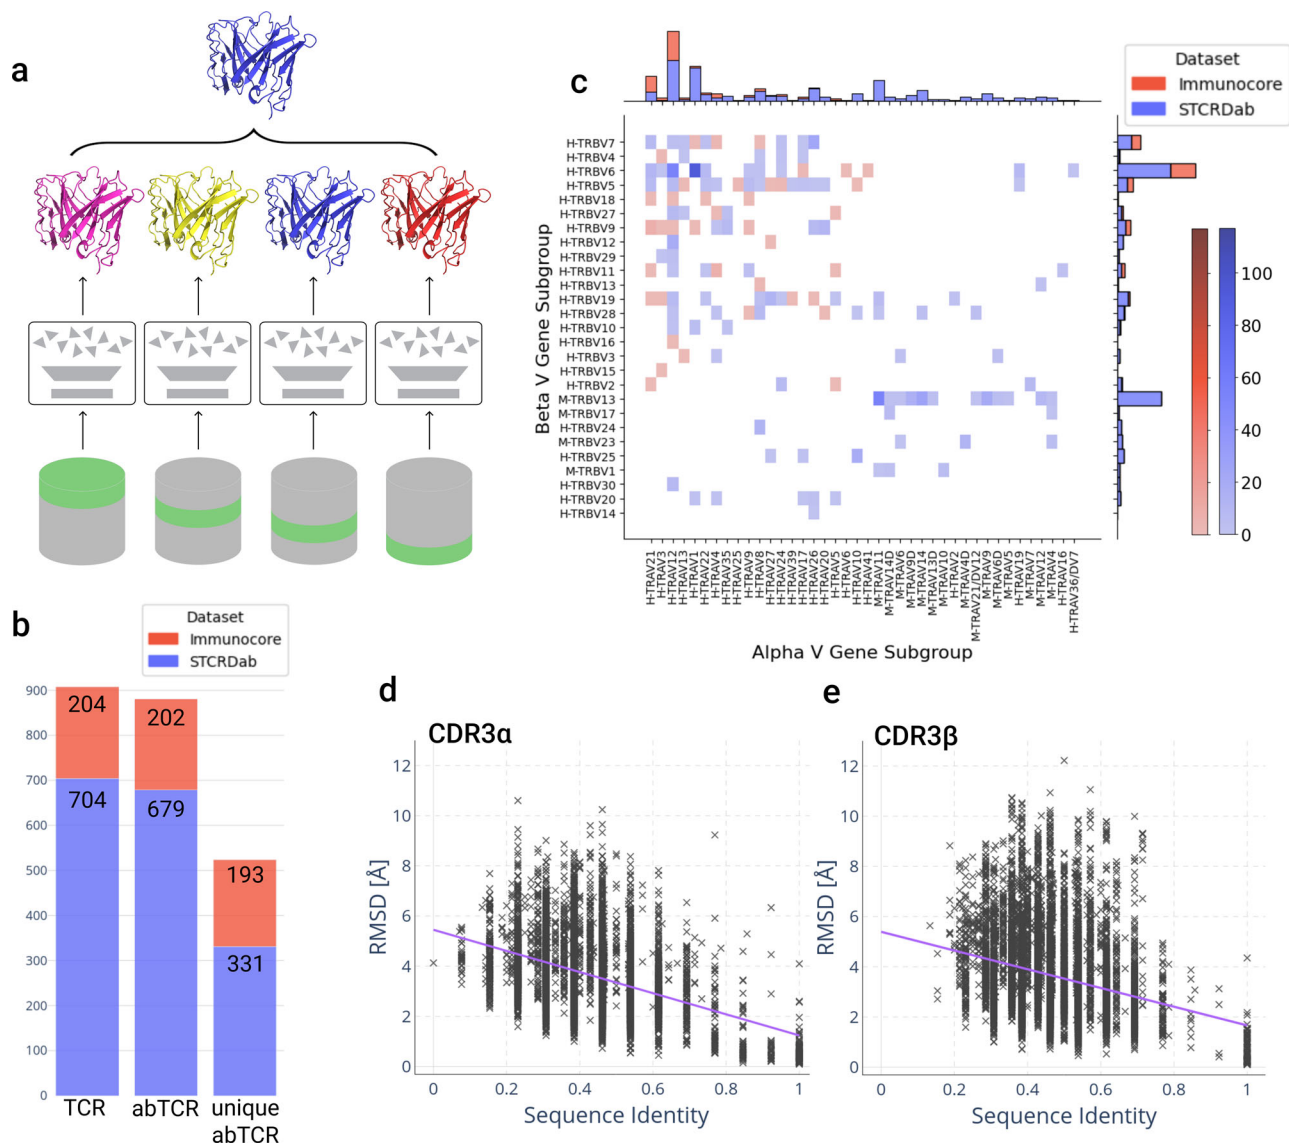


Fig. 1 | Overview of the architecture and training data distribution of TCRBuilder2+. **a** TCRBuilder2 architecture³⁶ for TCR structure prediction. An ensemble of models is trained using different training / validation splits to generate four predictions per sequence. The final model is chosen as the model closest to the mean prediction of the ensemble. **b** Adding Immunocore data to the STCRDab data expands the training set by 204 structures, and increases the number of non-redundant sequences by 58%. **c** The joint distribution over paired alpha and beta V gene subgroups of TCR structures in the training set; the accompanying gene level distribution is reported in SI Fig. 1. The paired space is biased and sparsely

sampled. Genes beginning with 'H-' are human, those beginning with 'M-' are murine. **d, e** Pairwise RMSD of CDR3 α (**d**) and CDR3 β (**e**) loops in the training data against pairwise sequence identity with a linear trend fit (purple, $PCC_{CDR3\alpha} = -0.41$, $PCC_{CDR3\beta} = -0.27$). For sequence identities > 0.65, sequence identity and RMSD of the CDR3 loops are somewhat correlated ($PCC_{CDR3\alpha} = -0.66$, $PCC_{CDR3\beta} = -0.52$). Sequence identities < 0.65 are substantially less correlated with RMSD for both CDR3 loops ($PCC_{CDR3\alpha} = -0.29$, $PCC_{CDR3\beta} = -0.17$). CDR3 β displays a slightly greater range of RMSD than CDR3 α loops.

TCRBuilder2 architecture, have been shown empirically to be able to capture more subtle and complex patterns from data, yielding improvements in predictive capability^{35,36,47}.

The sequence identity to structural distance distributions in Fig. 1d and e also provide useful background distributions against which to evaluate the predictive performance of TCR structure prediction methods, since they provide an upper bound of the expected RMSD for a given prediction if a template with the closest sequence identity to the query sequence were selected. We observed a wide range of structural distances between CDR3 loops across all sequence identities, many of which are substantially greater than 2.0 Å.

Finally, comparing the pairwise RMSD distributions of the CDR3 α and CDR3 β loops reveals that the CDR3 α loop (Fig. 1d) is as inherently structurally diverse as the CDR3 β loop (Fig. 1e, SI Fig. 8). We confirmed this finding by clustering each of the CDR3 loop types by RMSD (SI Table 2).

Surprisingly, despite their simpler genetic composition CDR3 α loops do not cluster into fewer structural clusters than CDR3 β loops; by contrast, the CDR3L loops of antibodies cluster into far fewer shapes than their CDR3H counterparts^{53–56}. This disparity between TCR and antibody structures hints that the structures adopted by CDR3 α loops may be more complex than suggested by their genetic recombination mechanism alone.

Benchmarking state-of-the-art TCR structure prediction models

We retrained TCRBuilder2 on our supplemented dataset to generate a new model ('TCRBuilder2+'), capturing the current performance ceiling of a TCR-specific deep learning model for high-throughput structure prediction. While the original TCRBuilder2 model followed from ABodyBuilder2³⁶ and selected the ensemble of trained models based on only the beta chain validation set accuracy, in this work we adapted the ensemble selection process to include the accuracy over both chains. This adaptation

Table 1 | Mean RMSD of TCR structure predictions by different prediction models evaluated on the test set

	α	FW α	CDR α 1	CDR α 2	CDR α 3	β	FW β	CDR β 1	CDR β 2	CDR β 3
TCRBuilder2	1.14	0.91	1.21	0.81	2.02	0.92	0.69	0.59	0.70	1.84
TCRBuilder2+	1.20	0.98	1.19	0.98	2.07	0.92	0.70	0.71	0.75	1.78
Alphafold Multimer	1.03	0.73	1.09	1.02	1.95	0.90	0.73	0.58	0.57	1.52
Alphafold3	0.97	0.71	0.93	0.83	1.81	0.82	0.59	0.51	0.58	1.57
TCRmodel	1.33	0.84	1.60	0.96	2.81	0.94	0.65	0.78	0.71	2.02
LYRA ^a	1.21	0.78	1.54	0.81	2.78	0.99	0.70	0.81	0.80	2.28
ABodyBuilder2	3.52	2.99	3.56	7.19	3.35	3.35	3.27	3.35	3.77	3.35

^aLYRA could only be evaluated on 38 of the 45 test structures for which a structure was generated. Statistically significantly differentiable RMSD values are italicised (for *p*-values see SI Table 3) and the best mean RMSD for any TCR region is emboldened. Backbone RMSD is reported for each region of the TCR. α/β —the entire alpha and beta chain respectively, FW α /FW β —the constant framework region of each chain, CDR α /CDR β —the three CDR loops of each chain.

better reflects TCR biology by accounting for the high structural diversity we observed in the alpha chain.

We benchmarked TCRBuilder2+ against the original TCRBuilder2³⁶, TCR-specific homology modellers (LYRA³⁴, and TCRmodel³³), an antibody-specific deep learning model (ABodyBuilder2³⁶), and the state-of-the-art general protein structure predictor Alphafold (Alphafold-Multimer⁵⁷ and Alphafold3⁵⁸). We report the mean RMSD and the RMSD distribution of the 45 test set structures predicted by different models in Table 1 and Fig. 2a, respectively.

As expected, predicting TCR structures with ABodyBuilder2 yields low accuracy predictions across all TCR regions despite the global similarities between TCRs and antibodies, confirming that structure prediction models trained on protein family subsets learn to make accurate predictions within the distribution of the data they are trained with, rather than generalising beyond it⁵⁹. The homology model LYRA³⁴ only produces predictions for sequences for which a suitable template is found, and so we are only able to benchmark 38 of the 45 test structures. We provide a comparison for those 38 predictions only in SI Figure 9. Overall, homology modelling with LYRA and TCRmodel³³ performs competitively for the framework regions (FW) and the germline encoded CDR2 α , CDR1 β and CDR2 β , but struggles with the CDR3 α , CDR3 β , and CDR1 α loops (Table 1, Fig. 2a).

TCRBuilder2, TCRBuilder2+, Alphafold Multimer, and Alphafold3 all exhibit competitive performance across the whole test set. The difference in mean RMSD errors as reported in Table 1 is less than 0.25 Å across all regions, although Alphafold3 exhibits a slightly lower mean across most regions. Similarly, the span of median RMSD across all four deep learning models is very tight (Fig. 2a) with overlapping error bars. It is worth noting that Alphafold3 has a later training data cutoff, and it is therefore not possible to exclude test/train overlap inflating its performance slightly.

TCRBuilder2+ performs similarly to TCRBuilder2 (Table 1), suggesting that current bottlenecks in TCR structure prediction are not alleviated by our supplemented training dataset. We have reported the *p*-values expressing the statistical significance of the differences between the models' performance (SI Table 3), and found that, with the exception of ABodyBuilder2 which significantly underperforms, it is not possible to distinguish the methods using a 5% threshold across most regions.

Based on the median RMSD, TCRBuilder2+ prediction accuracy is on par with TCRBuilder2 and is only slightly outperformed by Alphafold-Multimer (Fig. 2a). Examining how frequently each method makes a prediction of reasonable accuracy (<2 Å, see Methods), we found that the framework regions (FW) and germline-encoded loops (CDR1 and CDR2) are largely well predicted. As was observed for the homology model, CDR1 α is markedly more difficult to predict than CDR1 β , with mean TCRBuilder2+ prediction RMSDs of 1.25 ± 0.81 Å and 0.75 ± 0.41 Å respectively.

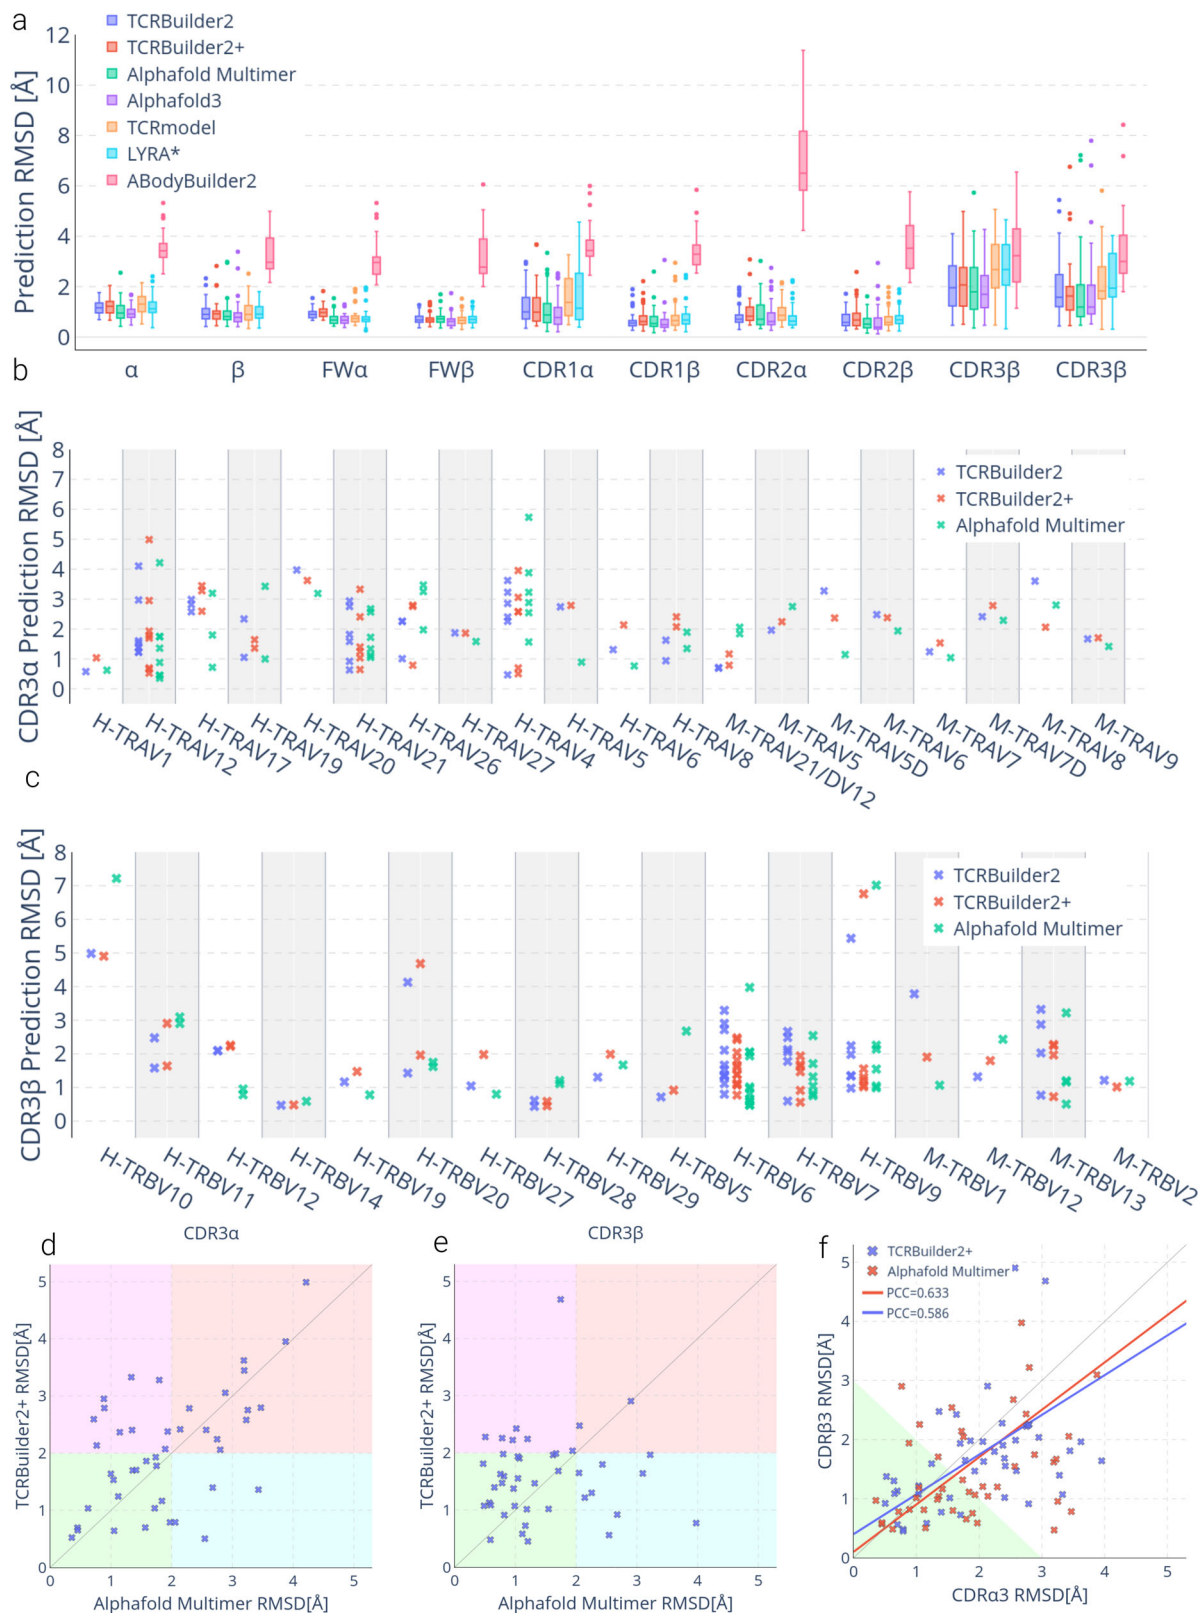
Meanwhile, the CDR3 β and CDR3 α loop predictions exhibit the largest errors, as is expected due to their greater sequence and structure diversity relative to the germline-encoded loops.

Specifically, TCRBuilder2+ predicts 34 (75.6%) CDR3 β structures with sub-2 Å RMSD, more than TCRBuilder2, Alphafold Multimer, and Alphafold3 which predict 26 (57.8%), 32 (71.1%), and 33 (73.3%) within that threshold, respectively (SI Table 4). On the other hand, Alphafold Multimer is able to predict 27 (60.0%) of CDR3 α structures within 2 Å RMSD, whereas TCRBuilder2+, TCRBuilder2, and Alphafold3 respectively retrieve 21 (46.6%), 23 (51.1%), and 25 (55.6%) CDR3 α structures within that threshold (SI Table 4). Every model benchmarked (with the exception of ABodyBuilder2) predicts CDR3 β loops with a higher mean accuracy than CDR3 α loops (Table 1, Fig. 2a), once again contrary to the conjecture that the greater junctional diversity of CDR3 β compared to CDR3 α ought to lead to higher structural complexity.

Due to the biases towards specific gene pairings in the training set (Fig. 1c), we then analysed the CDR3 α and CDR3 β predictive accuracy over different V-genes in the test set (Fig. 2b, c). Contrasting the two TCR-specific models, we find that predictive accuracy improves for genes that are better represented in the new training set (Fig. 1c, SI Fig. 10, 11), with TCRBuilder2+ resulting in, on average, higher accuracy CDR3 α predictions than TCRBuilder2 for human TRAV19 and -21 and CDR3 β predictions for human TRBV6, -7, and -9. This suggests that more data for a specific gene helps predict the structures of TCRs that derive from that gene, but that this improved performance might not necessarily translate to all TCRs generally. This in turn implies that increasing the amount and improving the genetic sampling of experimentally solved structure data with which models can be trained will be necessary to improve general inference of TCR structures.

To assess whether TCR-specific and general structure predictors struggle to model the same TCR sequences, we compared the RMSD of the TCRBuilder2+ and Alphafold Multimer CDR3 α and CDR3 β loop predictions for every test structure (Fig. 2d–e). Points in the green quadrant or red quadrant indicate both methods are able to (green) or unable to (red) predict the structure to within 2 Å, while points in the purple quadrant or cyan quadrant indicate that Alphafold Multimer (purple) or TCRBuilder2+ (cyan) is uniquely able to model the loop to within 2 Å. We observe for CDR3 α loops that a substantial number (13/45) fall into the red quadrant, indicating that neither TCRBuilder2+ nor Alphafold Multimer can accurately predict their conformation. In contrast, very few CDR3 β loops fall into the red quadrant (2/45), reiterating that CDR3 α structure is more challenging to predict regardless of the model used.

A hypothetical ensemble model able to perfectly identify whether the TCRBuilder2+ or Alphafold Multimer prediction is a better CDR3 β or CDR3 α loop model—independent of the other loop—would currently achieve 95.6% under 2 Å for CDR3 β (up from 75.6% for TCRBuilder2+ alone) or 71.1% under 2 Å for CDR3 α (up from 60.0% for Alphafold-Multimer alone). This suggests these models have captured different information, and that a method that can selectively incorporate knowledge of general protein loop conformations into TCR loop structure prediction may translate to larger performance boosts for CDR3 β than CDR3 α . Furthermore, the fact that even when combining the best predictions of both TCRBuilder2+ and Alphafold-Multimer the CDR3 α structure is more



elusive to predict emphasises its structural complexity despite its simpler genetic recombination.

As an indication of whether deep learning approaches have learned to model the TCR variable region globally rather than considering loops independently of one another, we investigated the correlation between CDR3 α and CDR3 β predictive accuracy (Fig. 2f). The Pearson’s correlation

coefficients (PCC) calculated for both AlphaFold Multimer (0.633) and TCRBuilder2+ (0.586) indicate a slight correlation between CDR3 α and CDR3 β prediction RMSD, suggesting that deep learning models consider loop conformations somewhat cooperatively. While comparing the RMSD of the CDR3 α and CDR3 β of each TCR in the test set, we found that 64.4% (29/45) and 60% (27/45) of TCRBuilder2+ and AlphaFold Multimer’s

Fig. 2 | Evaluation of TCR structure prediction models on our test set (SI Table 1). **a** Structure prediction RMSD by TCR region. Sample size of 45; boxes represent the inner quartiles, with the median value the dividing line; whiskers extend to points that lie within 1.5 times the interquartile range of the lower and upper quartile; observations beyond this range are displayed independently. *LYRA could only be evaluated on 38/45 structures for which a structure was generated. **b, c** Structure prediction accuracy over **(b)** CDR3 α by TRAV gene, and **(c)** CDR3 β by TRBV gene for TCRBuilder2 (blue), TCRBuilder2+ (red) and Alphafold Multimer (green). Genes beginning with 'H-' are human, those beginning with 'M-' are murine.

predictions respectively fall below the RMSD identity line (Fig. 2f), confirming that for the majority of test structures the CDR3 α is more difficult to model than the CDR3 β of a given TCR, regardless which model is used.

Finally, we explored whether TCR structural models have yet attained a level of accuracy where useful information can be extracted from TCR:pMHC interfaces predicted using physics-based docking⁶⁰ and present the extended results of these experiments in the Supplementary Information (SI Note 1). We found that a total RMSD over both the CDR3 α and CDR3 β < 3 Å strongly supports the retrieval of an accurate docked pose (interface RMSD over all heavy atoms < 4 Å in all cases we ran simulations for, indicating that structure predictions are beginning to show merit for TCR:pMHC interface prediction. Our evaluation of current state of the art structure prediction models suggests a third of TCRs are modelled within this threshold (green region in Fig. 2f), indicating further improvements are required to enable consistent interface prediction *via* docking. It is not possible to identify a consistent accuracy threshold for docking when considering the CDR3 α or CDR3 β individually (SI Fig. 12–14), highlighting once again the importance of accurately modelling both regions for TCR:pMHC interface prediction.

In summary, retraining TCRBuilder2 with substantially more training data did not yield large gains in average predictive accuracy, although the mean RMSD for CDR3 β reduced from 1.84 Å (TCRBuilder2) to 1.78 Å (TCRBuilder2+), narrowing the gap to Alphafold Multimer's benchmark of 1.52 Å. We did observe improvements across TRAV and TRBV genes for which the abundance of training examples increased, suggesting that increasing the diversity of training examples by experimentally solving as yet unsampled genes should yield accuracy improvements. Physics-based docking studies indicate that, while predicted TCR structures are beginning to show utility for predicting TCR:pMHC interfaces, improvements in TCR structure prediction accuracy should also yield improvements in docking. Furthermore, and in line with our observations of experimentally solved TCR structures, we found that the CDR3 α is consistently more challenging to predict than the CDR3 β , with evaluations of all models suggesting that the genetically simpler CDR3 junction region, CDR3 α , remains a key barrier to accurate TCR modelling.

Enabling structure-based deep learning research with 1.5 million TCR structure predictions

Next, we estimated to what degree the apparent diversity of alpha chain CDR structures is sampled by natural TCR repertoires. To assess this, we harnessed the computational efficiency of TCRBuilder2+ and predicted the structures of over 1.56M non-redundant TCR sequences from the newly released OTS database⁴⁸, yielding an approximately 2500-fold increase in TCR structures in the public domain. We then conducted a structural analysis, contrasting the TCRs in OTS to the BCRs/antibodies from OAS⁶¹. Our previous analysis of crystal structures suggested that the CDR3 α of TCRs is as structurally diverse as the CDR3 β ; here, we observe that TCR CDR3 α loops from T cell repertoire sequencing also exhibit this diversity.

Figure 3a and b show that the CDR1 and CDR2 loops of the TCR alpha and beta chains (a), and the antibody light and heavy chains (b) cluster into canonical forms, as has been shown to be the case for experimentally solved structures^{62,63}. Specifically, we find that for 1000 randomly sampled TCRs from our OTS structure predictions, the CDR1 α , CDR2 α , CDR1 β , and CDR2 β loops fall into 126, 200, 34, and 105 clusters, respectively, when applying a 1 Å RMSD threshold (SI Table 5). Repeating this clustering with

d, e TCRBuilder2+ vs Alphafold Multimer prediction RMSD of CDR3 α (d) and CDR3 β (e). Points in the green quadrant or red quadrant indicate both methods are able to (green) or unable to (red) predict the structure to within 2 Å, respectively, while points in the purple quadrant or cyan quadrant indicate that Alphafold Multimer (purple) or TCRBuilder2+ (cyan) is uniquely able to model the loop to within 2 Å. (f) CDR3 β vs CDR3 α prediction RMSD. For reference, subsequent docking case studies suggest that predictions within the green shaded region (CDR3 α RMSD + CDR3 β RMSD < 3 Å) can be reliably docked with an interface RMSD < 4 Å (see Section 2.4, SI Fig. 12a).

1000 antibodies from OAS, we find that the CDR1L, CDR2L, CDR1H, and CDR2H separate into 31, 9, 19, and 36 distinct loop conformations (SI Table 6).

In agreement with our analysis of experimentally solved TCR structures, neither the CDR3 α nor CDR3 β of naturally occurring OTS TCRs cluster into distinct structures: 1000 random TCRs fall into 840 and 828 clusters for CDR3 α and CDR3 β respectively (Fig. 3c). As expected, the antibody CDR3H also does not cluster into distinct shapes (880 clusters per 1000 antibodies), but the CDR3L loop does (131 clusters per 1000 antibodies). Comparing the clustering of antibody CDR3L to TCR CDR3 α structures exemplifies the increased structural diversity of the CDR3 α despite both loop sequences being the product of only VJ recombination.

The pairwise RMSD distributions of the CDR3 α and CDR3 β of the OTS predictions are indistinguishable (Fig. 3d), confirming that TCRBuilder2+ generates equally diverse conformations of these regions when predicting sequences in OTS. On the other hand, the RMSD distribution of the antibody CDR3L is shifted towards lower RMSDs relative to CDR3H and exhibits the bimodality expected of loops sampled from distinct canonical shapes; this bimodality is not observed for the TCR CDR3 α . These results are recapitulated if we consider loop length matched distributions of RMSD (SI Figs. 15–17).

Together, this analysis of predicted TCR and antibody structures demonstrates the difference in structural diversity of the VJ recombined CDR3 α and CDR3L loops. This concurs with our finding from the analysis of TCR crystal structures and the evaluation of TCR prediction models that CDR3 α structural complexity matches that of CDR3 β .

Potential explanations for CDR3 α structure diversity

Finally, we sought an explanation for the unexpectedly high structural diversity of CDR3 α by investigating the CDR3 α and CDR3 β sequence distributions (Fig. 4). Despite displaying similar overall sequence diversity (SI Fig. 8), CDR3 α sequences cluster by sequence identity, whereas the CDR3 β sequences do not (Fig. 4a, b). Greedy clustering confirms this: fewer CDR3 α than CDR3 β clusters emerge for sequence thresholds between 50% and 100% (Fig. 4c, SI Table 9), suggesting that CDR3 α sequences can be separated into sets of similar sequences.

To expand on this observation, we investigated the sequence logos of the CDR3 α (Fig. 4d) and CDR3 β (Fig. 4e), finding as expected that the central regions of both loops are highly diverse. However, the variability of the central CDR3 β residues is greater than that of the central CDR3 α residues, implying that diversity emerging from VDJ recombination with the additional junctions relative to the VJ recombination does increase the sequence diversity of the CDR3 β relative to CDR3 α . Both the clustering and the slightly lower variability of the CDR3 α sequences contrasts with our observations of CDR3 α and CDR3 β structures, where CDR3 α exhibits similar diversity to CDR3 β .

Both the CDR3 α and CDR3 β loop sequences have high glycine and serine content (Fig. 4d and e). Following from this observation we investigated the proportion of residues that could be conducive to flexibility or the emergence of multiple conformations per sequence, which could provide a potential explanation for structural diversity. Figure 4f and g suggest that the proportion of small residues conducive to loop flexibility is approximately 50% for both the CDR3 α (f) and CDR3 β (g). Furthermore, studies of TCR structure^{64–66} and *in-silico* simulations^{66,67} have suggested that both the

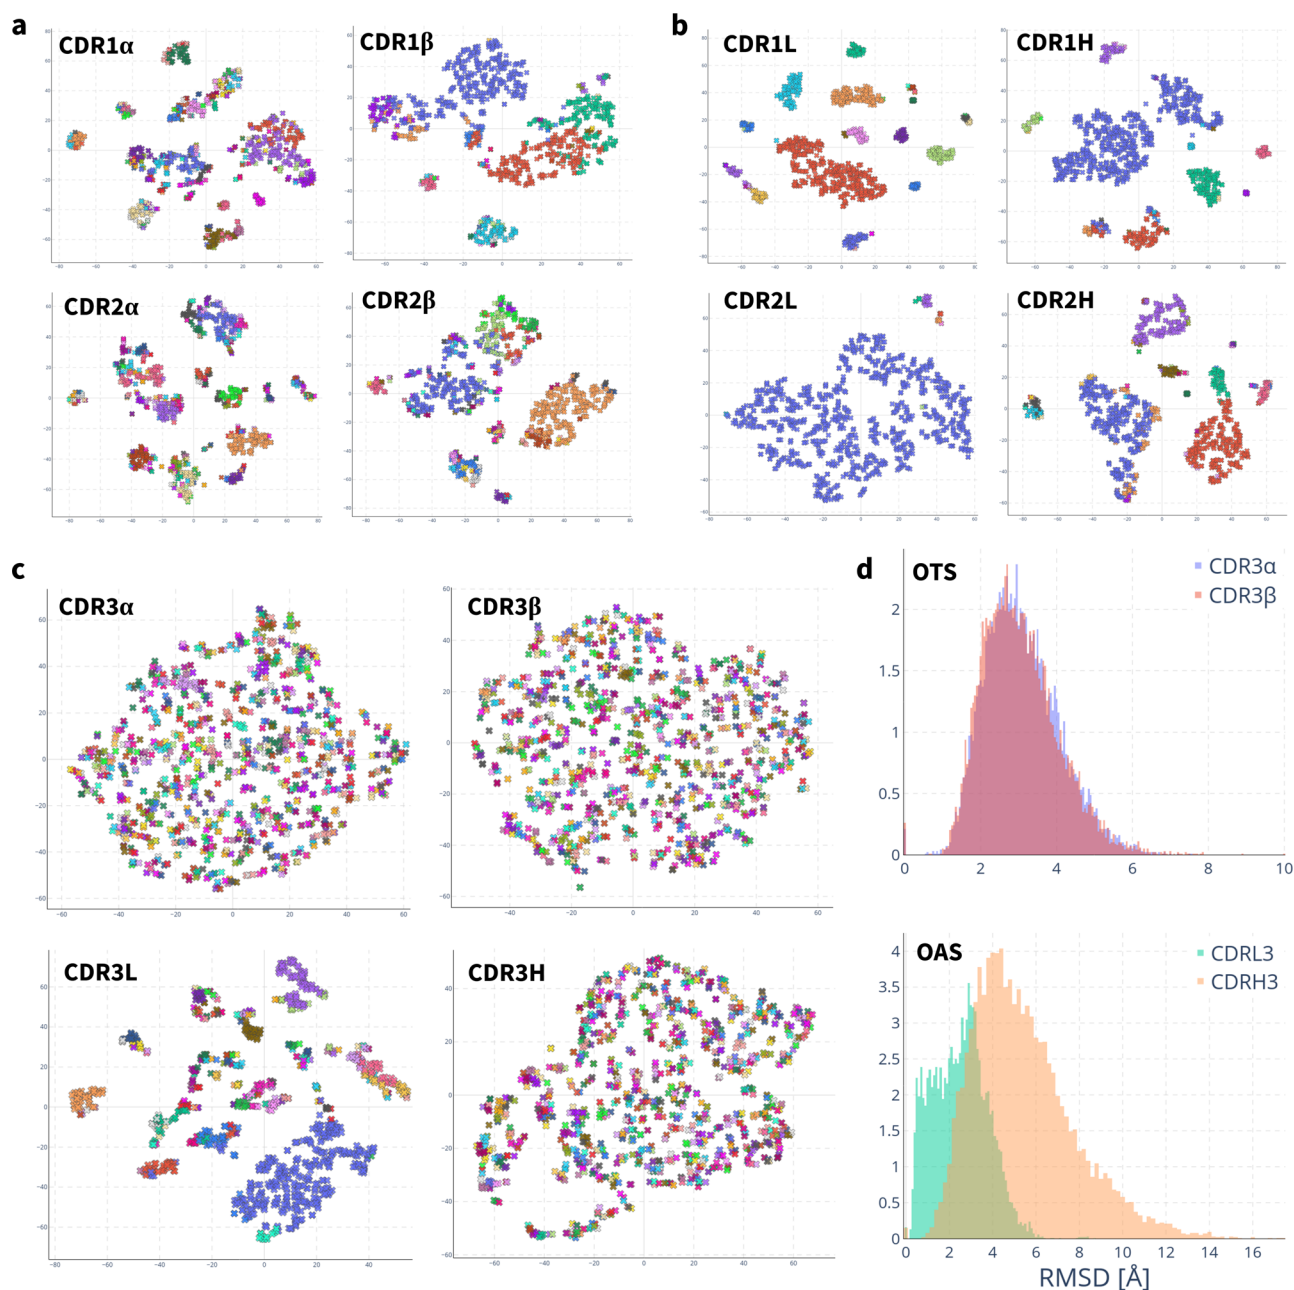


Fig. 3 | Structure distributions of 1000 randomly sampled TCRBuilder2+ predictions from the OTS database⁴⁸ and 1000 randomly sampled ABodyBuilder2³⁶ predictions from the OAS database⁶¹. All t-SNE plots are coloured by clustering with a 1 Å RMSD threshold. **a, b** t-SNE clustering of RMSD between canonical loops of TCRs (**a**) and antibodies (**b**). CDR1- and 2 visibly cluster into canonical forms. **c** t-SNE clustering of RMSD between TCR CDR3 α and CDR3 β , and antibody CDR3L and CDR3H. The antibody CDR3L loop clusters into canonical forms, whereas the TCR CDR3 α does not. No clustering is observed for the TCR CDR3 α ,

CDR3 β or the antibody CDR3H. **d** Distribution of pairwise RMSD between CDR3 α and CDR3 β predictions of OTS and between CDR3L and CDR3H predictions of OAS. The distribution of pairwise RMSD of predicted TCRs is indistinguishable between CDR3 α and CDR3 β , demonstrating equivalent structural diversity in natural repertoires. In antibodies the pairwise RMSD distribution of CDR3L is shifted towards lower RMSD than that of CDR3H, indicative of CDR3L structures with high structural similarity, as would be expected of canonical loops.

CDR3 α and CDR3 β structures may be flexible. Though it is difficult to analyse robustly computationally, we explored this hypothesis *via* a preliminary molecular dynamics case study of an apo TCR, the results of which we provide in the Supplementary Information (SI Note 2). Both TCRBuilder2+ and AlphaFold Multimer produce an ensemble of structure predictions (SI Fig. 18, 20), enabling us to initialise the simulations from distinct predicted coordinates. In our case study we found that an array of conformations of the CDR3 α and CDR3 β loops predicted by TCRBuilder2+ and AlphaFold Multimer appear to be physically plausible local energy minima of the simulations (SI Fig. 21). Therefore, while both

the amino acid composition and the molecular dynamics simulations support the possibility that both CDR3 α and CDR3 β loops are flexible, it is not possible to distinguish CDR3 α from CDR3 β or suppose that one is more or less rigid than the other.

To further explore whether the emergence of multiple conformations drives similar structural diversity despite reduced sequence diversity of the CDR3 α relative to the CDR3 β , we investigated how deterministic TRA or TRB VJ gene pairing is of CDR3 α or CDR3 β loop structure, respectively (Fig. 4h). For all TRAV-TRAJ and TRBV-TRBJ combinations for which at least five structures were observed in the TCRBuilder2+ training data, we

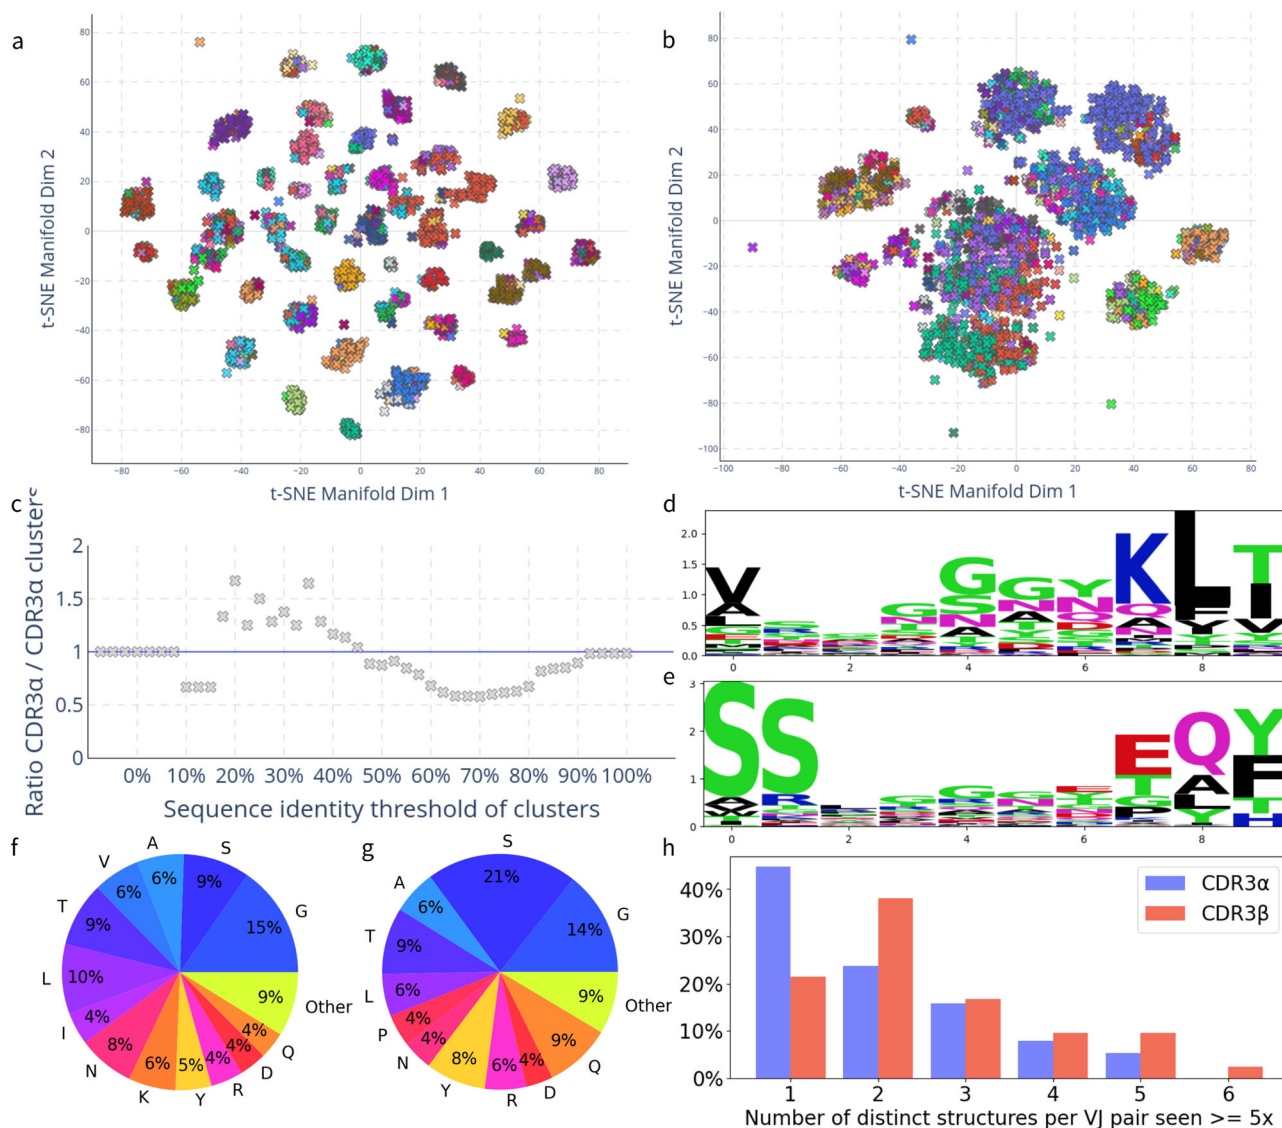


Fig. 4 | CDR3 α and CDR3 β sequence diversity exhibit different length scales.

a, b Sequence identity t-SNE distributions of 2000 OTS (**a**) CDR3 α and (**b**) CDR3 β loops coloured by 50% sequence identity cluster. CDR3 α sequences separate into fewer and more distinct clusters. **c** Ratio of CDR3 α sequence clusters to CDR3 β sequence clusters varied by sequence identity threshold. At high sequence identities more CDR3 β clusters emerge, whereas at lower sequence identities more CDR3 α clusters emerge. **d, e** Sequence logos of 2000 OTS samples' (**d**) CDR3 α and (**e**) CDR3 β loops of length ten. Central regions of the CDR3 α and CDR3 β sequence have

high entropy, and reveal high proportions of glycine. **f, g** Distribution of amino acids across 2000 OTS (**f**) CDR3 α and (**g**) CDR3 β samples. Blue-tinted segments indicate relatively small amino acids which may enable loop flexibility. **h** Histogram of the number of 2 Å structural clusters for variable (V) and joining (J) gene pairs (TRAV & TRAJ for CDR3 α and TRBV & TRBJ for CDR3 β) in training data structures for which at least five structures exist per VJ pair. VJ gene pairings map onto unique CDR3 α structures in 44% of cases and onto unique CDR3 β structures in 22% of cases.

calculated how many distinct CDR3 α and CDR3 β structures existed, respectively, using a 2 Å RMSD threshold. We observed that 44% of TRAV-TRAJ pairs map onto a single CDR3 α structure, while only 22% of TRBV-TRBJ pairs map onto a single CDR3 β structure. This suggests, in agreement with the expected effect of the insertion of the diversity gene and the additional junction in the CDR3 β loop, that the VJ gene pair is less deterministic of CDR3 β than CDR3 α structure. The high structural coherence of TRAV-TRAJ combination also suggests that CDR3 α structures are relatively well-defined for a given gene pair.

However, the better clustering of the CDR3 α sequences relative to the CDR3 β and high coherence of CDR3 α structure to TRAV-TRAJ gene pair raises the question where the equivalent CDR3 α and CDR3 β structural diversity we have observed originates from. Following our observation that the ratio between the number of CDR3 α and CDR3 β clusters is dependent on the sequence identity threshold, we investigated the combinatorial diversity of the CDR3 α and CDR3 β that arises from their genetic

recombination mechanisms. As is well established³, we found that the number of potential TRAV-TRAJ pairings outnumbers those of TRBV-TRBJ (SI Tables 7, 8). Even when accounting for the diversity gene, the potential combinatorial diversity of the beta chain is only about a third of that of the alpha chain, primarily due to the relatively low number of TRBJ genes compared to TRAJ^{3,68}. In conjunction with our observation that, when clustering CDR3 α and CDR3 β sequences permissively (20–40% sequence identity), more sequence clusters emerge for CDR3 α than CDR3 β (Fig. 4c, SI Table 9), this suggests that CDR3 α sequences may occupy tighter but more distinct regions in sequence space, as their primary diversification occurs at the level of genetic recombination. On the other hand, CDR3 β sequences appear to exhibit greater variability at the level of individual amino acids, but are constrained by a smaller pool of genetic recombinations. It is therefore possible that CDR3 α structural diversity is driven by the larger combinatorial pairing space provided by the larger number of functional TRAJ genes (51) relative to functional TRBJ genes (13).

Discussion

Incorporating structural information into TCR specificity prediction methods offers a compelling strategy towards alleviating their current limitations^{22,23} by encouraging models to learn more generalisable physicochemical principles of complementarity *in lieu* of amino acid sequence motifs^{25,26}. Given the current sparsity of solved TCR structure data³⁰ and the expense of generating it, *in-silico* structure prediction is a helpful tool to provide estimates of the 3D properties of TCRs.

In this work, we have retrained a state-of-the-art, high-throughput deep learning architecture for TCR structure prediction with the largest training set to date (TCRBuilder2+). We have made the resulting TCRBuilder2+ model parameters freely available (see Data Availability).

Evaluation of our model alongside AlphaFold Multimer and AlphaFold3 robustly characterises the current state-of-the-art of TCR structure prediction – all three models achieved median accuracies sub-2 Å RMSD across all regions, although the CDR3 α and CDR3 β loops remain discernibly harder to predict than the other loops. While performance was similar across the whole test set, on average TCRBuilder2+ showed CDR3 predictive performance improvements for genes for which the number of examples had increased in our new training set (Fig. 2b–c, SI Fig. 11). This aligns with related research exploring the generalisability of antibody structure prediction methods⁵⁹ as well as with more general work applying machine learning to biological data, which has shown that the quality, quantity, and sampling of the training data is critical for model performance^{49,69,70}. Furthermore, despite the large increase in non-redundant training data, we did not observe a statistically significant improvement in prediction accuracy across the whole test set compared to TCRBuilder2.

Supportive of the hypothesis that accurate TCR structure predictions may unlock additional information for downstream tasks such as specificity prediction²⁵, we found that the performance of docking simulations was linked to structure prediction accuracy (SI Note 1). We observed that a third of top-ranked TCR structure predictions exhibit sufficiently accurate CDR3 α and CDR3 β loops to enable docking within 4 Å interface RMSD (Fig. 2f), implying that, while an appreciable proportion of TCR structure predictions could already be useful for physics-based docking, there is still scope for improvements in prediction accuracy to benefit downstream tasks. Furthermore, disparities between poses with the best docking scores and the true best docks in our case studies suggest bespoke TCR:pMHC scoring functions may also be beneficial.

In agreement with prior findings, our simulation results suggest that the CDRs of TCRs may be flexible and may exhibit a degree of induced fit upon binding to the pMHC^{64–67}. This would imply that next-generation structure predictors would benefit from information about a partner antigen to generate more complementary TCR structures, and that a hypothetical, perfect structure predictor would require information about the antigen to predict an exact single prediction. Moreover, current structure prediction models are predominantly trained on TCRs in complex with antigen, since these are considerably more abundant in the existing data³⁰. If the antigen significantly influences the eventual bound structure of the TCR, then it is likely that current models will struggle to accurately predict apo TCR structures. While neither AlphaFold Multimer nor TCRBuilder2+ were trained to predict multiple conformations and evaluating flexibility of the predictions is challenging, both are comprised of ensembles, and as such multiple distinct predictions for a single TCR sequence can be generated (SI Fig. 20), at the expense of computational resource.

The inference speed of TCRBuilder2+ (approximately 250x faster than AlphaFold Multimer in our experiments (SI Table 10)) enabled the prediction of over 1.5 million sequence non-redundant TCR structures from repertoire studies, the first dataset of its kind⁷¹. Structurally characterising repertoire data has been shown to yield clinically relevant insights for BCRs^{72–74} and has enabled the training of deep learning architectures on structural data at previously inaccessible scales^{75,76}. Our release of the OTS structure predictions has the potential to enable similar advances in computational TCR research.

We found that both TCR-specific and general protein structure predictors modelled the VJ-recombined CDR3 α loop with similar accuracy to

the VDJ-recombined CDR3 β loop. Our analysis of the landscape of solved TCR structures also found that the VJ-recombined junction in the alpha chain appears to be as structurally diverse as the VDJ-recombined junction of the beta chain, and suggests that CDR3 α loop structures do not readily cluster into canonical forms. Prior works that assigned CDR3 α loops to canonical structures presented evidence supportive of this finding, showing that CDR3 α is not assigned to a cluster by sequence as consistently as CDR1 α & 2 α ⁶², and that structure inference based on class assignment to a canonical loop yields less accurate results for CDR3 α than CDR3 β ³⁴. For example, our 2019 paper⁶² shows that while 76.8% (43/(43+13)) of CDR1 α loop structures and 72.0% (36/(36+14)) of CDR2 α loop structures were assigned to a canonical cluster, only 22.4% (24/(24 + 83)) of CDR3 α loops could be assigned. This contrasts strongly with the trends seen across B cell receptors and antibodies: despite analogous CDR3 diversification mechanisms, CDR3L loops predominantly fold into canonical structures, while CDR3H structures are substantially more structurally diverse^{36,47,62}.

One potential functional rationale for antibodies and TCRs to have evolved these different levels of structural diversity in their VJ junction is that, in antibodies, the CDR3L tends to play a supportive role in binding, with the CDR3H driving antigen specificity and binding⁴⁵, whereas current evidence suggests that the TCR CDR3 α and CDR3 β play a more equal role in pMHC binding. For antibodies this has been substantiated through recent observations that the heavy chain clonotype is predictive of binding specificity but not the light chain clonotype⁴⁴ and metadynamics studies suggesting that CDR3H loops sample conformations in the space left unoccupied by the CDR3L loop and not vice versa⁷⁷.

While elucidating the precise rules of TCR to pMHC binding remains an active topic of research, structural analyses have shown that the CDR3 α loop and the CDR1 α loop, which is also difficult to predict despite being germline encoded (Table 1), frequently make direct peptide contact in TCR-pMHC interfaces^{30,78,79}, suggesting that additional information about the antigen or binding mode may be useful for attaining accurate structure predictions. Furthermore, unlike CDR3L loops relative to CDR3H loops, the CDR3 α samples a similar length distribution to CDR3 β ⁷⁸ (SI Fig. 15). Given the frequent contact between the peptide and both the CDR3 α and CDR3 β , it would be coherent for evolutionary pressure to have increased CDR3 α structural diversity alongside CDR3 β diversity to achieve better peptide sensitivity.

Supportive of this hypothesis, the number of human TRA and TRB locus genes indicate that the potential combinatorial diversity of the alpha chain is greater than that of the beta chain, compensating somewhat for the presence of only one junction region rather than two (SI Tables 7, 8)⁶⁸. This appears to be driven primarily by the greater number of TRAJ genes (61) relative to both TRBJ (14) and IGLJ/KJ (11/5). Alpha chain combinatorial and structural diversity substantially exceeds that of the antibody light chain, which remains less structurally diverse despite undergoing additional somatic hypermutation diversification. Overall, this outsized diversity in the alpha chain VJ junction, coupled with relative data paucity, likely contributes substantially towards the CDR3 α prediction errors being on-par with those of the CDR3 β across all models tested. However, the high structural coherence of TRAJ-TRAV gene pairs and observed predictive improvements over better sampled genes suggests that supplementing existing TCR structure data with as yet unsolved crystal structures could yield substantial improvements in prediction accuracy.

TCRBuilder2+ allows TCR structures to be predicted at high-throughput with sub-2 Å median RMSD across all regions, but near-term improvements on these values may require fundamental methodological changes that better reflect TCR biology. Our analysis highlights the structural variability of the TCR CDR3 α , on par with that of the CDR3 β , differentiating T cells from their B cell/antibody cousins and reinforcing the value of paired-chain sequencing data and incorporating both chains in TCR:pMHC specificity models. Finally, our high-throughput model, TCRBuilder2+, trained on a supplemented dataset, enables the integration of TCR structure predictions into *in-silico* TCR screening and design pipelines, and, in combination with our release of ca. 1.5 million predicted

Table 2 | Resolution statistics of training data

	Mean ± Std	Median	Range
STCRDab	2.56 ± 0.47	2.60	(1.26, 3.48)
Immunocore	2.43 ± 0.51	2.36	(1.27, 3.81)

structures of native TCRs, will help unlock the field of repertoire structural immunoinformatics.

Methods

Data curation

The original training data of TCRBuilder2 contained 704 experimentally resolved TCR structures sourced from STCRDab^{30,36}. These structures were obtained from 463 distinct PDB files, and 349 of the 704 structures are unique in sequence. We extended this training set with 204 Immunocore TCR structures, of which 201 sequences are unique and only 6 overlap with the existing training set.

Further filtering for alpha/beta TCRs yielded a total of 881 structures, of which 524 are unique in sequence. As a quality control we applied a resolution threshold of 3.5 Å to the TCRs from STCRDab during training. We did not apply this threshold to the supplemented Immunocore data as all structures had a resolution of < 3.81 Å (Table 2). Of the STCRDab training TCRs, 38.8 % are apo; 4.4% of Immunocore training TCRs are apo.

TCRBuilder2 was evaluated on a relatively small test set of 21 structures⁵². Since the original work was published more TCR structures have become available, allowing for the curation of a larger and more robust test set. We updated the database of TCR structures, STCRDab³⁰, and generated a new, independent test set with the following criteria: TCRs must be paired alpha-beta chains, have lower than 3.5 Å resolution, be unique within the test set and be non-identical to a paired sequence in the training set. This process yielded 45 TCR structures, the PDB codes of which we have provided in the Supplementary Information (SI Table 1).

Gene annotation and numbering

We used ANARCI⁸⁰ with the IMGT gene set^{68,81} retrieved 2023/09/04 to annotate and number TCR sequences.

TCRBuilder2+ model and training

We retrained the deep learning model TCRBuilder2³⁶ on our new dataset to create TCRBuilder2+. Like the original model, the neural network architecture of TCRBuilder2+ is based on the 'structure module' of Alphafold³⁵. The structure module is an SE(3)-invariant neural network architecture through which we pass the one-hot encoding of the TCR sequence as input. This encoding is then passed through eight sequential layers, each consisting of an invariant point attention module and a backbone update, which are used to update the translations and rotations of amino-acids represented as unconstrained planar triangles in Euclidian space. The model is trained *via* gradient descent with the RAdam optimiser⁸² using a variation of the Frame-Aligned Point Error (FAPE loss). We used early stopping on a withheld validation set to determine the point we consider the parameters to be sufficiently optimised. For further details please see the ImmuneBuilder methods section³⁶. However, because TCRBuilder2 followed from ABodyBuilder2, in the original work the model ensemble was selected to minimise the RMSD over the beta chain, whose genetic recombination is analogous to the heavy chain. We adapted the ensemble selection in TCRBuilder2+ to account for the structural diversity of the alpha chain by selecting the ensemble of models that minimises the RMSD over both the alpha and beta chain on the validation set. Specifically, ten models were independently trained, each with a partially randomised validation set. The validation sets each contained 30 TCRs which were held constant across all ten splits as well as 20 randomly selected TCRs. After training all ten models to convergence, we evaluated them on the constant subset of the validation set to select the four models with the lowest total RMSD over both the alpha and the beta chain. During inference we predicted four structures and selected the model

Table 3 | Definitions for anchors and loops based on IMGT numbering

Region	Anchor residue numbers	Loop residue numbers
CDR1	22-27, 39-44	28-38
CDR2	51-56, 66-71	57-65
CDR3	100-105, 118-123	106-117

Structures were aligned by anchors and RMSD was then calculated over residues in the loops.

closest to the ensembled average. The selected structure was then refined using the OpenMM implementation of the AMBER14 protein force field and returned as the final prediction^{83,84}. To produce an uncertainty estimate, we also calculated the sum of the squared distance of each model's prediction x_i to the mean prediction μ : $\epsilon = \sum_{i=0}^4 (x_i - \mu)^2$. This provides an uncertainty proxy for the coordinates of every amino acid. To produce multiple conformations for both the molecular dynamics and docking simulations we returned all four models' predictions, and minimised each of them with the AMBER14 force field.

TCR structure predictions

ImmuneBuilder. TCRBuilder2+, TCRBuilder2, and ABodyBuilder2 are all installable via the ImmuneBuilder package. We used the previously published versions of the weights for TCRBuilder2 and ABodyBuilder2, and the weights we trained as described above for TCRBuilder2+.

Alphafold Multimer. We used an instance of Alphafold Multimer⁵⁷ installed on our servers, and used version 2.3.1 of the weights. Due to constraints on the memory of the installation, we used the reduced version of the BFD database, as provided accompanying the Alphafold weights. We set a cut-off date of 2020/12/30 for templates, as this excludes all the test structures. However, it is possible that some of the test structures were included in the training data for Alphafold Multimer version 2.3.1, potentially giving the model a slight advantage.

Alphafold3. We used the Alphafold Server (alphafoldserver.com) accessed on September 10th, 11th and 12th 2024 (due to the 20 structure per day limit) to predict the TCR structures in our test set. It is possible that train test overlap is inflating the performance of Alphafold3 due to a training cut-off of 2023/01/12, and a template cut-off of 2022/09/28, which does not exclude our entire test set⁵⁸.

LYRA & TCRmodel. We installed LYRA³⁴ as part of the TCRpMHCmodel package⁸⁵ and ran inference withholding the target structure from the templates. We installed TCRmodel³³ with ROSETTA⁸⁶ and made predictions using the default settings.

Sequence numbering and gene annotation

We used IMGT numbering and gene annotations computed on amino acids by ANARCI^{80,81}. Table 3 shows the definitions of loops and anchors used. We refer to the CDR3 loops of the alpha and beta chain as CDR3 α and CDR3 β respectively.

Calculating structural deviation

We used the root-mean-squared-distance (RMSD) as the evaluation metric for the predictions, which was calculated over the amino acid backbone atoms (N, C ω , C, and O). We aligned each CDR by the backbone atoms of the anchor residues and calculated RMSD over the backbone atoms of residues in the loops. The loop and anchor residue numbers are provided in Table 3, and were annotated using ANARCI⁸⁰ according to the IMGT numbering scheme⁸¹.

To compare CDR3 loop structures of different lengths, we implemented spline up-sampling, the parameters of which we set such that the distance calculated is RMSD equivalent when the loops are of the same length. Further details can be found in the Supplementary Information (SI Note 3).

An RMSD threshold for an acceptable TCR structure prediction was based on our observations of the sequence to structure identity distributions of CDR3 α and CDR3 β (Fig. 1d & e). Since RMSD will naturally be greater for longer CDR3 loops due to the increased degrees of freedom in the backbone, we also investigated the background distribution within each loop length (SI Fig. 6). This confirmed 2.0 Å as a sensible threshold across the range of TCR CDR3 lengths.

Statistical significance of RMSD distributions

We used the p -values of a two-tailed statistical test on the t -values to accept or reject the hypothesis that the distributions of RMSD of different models are different from one another. The null-hypothesis is therefore that the RMSD distribution of models is the same. We used the commonly defined t -score: $t = |\mu_1 - \mu_2| / \sqrt{(\sigma_1^2 + \sigma_2^2) / N}$, and used the survival function as implemented in *scipy*, with $N - 1$ degrees of freedom, to calculate the p -value, which we doubled, accounting for the test being two-tailed: $p = 2 * \text{survival_function}(t, N - 1)$. $N = 45$, the number of structures in the test set.

Calculating LDDT scores

We calculated the local-difference-distance test (LDDT) metric⁸⁷ using the implementation provided within the ‘compare-structures’ interface of the OpenStructure software package⁸⁸.

Sequence identity

We calculated the pairwise sequence identity between sequences with the following formula:

$$\text{alignment}(seq_i, seq_j) / \max(\text{len}(seq_i), \text{len}(seq_j))$$

We used *alignment()* as implemented in *Biopython* as *pairwise2.align.globalxx*⁸⁹. The function returns the sum of the alignment of two sequences where the alignment contains 1 for every matched pair of amino acids and zero for every mismatched amino acid pair. No gap or insertion penalties were used.

Clustering

To cluster CDR loops by structure and sequence we implemented a greedy clustering algorithm which we applied to pairwise distance matrices (Alg. 1). For structures the distance matrices contained RMSD between structures, and upsampled RMSD in cases where loop lengths were different. For sequences we calculated a sequence distance matrix from the sequence identity matrix as $d_{ij} = 1 - i_{ij}$.

Algorithm 1. Greedy clustering

```

1: function ClusterD, N,  $\tau$ 
2:   Let clusters  $\leftarrow$  {}
3:   Let clustered_indices  $\leftarrow$  []
4:   Let count  $\leftarrow$  0
5:   for  $i = 1$  to  $N$  do
6:     if  $i \notin$  clustered_indices then
7:       clusters[count]  $\leftarrow$  [i]
8:       append  $i$  to clustered_indices
9:       for  $j = i + 1$  to  $N$  do
10:        if  $j \notin$  clustered_indices and  $D_{ij} < \tau$  then
11:          append  $D_{ij}$  to clusters[count]
12:          append  $j$  to clustered_indices
13:        end if
14:      end for
15:      count  $\leftarrow$  count + 1
16:    end if
17:  end for
18:  return clusters
19: end function

```

To visualise clusters we used the t-SNE dimension reduction algorithm⁹⁰ as implemented in *scikit-learn*⁹¹, with a perplexity parameter of 10. We used the t-SNE plots for visualisation purposes only and used the clustering algorithm as outlined above to colour the clusters.

Statistics and reproducibility

All TCR structure predictions were evaluated using a publicly available test set containing 45 TCRs deposited in the PDB (SI Table 1). We report the sample mean of the RMSD of the aligned structures against the crystal structures.

The box plots reported in Fig. 2, which visualise the RMSD evaluation, are defined as follows: boxes represent the inner quartiles of the data, with the median value the dividing line; whiskers extend to points that lie within 1.5 times the interquartile range of the lower and upper quartile; observations beyond this range are displayed independently. $N = 45$, the number of structures in the test set.

Data availability

TCRBuilder2+ weights: <https://doi.org/10.5281/zenodo.10892159>⁹². OTS TCR structure predictions made with TCRBuilder2+: <https://doi.org/10.5281/zenodo.10854757>⁷¹. Data used to generate manuscript figures is provided in the Supplementary Data. All other non-proprietary data is readily available upon request from the corresponding author. Proprietary data pertains only to the crystal structures solved by Immunocore used for model training, all other crystal structures used for training are available from STCRDab³⁰.

Code availability

TCRBuilder2+ github (integrated into ImmuneBuilder)⁹³: <https://github.com/oxpig/ImmuneBuilder>. Version release with DOI accompanying publication: <https://doi.org/10.5281/zenodo.14727165>.

Received: 3 June 2024; Accepted: 9 February 2025;

Published online: 04 March 2025

References

- van der Merwe, P. A. & Dushek, O. Mechanisms for T cell receptor triggering. *Nat. Rev. Immunol.* **11**, 47–55 (2011).
- Birnbaum, M. E. et al. Deconstructing the peptide-MHC specificity of T cell recognition. *Cell* **157**, 1073–1087 (2014).
- Murphy, K. & Casey, W. *Janeway's Immunobiology* 9th edn (2017).
- Cai, M., Bang, S., Zhang, P. & Lee, H. ATM-TCR: TCR-epitope binding affinity prediction using a multi-head self-attention model. *Front. Immunol.* **13**, 893247 (2022).
- Montemurro, A. et al. NetTCR-2.0 enables accurate prediction of TCR-peptide binding by using paired TCR α and β sequence data. *Commun. Biol.* **4**, 1060 (2021).
- Ehrlich, R. et al. SwarmTCR: a computational approach to predict the specificity of T cell receptors. *BMC Bioinform.* **22**, 422 (2021).
- Moris, P. et al. Current challenges for unseen-epitope TCR interaction prediction and a new perspective derived from image classification. *Brief. Bioinform.* **22**, bbaa318 (2021).
- Luu, A., Leistico, J., Miller, T., Kim, S. & Song, J. Predicting TCR-Epitope binding specificity using deep metric learning and multimodal learning. *Genes* **12**, 572 (2021).
- Jokinen, E., Huuhtanen, J., Mustjoki, S., Heinonen, M. & Lähdesmäki, H. Predicting recognition between T cell receptors and epitopes with TCRGP. *PLOS Comput. Biol.* **17**, e1008814 (2021).
- Fischer, D. S., Wu, Y., Schubert, B. & Theis, F. J. Predicting antigen specificity of single T cells based on TCR CDR 3 regions. *Mol. Syst. Biol.* **16**, e9416 (2020).
- Tong, Y. et al. SETE: Sequence-based ensemble learning approach for TCR Epitope binding prediction. *Comput. Biol. Chem.* **87**, 107281 (2020).
- Lu, T. et al. Deep learning-based prediction of the T cell receptor-antigen binding specificity. *Nat. Mach. Intell.* **3**, 864–875 (2021).

13. Weber, A., Born, J. & Rodriguez Martínez, M. TITAN: T-cell receptor specificity prediction with bimodal attention networks. *Bioinformatics* **37**, i237–i244 (2021).
14. Gao, Y. et al. Pan-peptide meta learning for T-cell receptor–antigen binding recognition. *Nat. Mach. Intell.* **5**, 236–249 (2023).
15. Gielis, S. et al. Detection of enriched T cell epitope specificity in full T cell receptor sequence repertoires. *Front. Immunol.* **10**, 2820 (2019).
16. Chen, J. et al. TEPCAM: prediction of T-cell receptor–epitope binding specificity via interpretable deep learning. *Protein Sci.* **33**, e4841 (2024).
17. Banerjee, A. et al. BATMAN: Improved T cell receptor cross-reactivity prediction benchmarked on a comprehensive mutational scan database <https://doi.org/10.1101/2024.01.22.576714> (2024).
18. Zahid, H. J. et al. Large-scale statistical mapping of T-cell receptor β sequences to Human Leukocyte Antigens <https://doi.org/10.1101/2024.04.01.587617> (2024).
19. Glanville, J. et al. Identifying specificity groups in the T cell receptor repertoire. *Nature* **547**, 94–98 (2017).
20. Wang, A. et al. RACER-m leverages structural features for sparse T cell specificity prediction. *Sci. Adv.* **10**, 161 (2024).
21. Schaap-Johansen, A. L., Vujović, M., Borch, A., Hadrup, S. R. & Marcatili, P. T. Cell epitope prediction and its application to immunotherapy. *Front. Immunol.* **12**, 712488 (2021).
22. Hudson, D., Fernandes, R. A., Basham, M., Ogg, G. & Koohy, H. Can we predict T cell specificity with digital biology and machine learning? *Nat. Rev. Immunol.* **23**, 511–521 (2023).
23. Meysman, P. et al. Benchmarking solutions to the T-cell receptor epitope prediction problem: IMMREP22 workshop report. *bioRxiv* <https://doi.org/10.1101/2022.10.27.514020v1> (2022).
24. Cancer Grand Challenges. T-cell receptors: Decipher the T-cell receptor cancer- recognition code <https://cancergrandchallenges.org/challenges/t-cell-receptors> (2023).
25. Ghoreyshi, Z. S. & George, J. T. Quantitative approaches for decoding the specificity of the human T cell repertoire. *Front. Immunol.* **14**, 1228873 (2023).
26. Zhao, Y. et al. DeepAIR: A deep learning framework for effective integration of sequence and 3D structure to enable adaptive immune receptor analysis. *Sci. Adv.* **9**, eabo5128 (2023).
27. McMaster, B., Thorpe, C., Ogg, G., Deane, C. M. & Koohy, H. Can AlphaFold's breakthrough in protein structure help decode the fundamental principles of adaptive cellular immunity? *Nat. Methods* <https://www.nature.com/articles/s41592-024-02240-7> (2024).
28. Bradley, P. Structure-based prediction of T cell receptor:peptide-MHC interactions. *eLife* **12**, 2022.08.05.503004 (2023).
29. Berman, H. M. et al. The Protein Data Bank. *Techn. Rep. 1* <http://www.rcsb.org/pdb/status.html> (2000).
30. Leem, J., De Oliveira, S. H., Krawczyk, K. & Deane, C. M. STCRDab: the structural T-cell receptor database. *Nucleic Acids Res.* **46**, D406–D412 (2018).
31. Shugay, M. et al. VDjdb: a curated database of T-cell receptor sequences with known antigen specificity. *Nucleic Acids Res.* **46**, D419–D427 (2018).
32. Schmitt, D. et al. Repertoire Builder: High-throughput structural modeling of B and T cell receptors. *Mol. Syst. Des. Eng.* **4**, 761–768 (2019).
33. Gowthaman, R. & Pierce, B. G. TCRmodel: high resolution modeling of T cell receptors from sequence. *Nucleic Acids Res.* **46**, W396–W401 (2018).
34. Klausen, M. S., Anderson, M. V., Jespersen, M. C., Nielsen, M. & Marcatili, P. LYRA, a webserver for lymphocyte receptor structural modeling. *Nucleic Acids Res.* **43**, W349 (2015).
35. Jumper, J. et al. Highly accurate protein structure prediction with AlphaFold. *Nature* **596**, 583–589 (2021).
36. Abanades, B. et al. ImmuneBuilder: Deep-Learning models for predicting the structures of immune proteins. *Commun. Biol.* **6**, 575 (2023).
37. Hanser, T. Federated learning for molecular discovery. *Curr. Opin. Struct. Biol.* **79**, 102545 (2023).
38. Brinkhaus, H. O., Rajan, K., Schaub, J., Zielesny, A. & Steinbeck, C. Open data and algorithms for open science in AI-driven molecular informatics. *Curr. Opin. Struct. Biol.* **79**, 102542 (2023).
39. Zhang, H., Zhan, X. & Li, B. GIANA allows computationally-efficient TCR clustering and multi-disease repertoire classification by isometric transformation. *Nat. Commun.* **12**, 4699 (2021).
40. Chronister, W. D. et al. TCRMatch: predicting T-cell receptor specificity based on sequence similarity to previously characterized receptors. *Front. Immunol.* **12**, 640725 (2021).
41. Gauss, G. & Lieber, M. R. Mechanistic constraints on diversity in human V(D)J recombination. *Mol. Cell. Biol.* **16**, 258–269 (1996).
42. Hou, X. et al. Analysis of the repertoire features of TCR beta chain CDR3 in human by high-throughput sequencing. *Cell. Physiol. Biochem.* **39**, 651–667 (2016).
43. Cabaniols, J. P., Fazilleau, N., Casrouge, A., Kourilsky, P. & Kanellopoulos, J. M. Most $\alpha\beta$ T cell receptor diversity is due to terminal deoxynucleotidyl transferase. *J. Exp. Med.* **194**, 1385–1390 (2001).
44. Jaffe, D. B. et al. Functional antibodies exhibit light chain coherence. *Nature* **611**, 352–357 (2022).
45. Regep, C., Georges, G., Shi, J., Popovic, B. & Deane, C. M. The H3 loop of antibodies shows unique structural characteristics. *Proteins: Struct., Funct., Bioinform.* **85**, 1311–1318 (2017).
46. Mayer, A. & Callan, C. G. Measures of epitope binding degeneracy from T cell receptor repertoires. *Proc. Natl Acad. Sci. USA* **120**, e2213264120 (2023).
47. Ruffolo, J. A., Chu, L. S., Mahajan, S. P. & Gray, J. J. Fast, accurate antibody structure prediction from deep learning on massive set of natural antibodies. *Nat. Commun.* **14**, 2389 (2023).
48. Raybould, M. I. et al. The Observed T cell receptor Space database enables paired-chain repertoire mining, coherence analysis and language modelling. *BioRxiv* <https://doi.org/10.1101/2024.05.20.594960> (2024).
49. Durant, G., Boyles, F., Birchall, K., Marsden, B. & Deane, C. M. Robustly interrogating machine learning-based scoring functions: what are they learning? *Bioinformatics* **2023**, 1–7 (2023).
50. Klarner, L. et al. Drug discovery under covariate shift with domain-informed prior distributions over functions. In: *Proc. 40th International Conference on Machine Learning* (2023) <https://github.com/leojklarner/Q-SAVI>.
51. Nauta, M., Walsh, R., Dubowski, A. & Seifert, C. Uncovering and correcting shortcut learning in machine learning models for skin cancer diagnosis. *Diagnostics* **12**, 40 (2021).
52. Abanades, B. et al. Supplementary Information ImmuneBuilder: Deep-Learning models for predicting the structures of immune proteins. *Commun. Biol.* **6**, 575 (2023).
53. Chothia, C. & Lesk, A. M. Canonical structures for the hypervariable regions of immunoglobulins. *J. Mol. Biol.* **196** (1987).
54. Wong, W. K. et al. SCALOP: Sequence-based antibody canonical loop structure annotation. *Bioinformatics* **35**, 1774–1776 (2019).
55. North, B., Lehmann, A. & Dunbrack, R. L. A new clustering of antibody CDR loop conformations. *J. Mol. Biol.* **406**, 228–256 (2011).
56. Nowak, J. et al. Length-independent structural similarities enrich the antibody CDR canonical class model. *mAbs* **8** (2016).
57. Evans, R. et al. Protein complex prediction with AlphaFold-Multimer. *bioRxiv* <https://doi.org/10.1101/2021.10.04.463034v2> (2022).
58. Mikhaylov, V. & Levine, A. J. Accurate modeling of peptide-MHC structures with AlphaFold. <https://doi.org/10.1101/2023.03.06.531396>.
59. Greenshields-Watson, A., Abanades, B. & Deane, C. M. Investigating the ability of deep learning-based structure prediction to extrapolate and/or enrich the set of antibody CDR canonical forms. *Front. Immunol.* **15**, 1352703 (2024).

60. Van Zundert, G. C. et al. The HADDOCK2.2 web server: user-friendly integrative modeling of biomolecular complexes. *J. Mol. Biol.* **428**, 720–725 (2016).
61. Olsen, T. H., Boyles, F. & Deane, C. M. Observed antibody space: a diverse database of cleaned, annotated, and translated unpaired and paired antibody sequences. *Protein Sci.* **31**, 141–146 (2022).
62. Wong, W. K., Leem, J. & Leem, J. Comparative analysis of the CDR loops of antigen receptors. *Front. Immunol.* **10**, 2454 (2019).
63. Al-Lazikani, B., Lesk, A. M. & Chothia, C. Canonical structures for the hypervariable regions of T cell $\alpha\beta$ receptors. *J. Mol. Biol.* **295**, 979–995 (2000).
64. Baker, B. M., Scott, D. R., Blevins, S. J. & Hawse, W. F. Structural and dynamic control of T-cell receptor specificity, cross-reactivity, and binding mechanism. *Immunol. Rev.* **250** (2012).
65. Gras, S. et al. A structural voyage toward an understanding of the MHC-I-restricted immune response: Lessons learned and much to be learned. *Immunol. Rev.* **250** (2012).
66. Holland, C. J. et al. In silico and structural analyses demonstrate that intrinsic protein motions guide t cell receptor complementarity determining region loop flexibility. *Front. Immunology* **9** <https://doi.org/10.3389/fimmu.2018.00674/full> (2018).
67. Crooks, J. E., Boughter, C. T., Scott, L. R. & Adams, E. J. The hypervariable loops of free TCRs sample multiple distinct metastable conformations in solution. *Front. Mol. Biosci.* **5** <https://doi.org/10.3389/fmolb.2018.00095/full> (2018).
68. Giudicelli, V. IMGT/GENE-DB: a comprehensive database for human and mouse immunoglobulin and T cell receptor genes. *Nucleic Acids Res.* **33**, D256–D261 (2004).
69. Hummer, A. M., Schneider, C., Chinery, L. & Deane, C. M. Investigating the volume and diversity of data needed for generalizable antibody-antigen $\Delta\Delta G$ prediction. *BioRxiv* <https://doi.org/10.1101/2023.05.17.541222> (2023).
70. Klamer, L., Reutlinger, M., Schindler, T., Deane, C. & Morris, G. Bias in the Benchmark: Systematic experimental errors in bioactivity databases confound multi-task and meta-learning algorithms. In *2nd ICML AI for Science Workshop* (OpenReview.net, 2022).
71. Quast, N. P. et al. TCRBuilder2+ predictions of the Observed T-cell receptor Space (OTS). zenodo <https://doi.org/10.5281/zenodo.10854757> (2024).
72. Raybould, M. I. et al. Public Baseline and shared response structures support the theory of antibody repertoire functional commonality. *PLoS Comput. Biol.* **17**, e1008781 (2021).
73. Robinson, S. A. et al. Epitope profiling using computational structural modelling demonstrated on coronavirus-binding antibodies. *PLoS Comput. Biol.* **17** (2021).
74. Spöndlin, F. C. et al. Improved computational epitope profiling using structural models identifies a broader diversity of antibodies that bind to the same epitope. *Front. Mol. Biosci.* **10** (2023).
75. Dreyer, F. A., Cutting, D., Schneider, C., Kenlay, H. & Deane, C. M. Inverse folding for antibody sequence design using deep learning. In *2023 ICML Workshop on Computational Biology* (ICML, 2023).
76. Hummer, A. M., Haraldson Høie, M., Olsen, T. H., Nielsen, M. & Deane, C. M. AntiFold: Improved antibody structure design using inverse folding. *NeurIPS Machine Learning for Structural Biology Workshop* <https://opig.stats.ox.ac.uk/data/downloads/AntiFold> (2023).
77. Guloglu, B. & Deane, C. M. Specific attributes of the VL domain influence both the structure and structural variability of CDR-H3 through steric effects. *Front. Immunol.* **14** <https://doi.org/10.3389/fimmu.2023.1223802/full> (2023).
78. Raybould, M. I. J., Nissley, D. A., Kumar, S. & Deane, C. M. Computationally profiling peptide:MHC recognition by T-cell receptors and T-cell receptor-mimetic antibodies. *Front. Immunol.* **13**, 2022.07.15.500208 (2023).
79. McMaster, B., Thorpe, C., Rossjohn, J., Deane, C. M. & Koohy, H. Quantifying conformational changes in the TCR:pMHC-I binding interface. *Front. Immunol.* **15**, 1491656 (2024).
80. Dunbar, J. & Deane, C. M. ANARCI: antigen receptor numbering and receptor classification. *Bioinformatics* **32**, 298–300 (2016).
81. Lefranc, M. P. et al. IMGT, the international ImMunoGeneTics information system*. *Nucleic Acids Res.* **33** (2005).
82. Liu, L. et al. On the variance of the adaptive learning rate and beyond. In *8th International Conference on Learning Representations, ICLR 2020* (OpenReview.net, 2020).
83. Eastman, P. et al. OpenMM 7: Rapid development of high performance algorithms for molecular dynamics. *PLoS Comput. Biol.* **13** (2017).
84. Maier, J. A. et al. ff14SB: improving the accuracy of protein side chain and backbone parameters from ff99SB. *J. Chem. Theory Comput.* **11**, 3696–3713 (2015).
85. Jensen, K. K. et al. TCRpMHCmodels: structural modelling of TCR-pMHC class I complexes. *Sci. Rep.* **9**, 14530 (2019).
86. Leaver-Fay, A. et al. Rosetta3: An object-oriented software suite for the simulation and design of macromolecules. *Methods Enzymol.* **487** (2011).
87. Mariani, V., Biasini, M., Barbato, A. & Schwede, T. IDDT: A local superposition-free score for comparing protein structures and models using distance difference tests. *Bioinformatics* **29** (2013).
88. Biasini, M. et al. OpenStructure: An integrated software framework for computational structural biology. *Acta Crystallogr. D: Biol. Crystallogr.* **69** (2013).
89. Cock, P. J. A. et al. Biopython: freely available Python tools for computational molecular biology and bioinformatics. *Bioinformatics* **25**, 1422–1423 (2009).
90. Van Der Maaten, L. & Hinton, G. Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**, 2579–2605 (2008).
91. Pedregosa, F. et al. Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
92. Quast, N. P. et al. TCRBuilder2+ weights (v1.0). zenodo <https://doi.org/10.5281/zenodo.10892159> (2024).
93. Quast, N. et al. ImmuneBuilder version 1.2 (accompanying TCRBuilder2+ release) (v1). zenodo <https://doi.org/10.5281/zenodo.14727164> (2024).

Acknowledgements

N.P.Q.: Engineering and Physical Sciences Research Council (EPSRC) grant number EP/S024093/1 and Immunocore via the SABS:R³ doctoral training programme. B.A.: EPSRC grant number EP/S024093/1 and Roche via the SABS:R³ doctoral training programme. B.G.: Wellcome Trust grant number: 102164/Z/13/Z. M.I.J.R.: Postdoctoral Fellowship from Boehringer Ingelheim.

Author contributions

N.P.Q., M.I.J.R., S.H., and C.M.D. conceived the project and designed the study with input from all authors. N.P.Q. trained the final TCRBuilder2+ model, optimised the model version, conducted the analysis, and compiled the results. B.A. designed and implemented the original ImmuneBuilder model and provided training code for TCRBuilder2+. N.P.Q. compared TCRBuilder2+ against other methods. B.G. conducted and analysed the molecular dynamics study. N.P.Q. ran the docking simulations. V.K. and S.H. provided the additional TCR structure data from Immunocore and advised the project. N.P.Q., M.I.J.R., and C.M.D. wrote the manuscript with input from all authors. M.I.J.R., S.H., and C.M.D. supervised the project.

Competing interests

S.H. and V.K. are employees of Immunocore Ltd and may hold Immunocore stock options. C.D. discloses membership of the Scientific Advisory Board of Fusion Antibodies and AI proteins. All other authors declare no conflict of interest.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42003-025-07708-6>.

Correspondence and requests for materials should be addressed to Charlotte M. Deane.

Peer review information *Communications Biology* thanks Xingcheng Lin, Brian Pierce, and the other, anonymous, reviewer for their contribution to the peer review of this work. Primary Handling Editors: Laura Rodríguez Pérez and Ophelia Bu.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025