

---

# 000 001 QUANTIFYING CROSS-ATTENTION INTERACTION IN 002 TRANSFORMERS FOR INTERPRETING TCR-PMHC 003 BINDING 004 005

006 **Anonymous authors**

007 Paper under double-blind review  
008  
009  
010  
011

## ABSTRACT

013 CD8+ “killer” T cells and CD4+ “helper” T cells play a central role in the adaptive  
014 immune system by recognizing antigens presented by Major Histocompatibility Complex (pMHC)  
015 molecules via T Cell Receptors (TCRs). Modeling binding between T cells and the pMHC complex is fundamental to understanding  
016 basic mechanisms of human immune response as well as in developing therapies.  
017 While transformer-based models such as TULIP have achieved impressive  
018 performance in this domain, their black-box nature precludes interpretability and  
019 thus limits a deeper mechanistic understanding of T cell response. Most existing  
020 post-hoc explainable AI (xAI) methods are confined to encoder-only, co-attention,  
021 or model-specific architectures and cannot handle encoder-decoder transformers  
022 used in TCR-pMHC modeling. To address this gap, we propose Quantifying  
023 Cross-Attention Interaction (QCAI), a new post-hoc method designed to interpret  
024 the cross-attention mechanisms in transformer decoders. Quantitative evaluation  
025 is a challenge for XAI methods; we have compiled TCR-XAI, a benchmark  
026 consisting of 274 experimentally determined TCR-pMHC structures to serve as  
027 ground truth for binding. Using these structures we compute physical distances  
028 between relevant amino acid residues in the TCR-pMHC interaction region and  
029 evaluate how well our method and others estimate the importance of residues in  
030 this region across the dataset. We show that QCAI achieves state-of-the-art per-  
031 formance on both interpretability and prediction accuracy under the TCR-XAI  
032 benchmark.  
033

## 034 1 INTRODUCTION

035 T cells play a pivotal role in the adaptive immune system by identifying and responding to antigenic  
036 proteins, both from pathogens such as viruses, bacteria and cancer cells (Joglekar & Li, 2021) as  
037 well as in the context of autoimmunity. The final and arguably most critical component of T cell  
038 response is binding between the peptide Major Histocompatibility Complex (pMHC) which con-  
039 tains an antigenic peptide bound to a MHC molecule and the surface receptor on T cells (TCR).  
040 The specificity of this interaction underpins T cell-mediated immunity and is an intense area of re-  
041 search in both the development of therapies and fundamental understanding of immune response.  
042 Understanding T cell response is the key to vaccines that confer long-lasting immunity, and can also  
043 enable effective personalized cancer therapies (Rojas et al., 2023; Poorebrahim et al., 2021).  
044

045 Transformer models have recently been used to analyze T cell immunity (Hudson et al., 2023; Li  
046 et al., 2023; Karthikeyan et al., 2023; Driessens et al., 2024; Cornwall et al., 2023). Specifically,  
047 models have been developed to predict TCR-pMHC binding such as TULIP (Meynard-Piganeau  
048 et al., 2024), Cross-TCR-Interpreter (Koyama et al., 2023), TCR-BERT (Wu et al., 2024b), and  
049 BERTrand (Myronov et al., 2023)<sup>1</sup>. However these models are black boxes and suffer from a lack  
050 of interpretability, which is critically important in elucidating the mechanisms involved in T cell  
051 response. To address this challenge, post-hoc explanation techniques (Kenny et al., 2021) have been  
052 developed to connect elements of the input and model to the outputs. However, current existing  
053 post-hoc methods (e.g., AttnLRP (Achitbat et al., 2024), TokenTM (Wu et al., 2024a), and TEP-

---

<sup>1</sup>For a comprehensive discussion, please consult Section A.4 of the Appendix.

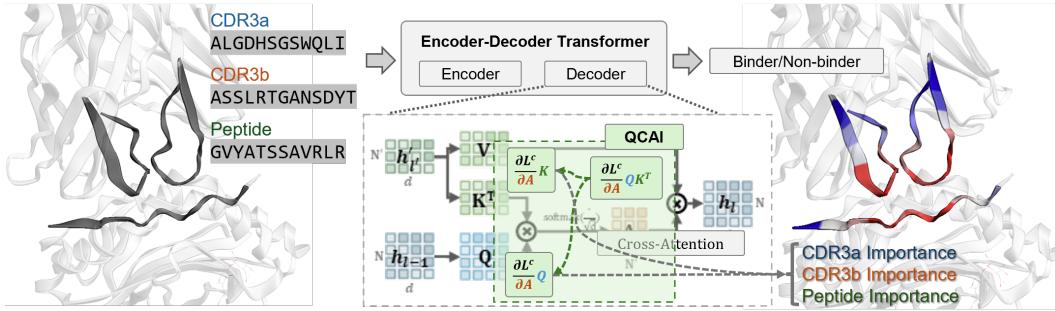


Figure 1: **Quantifying Cross-Attention Interaction (QCAI)** is a post-hoc explanation method designed for cross-attention mechanisms. In this paper, we show that QCAI enables insight into the structural basis for TCR-pMHC binding.

CAM (Chen et al., 2024)) are designed for encoder-only transformer or convolution neural network (CNN) models, while state-of-the-art TCR-pMHC binding predictors adopt encoder-decoder architectures.

The main contribution of this paper is to fill this gap with a novel post hoc explanation method that we call **Quantifying Cross-Attention Interaction (QCAI)** that enables interpretation of any encoder-decoder transformer model while taking cross-attention into account. Our motivation is the application of TCR-pMHC modeling, but cross-attention is used extensively in vision and NLP applications as well and thus QCAI is potential to be applied to other fields. Figure 1 shows how QCAI is used to analyze decoder blocks; cross-attention between the CDR3 and peptide sequences is captured used to generate importance scores by residue position. Another important contribution of this paper is to provide a way to quantify the performance of XAI methods for TCR-pMHC binding. Typically XAI methods are evaluated qualitatively (e.g. in image analysis), but in the context of immunology, interpretations that match intuition are challenging to justify. We introduce **TCR-XAI**, a compilation of 274 experimentally-determined X-ray crystallography structures of TCR-pMHC complexes. For each complex we determine interaction distance between the CDR3 regions and peptide. Using this benchmark, we can determine whether the importance scores over the input produced by any particular method matches the expected interaction shown in the experimental structure.

We performed extensive evaluation of QCAI against other post hoc methods and demonstrate the benefit of incorporating cross-attention. We conduct an extensive comparison with other methods over the TCR-XAI benchmark and demonstrate that QCAI achieve state-of-the-art performance. We also analyze two case studies of TCR-pMHC systems to highlight mechanisms identified by QCAI.

## 2 PRELIMINARIES

In this section we first outline some basic concepts for self-attention and cross-attention and then discuss limitations in existing post-hoc methods. Transformer-based architectures typically consist of  $L$  stacked encoder layers, or a combination of encoder and decoder layers. Each layer comprises two primary components: Multi-Head Attention (MHA) and a Feed-Forward Network (FFN), each followed by layer normalization and residual connections (Vaswani, 2017). The distinction between encoder and decoder modules lies in their input structure and the type of attention mechanism employed.

In the encoder, each layer takes the output of the previous layer  $h_{l-1} \in \mathbb{R}^{N \times d}$  and computes  $h_l \in \mathbb{R}^{N \times d}$ , where  $N$  is the number of tokens and  $d$  is the hidden dimension. In contrast, the decoder layer integrates two inputs:  $h_{l-1} \in \mathbb{R}^{N \times d}$  from the previous decoder layer, and  $h'_l \in \mathbb{R}^{N' \times d}$  from the corresponding encoder layer. The decoder output remains  $h_l \in \mathbb{R}^{N \times d}$ , with  $N'$  denoting the number of source tokens.

These inputs are linearly projected into query ( $Q_l$ ), key ( $K_l$ ), and value ( $V_l$ ) matrices for the MHA computation. For encoder, it could be computed following  $Q_l = W_l^Q h_{l-1}$ ,  $K_l = W_l^K h_{l-1}$ , and  $V_l = W_l^V h_{l-1}$ . For decoder, it could be computed following  $Q_l = W_l^Q h_{l-1}$ ,  $K_l = W_l^K h'_l$ , and

108  $V_l = W_l^V h_l'$ . Where  $W_l^Q, W_l^K, W_l^V \in \mathbb{R}^{d \times d}$  are trainable projection matrices. For brevity, bias  
 109 terms are omitted. Considering a single attention head for simplicity, the attention matrix  $A_l$  for  
 110 layer  $l$  is computed as:

$$111 \quad 112 \quad 113 \quad A_l = \text{softmax} \left( \frac{Q_l K_l^\top}{\sqrt{d}} \right).$$

114 The shape of  $A_l$  is  $\mathbb{R}^{N \times N}$  for the encoder and  $\mathbb{R}^{N \times N'}$  for the decoder. The output of the attention  
 115 module is computed as:  $h_l = W_l^O (A_l V_l + h_{l-1}) \in \mathbb{R}^{N \times d}$ , where  $W_l^O \in \mathbb{R}^{d \times d}$  is a learnable output  
 116 projection matrix. Outputs from multiple attention heads are concatenated before being linearly  
 117 transformed.

## 118 2.1 LIMITATIONS OF CURRENT INTERPRETABILITY METHODS FOR TRANSFORMERS

120 Several post-hoc interpretability methods, such as TokenTM (Wu et al., 2024a), AttnLRP (Achitbat  
 121 et al., 2024), and AttCAT (Qiang et al., 2022), have demonstrated reliable performance on encoder-  
 122 only transformer models that rely on self-attention. However, these methods are not designed to  
 123 extract the interaction information from cross-attention found in decoder layers. As a result, their ap-  
 124 plicability is limited in models that include decoder components, such as TULIP (Meynard-Piganeau  
 125 et al., 2024) and MixTCRpred (Croce et al., 2024).

126 The core distinction between self-attention and cross-attention lies in the source of the key ( $K$ ) and  
 127 value ( $V$ ) matrices. While self-attention derives  $Q$ ,  $K$ , and  $V$  from the same input, cross-attention uses  
 128 separate inputs for  $Q$  and  $(K, V)$ . Consequently, the attention matrix  $A$  in cross-attention has  
 129 dimensions  $\mathbb{R}^{N \times N'}$  instead of  $\mathbb{R}^{N \times N}$ , where  $N$  is the number of query tokens and  $N'$  is the number  
 130 of key tokens. Additionally,  $A$  now represents the fused information from both modalities. This  
 131 asymmetry poses a challenge for interpretability: the attention matrix no longer provides a direct  
 132 measure of query token importance of one input modality, making it difficult to attribute model  
 133 predictions to input query tokens.

## 135 3 QUANTIFYING CROSS-ATTENTION INTERACTION

137 In this section we present our main contribution, which is a way to handle the aforementioned  
 138 asymmetry so that cross-attention can be captured. Since the attention matrix is computed as a scaled  
 139 dot-product  $QK^\top$ , which captures the cosine similarity between query and key representations,  
 140 interpreting the cross-attention mechanism can be structured into three key steps: 1. identifying  
 141 which components of the attention matrix contribute most significantly to the model’s prediction, 2.  
 142 decomposing these importance values into contributions from the query and key inputs, respectively,  
 143 and 3. aggregating the cross-attention explanation with other layers’ explanation.

144 Inspired by GradCAM (Selvaraju et al., 2017), we propose to compute the importance of the attention  
 145 matrix  $A_l$  at layer  $l$  using the gradient of the loss  $L^c$  with respect to  $A_l$  for a target class  $c$ , in  
 146 conjunction with the attention weights themselves. Specifically, we define the importance score map  
 147 as:

$$148 \quad 149 \quad \mathbf{S}(A_l) = \mathbb{E}_H \left( \text{ReLU} \left( \frac{\partial L^c}{\partial A_l} \odot A_l \right) \right) + I \in \mathbb{R}^{N \times N'},$$

150 where  $\mathbb{E}_H(\cdot)$  denotes averaging across all attention heads,  $\odot$  represents element-wise multiplication,  
 151 and  $I$  denotes the identity matrix for residue connection. This formulation highlights the attention  
 152 entries that both have high weights and contribute significantly to the class-specific loss. The next  
 153 step is to quantify this attention importance map into contributions from the query and key inputs.  
 154 By analyzing the structure of the attention matrix, which serves as a soft alignment between queries  
 155 and keys, we aim to attribute the importance scores back to the input tokens in both sequences.

### 156 3.1 QUANTIFYING QUERY IMPORTANCE FROM CROSS-ATTENTION

158 For the query input  $Q_l$  at layer  $l$ , its importance scores with respect to the loss  $L^c$  for class  $c$  can be  
 159 estimated in a GradCAM-style fashion as:

$$160 \quad 161 \quad \mathbf{S}(Q_l) = \text{ReLU} \left( \frac{\partial L^c}{\partial Q_l} \odot Q_l \right),$$

---

162 where  $\odot$  denotes element-wise multiplication. To obtain token-level importance scores from this  
 163 matrix, we compute the column-wise maximum:

164

$$\omega_l^Q = \arg \max_i \mathbf{S}(Q_l)_{i,j} \in \mathbb{R}^N,$$

165

166 where  $i$  indexes the feature dimension,  $j$  indexes the query tokens, and  $\arg \max_i$  denotes the maximum  
 167 across the feature dimension. However, this importance score is intrinsic to  $Q_l$  itself and  
 168 does not reflect how  $Q_l$  is influenced by the attention mechanism. Explaining the attention matrix  
 169 is a key component of post-hoc methods for interpreting transformer models (Wu et al., 2024a). In  
 170 the case of cross-attention, the query and key matrices originate from different inputs, and thus the  
 171 resulting attention matrix is not necessarily square. To better understand how  $Q_l$  contributes within  
 172 the attention process, we define its attention-conditioned importance scores as  $\mathbf{S}(Q_l; A_l)$ , the query  
 173 importance modulated by the attention matrix  $A_l$ . We approximate this as:

174

175

$$\mathbf{S}(Q_l; A_l) \propto \frac{\partial L^c}{\partial A_l} \cdot Q_l,$$

176

177 **where  $\cdot$  is matrix product.** From the previous step, we have already obtained the attention importance  
 178 map  $\frac{\partial L^c}{\partial A_l} \odot A_l$ . We now seek a transformation that allows us to infer  $\mathbf{S}(Q_l; A_l)$  from this. The  
 179 attention matrix is computed via scaled dot-product as  $A_l = Q_l K_l^\top$  with softmax and  $\sqrt{d}$  ignored  
 180 for simplicity. We can express with linear operations (e.g., ReLU,  $\mathbb{E}_H$ , and  $(\cdot) + I$ ) ignored for  
 181 simplicity.:

182

$$\mathbf{S}(A_l) = \frac{\partial L^c}{\partial A_l} \cdot Q_l K_l^\top,$$

183

184 To isolate the influence of  $Q_l$ , we need to eliminate  $K_l^\top$  from the right-hand side. Since  $K_l$  is not  
 185 guaranteed to be square or invertible, we employ the Moore-Penrose pseudoinverse:

186

187

$$\frac{\partial L^c}{\partial A_l} \cdot Q_l K_l^\top = \mathbf{S}(A_l)$$

188

$$\frac{\partial L^c}{\partial A_l} \cdot Q_l = \mathbf{S}(A_l) \cdot K_l (K_l^\top K_l)^{-1} \in \mathbb{R}^{N \times d}.$$

189

190 This yields a decomposition of attention importance into the query space. Then, the importance  
 191 scores corresponding to the token part can be obtained following:

192

193

$$\omega_l^A = \arg \max_i \left( \frac{\partial L^c}{\partial A_l} \cdot Q_l \right)_{i,j} \in \mathbb{R}^N,$$

194

195 where  $i$  indexes the feature dimension,  $j$  indexes the query tokens, and  $\arg \max_i$  denotes maximum  
 196 taken over feature dimension. However, to ensure robustness, particularly in cases where  $Q_l$  is  
 197 also influenced by other layers. We conservatively combine this result with the intrinsic query  
 198 importance:

199

200

$$\mathbf{S}(Q_l; A_l) = \max \left( \omega_l^A, \omega_l^Q \right).$$

201

202 Here, the maximum is applied element-wise to capture the strongest importance attribution from  
 203 either source.

204

### 3.2 QUANTIFYING KEY IMPORTANCE FROM CROSS-ATTENTION

205

206 Similar to the approach used to extract query importance scores, the key matrix importance can  
 207 also be quantified into two components: (1) the intrinsic importance of the key matrix, denoted as  
 208  $\mathbf{S}(K_l)$ , and (2) the attention-conditioned importance,  $\mathbf{S}(K_l; A_l)$ , which reflects how the key matrix  
 209 contributes to the attention mechanism.

210

211 The intrinsic importance of the key matrix with respect to the loss  $L^c$  for class  $c$  can be estimated  
 212 using a GradCAM-style formulation:

213

214

$$\mathbf{S}(K_l) = \text{ReLU} \left( \frac{\partial L^c}{\partial K_l} \odot K_l \right).$$

215

216 To obtain token-level importance scores from this matrix, we compute the column-wise maximum:  
217

$$218 \quad \omega_l^K = \arg \max_i \mathbf{S}(K_l)_{i,j} \in \mathbb{R}^{N'}. \\ 219$$

220 where  $i$  indexes the feature dimension,  $j$  indexes the key tokens, and  $\arg \max$  denotes the maximum  
221 across the feature dimension ( $d$ ). Similar to the issue encountered in query attention quantification,  
222 the attention matrix is no longer necessarily square for key attention quantification. However,  
223 compared to decomposing query importance, extracting key importance from the attention matrix  
224 is more straightforward, since attention explicitly maps queries into the key space. Thus, we can  
225 directly analyze the attention matrix to determine which key tokens exert the strongest influence on  
226 the query representations. Because transformer models rely primarily on token-level outputs, we  
227 focus on interpreting token-level activations. The attention matrix  $A \in \mathbb{R}^{N \times N'}$  indicates how each  
228 query token (rows) attends to the key tokens (columns). To evaluate the overall importance of each  
229 key token in guiding the query representations, we compute the maximum relevance of each key  
230 across all queries and attention heads:

$$231 \quad \omega_l^{A'} = \arg \max_i \left( \mathbb{E}_H \left( \text{ReLU} \left( \frac{\partial L^c}{\partial A_{i,j}} \cdot A_{i,j} \right) \right) \right) \in \mathbb{R}^{N'}, \\ 232$$

233 where  $\mathbb{E}_H$  denotes averaging over attention heads and  $i$  and  $j$  index the queries and keys respectively,  
234 and  $\arg \max_i$  denotes the maximum across the feature dimension.

235 Finally, we combine this attention-derived importance with the intrinsic importance to produce a  
236 robust estimate of key token relevance:

$$237 \quad \mathbf{S}(K_l; A_l) = \max \left( \omega_l^{A'}, \omega_l^K \right), \\ 238$$

239 where the maximum is taken element-wise to reflect the highest attribution signal from either source.

### 240 3.3 AGGREGATION OF LAYER IMPORTANCE SCORES

241 Inspired by the attention flow perspective (Abnar & Zuidema, 2020), we aggregate token-level im-  
242 portance scores across layers to track how relevance propagates from the final output back through  
243 the decoder and encoder layers. Let  $k$  denote the index of the first decoder layer (with cross-  
244 attention) encountered when traversing the model from the output layer backwards. All subsequent  
245 layers with smaller indices are assumed to be encoder layers with self-attention. To capture how  
246 importance propagates through these layers, we define the aggregated token-level importance scores  
247 at layer  $k$ , denoted by  $\tilde{\mathbf{S}}_k$ , recursively as follows:

$$248 \quad \tilde{\mathbf{S}}_k = \begin{cases} \mathbf{S}(Q_k; A_k) \cdot \tilde{\mathbf{S}}_{k+1} & \text{(query)} \\ \mathbf{S}(K_k; A_k) \cdot \tilde{\mathbf{S}}_{k+1} & \text{(key)} \end{cases} \\ 249$$

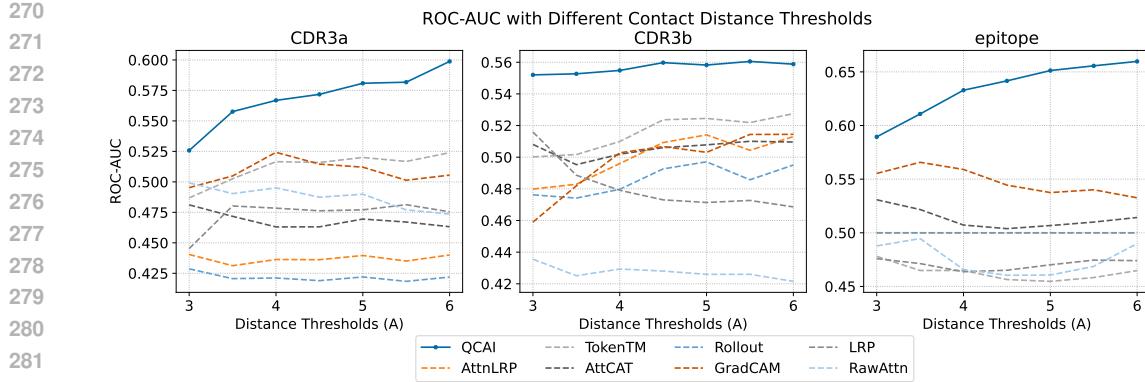
250 In models with multiple decoder blocks that contain cross-attention, importance interactions may  
251 diverge and converge at different points. To handle such cases, we adopt a conservative strategy and  
252 aggregate importance via element-wise maximum to retain the most influential attribution signal:

$$253 \quad \tilde{\mathbf{S}}_k = \begin{cases} \max \left( \mathbf{S}(Q_k; A_k), \tilde{\mathbf{S}}_{k+1} \right) & \text{(query)} \\ \max \left( \mathbf{S}(K_k; A_k), \tilde{\mathbf{S}}_{k+1} \right) & \text{(key)} \end{cases} \\ 254$$

255 These recursive rules ensure that attention importance is correctly traced through both cross-  
256 attention and self-attention components. Consequently, if the explanation path contains any decoder  
257 block with cross-attention, the final output of our QCAI method will be a vector of token-level  
258 importance scores, indicating the contribution of each input token to the model’s prediction.

## 259 4 EXPERIMENTAL RESULTS AND DISCUSSION

260 We first evaluate our proposed QCAI method using a state-of-art BERT-based model named TULIP,  
261 a transformer architecture tailored for predicting TCR-pMHC binding, which focuses on the role



270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
300  
301  
302  
303  
304  
305  
306  
307  
308  
309  
310  
311  
312  
313  
314  
315  
316  
317  
318  
319  
320  
321  
322  
323

Figure 2: ROC-AUC of predicted importance scores for TCR-pMHC binding site identification across a threshold of interaction distances demonstrates that QCAI surpasses competing methods in all cases.

of cross-attention and outperforms one of the widely used baseline models, NetTCR-2.2 (Jensen & Nielsen, 2023). TULIP adopts an encoder-decoder design and processes three modalities in parallel: CDR3a, CDR3b, and peptide sequences (Meynard-Piganeau et al., 2024). Each encoder independently transforms input sequences into latent feature representations, while decoder layers model inter-sequence interactions (Devlin, 2018; Vaswani, 2017). As a self-regressive generative model, TULIP estimates the conditional probability distribution of each sequence (e.g., peptide) conditioned on the others (e.g., CDR3a, CDR3b) (Meynard-Piganeau et al., 2024).

We compare our QCAI method against several existing post-hoc interpretability techniques, including AttnLRP (Achitbat et al., 2024), TokenTM (Wu et al., 2024a), AttCAT (Qiang et al., 2022), Rollout (Abnar & Zuidema, 2020), GradCAM (Selvaraju et al., 2017), LRP (Binder et al., 2016), and RawAttn (Wiegreffe & Pinter, 2019). For methods that require aggregation across all layers and compute on attention matrix, we apply them exclusively to the self-attention layers and omit cross-attention components, as these competing methods do not support cross-attention explanations. All experiments were implemented in Python using the PyTorch framework. Testing was conducted on a local workstation equipped with two NVIDIA A2000 GPUs, 16 Intel Xeon E5 CPU cores, and 64 GB of RAM.

#### 4.1 A BENCHMARK FOR TCR-pMHC BINDING INTERPRETATION

To quantitatively assess the quality of interpretability methods, we constructed a benchmark that we call TCR-XAI using structural data from TCR-pMHC complexes. We collected 274 valid samples from the STCRDab (Leem et al., 2018) and TCR3d 2.0 (Lin et al., 2025) datasets, which consist of 213 (77.7%) MHC-I samples and 61 (22.3%) MHC-II samples. Only samples with complete TCR  $\alpha$  and  $\beta$  chains, full peptide sequences, intact CDR3 regions, and non-overlapping MHC and peptide chain IDs were retained. Statistics of the benchmark set are discussed in Section A.9 of the Appendix. For each sample, we computed residue-level distances: (1) from each CDR residue to the closest atom in the peptide, and (2) from each peptide residue to the closest atom in any CDR region. A smaller distance indicates a stronger interaction, and atomic distance as a proxy for ground-truth importance. We believe this is a simple, yet highly useful assumption since protein folding and, by extension, protein-protein interactions are most routinely evaluated by the closeness of packing, which signals the exclusion of water molecules and demands the formation of all possible hydrogen bonds (without water molecules). Other types of interatomic interactions such as hydrophobic contacts and ionic bonds contribute to binding, but they generally cannot be realized without exclusion of water. Thus, the formation of a stable protein-protein interface has a sharp distance threshold, above-which the interaction is not likely to be stable. The importance of residue-level distance is evident in prior work, starting with TCRdist Mayer-Blackwell et al. (2021), a classic unsupervised method for TCR-pMHC binding prediction. It defines the TCRdist distance as “the similarity-weighted mismatch distance between the potential pMHC-contacting loops of the two receptors.” Using distance as an indicator is also common in TCR-pMHC transformer model

324 explanations, though typically for qualitative rather than quantitative evaluation (e.g., PISTE (Feng  
325 et al., 2024) and TCR-BERT (Wu et al., 2024b)).  
326

## 327 4.2 ROC ANALYSIS AND PERTURBATION EXPERIMENTS 328

329 To evaluate the explainability of different post-hoc interpretation methods, we quantitatively assess  
330 their ability to identify true TCR-pMHC binding sites using the TCR-XAI benchmark. We computed  
331 ROC curves (ROC curves for each threshold are shown in Section A.7 in the Appendix.) by compar-  
332 ing predicted residue importance against ground-truth binding site annotations derived from struc-  
333 tural data, where the ground-truth was defined according to distance threshold between 3 and 6 Å,  
334 with the ROC using predicted importance as the threshold. As shown in Figure 2, QCAI achieves  
335 AUCs of 0.5492, 0.5489, and 0.6024 for CDR3a, CDR3b, and peptide respectively and consistently  
336 outperforms other methods. Notably, QCAI exceeds 0.6 on the peptide chain, demonstrating strong  
337 alignment between its predicted importance scores and the underlying structural binding interac-  
338 tions.

339 We also conducted perturbation studies on to assess whether each method identifies important  
340 residues; we adopt two commonly used metrics: Log-Odds Score (LOdds) and Area Over the Per-  
341 turbation Curve (AOPC). AOPC measures explanation quality by averaging the drop in model con-  
342 fidence as top- $k$  important features are removed. Higher AOPC indicates better alignment between  
343 explanation and model behavior. LOdds computes the change in log-odds of the model’s prediction  
344 before and after perturbing a feature. A larger LOdds value indicates greater importance of the per-  
345 turbed feature. Perturbation is implemented by replacing the  $k$  highest-scoring tokens with padding  
346 tokens (<PAD>). We evaluate the CDR3a, CDR3b and peptide chains separately, with  $k=4$   
347 for the CDR3a and CDR3b chains, and  $k=7$  for peptides to match the average number of predicted  
348 binding residues across TCR-XAI.

Chains	CDR3a $k=4$		CDR3b $k=4$		Peptide $k=7$	
	LOdds	AOPC	LOdds	AOPC	LOdds	AOPC
<b>QCAI (Ours)</b>	<b>-3.328</b>	0.014	<b>-3.498</b>	<b>0.045</b>	<b>-1.470</b>	<b>0.013</b>
AttnLRP (Achtibat et al., 2024)	-2.481	0.020	-2.662	0.032	-0.017	0.000
TokenTM (Wu et al., 2024a)	-2.195	0.021	-2.383	0.032	-0.736	0.012
AttCAT (Qiang et al., 2022)	-2.825	0.020	-3.131	0.044	-0.694	0.006
Rollout (Abnar & Zuidema, 2020)	-2.356	<b>0.022</b>	-2.653	0.032	-0.044	-0.001
GradCAM (Selvaraju et al., 2017)	-2.700	0.019	-3.112	0.038	-1.004	0.009
LRP (Binder et al., 2016)	-2.938	0.020	-3.127	0.043	-1.167	0.011
RawAttn (Wiegreffe & Pinter, 2019)	-2.734	0.015	-3.250	0.039	-0.691	0.010

360 Table 1: Perturbation experiment results using fixed thresholds. Thresholds for the  $\alpha$  and  $\beta$  chains  
361 are  $k=4$ , and for the peptide chain  $k=7$ . The average number of binding regions are 3.64, 4.12, and  
362 7.05 for  $\alpha$ ,  $\beta$ , and peptide chains respectively.  
363

364 Table 1 shows that QCAI consistently outperforms other methods across most metrics. QCAI  
365 achieves the most negative LOdds and highest AOPC scores in the CDR3b and peptide chains,  
366 indicating greater disruption to the model’s confidence when informative residues are perturbed.  
367 Although Rollout outperforms QCAI in AOPC on the CDR3a chain, QCAI still achieves the best  
368 LOdds score.  
369

## 370 4.3 IDENTIFICATION OF BINDING REGION RESIDUES WITH IMPORTANCE SCORES 371

372 Using the TCR-XAI benchmark we construct an evaluation metric that we call *Binding Region Hit*  
373 *Rate* (BRHR). To compute BRHR, we first choose a percentile threshold  $t \in (0, 1]$  and identify the  
374 top  $t$  fraction of residues with respect to highest importance scores  $S$ . Each of these residues is  
375 marked a hit if its interaction distance is in the top  $t$  fraction of interaction distances. We compute  
376 the hit rate for each input sequence type in each sample and take the mean across TCR-XAI to obtain  
377 the final BRHR. This metric reflects the proportion of true binding residues (according to structural  
proximity) that are successfully identified by the explanation method.

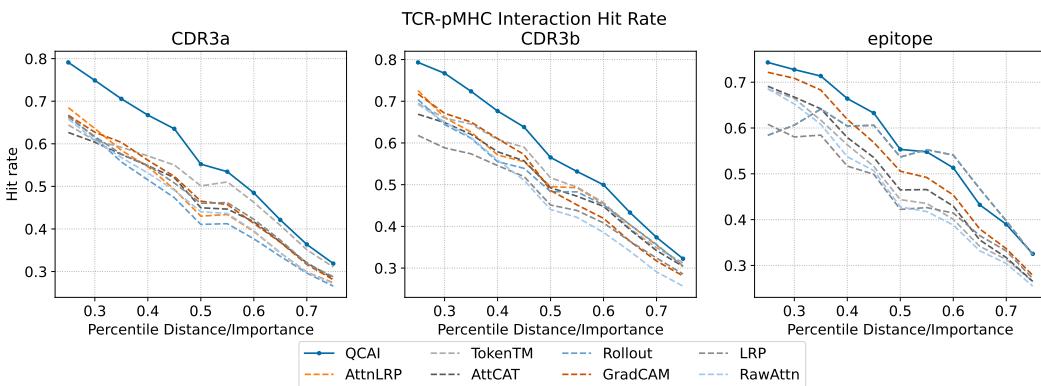


Figure 3: Comparison of TCR-pMHC Binding Region Hit Rate (BRHR) across different methods on different chains. At any selected percentile of distance/importance, the higher the hit rate the more closely the importance tracks physical interaction distance. QCAI surpasses other methods in all practical cases.

As shown in Figure 3, our method achieves state-of-the-art performance compared to all other explanation methods. For the peptide chain, our method consistently outperforms all other methods before the 50th percentile. After this threshold other methods prevail but have high false positive rate of other methods (as seen in the ROC analysis). We postulate that the latter effect is due to the fact that these methods can only access self-attention weights from the encoder and cannot benefit from the regulatory influence of cross-attention layers.

#### 4.4 CASE STUDIES

To highlight the ability of QCAI to assist in the interpretation of TCR-pMHC binding we discuss two specific examples, one for CD8+ T cells and one for CD4+ T cells. In both cases the analysis of importance using QCAI finds residue positions in CDR3s that form critical contacts with **epitope** peptides and, by revealing unconstrained positions in longer CDR loops, can explain large differences in TCR-peptide-HLA binding affinity.

In the first case study (Figure 4(a)) we consider the immunodominant CD8+ T-cell peptide from the influenza matrix protein which has been used to understand influenza T cell response. Multiple crystal structures (1OGA (Stewart-Jones et al., 2003) and 5TEZ (Yang et al., 2017)) of different TCRs recognizing this peptide have revealed a common mode of binding that involves the insertion of a single CDR3b sidechain (R98 in the 1OGA structure) into a notch between the peptide and the HLA-A2 alpha-2 helix and, otherwise, makes numerous contacts with the HLA-A2, whose shape depends on the peptide. In one distinct example, the TCR in the 5TEZ structure is rotated by 40 degrees around the HLA-TCR axis to create a very different group of TCR-HLA-A2 contacts, but this TCR also places a CDR3b sidechain (W99 in the 5TEZ structure) in the notch between peptide and HLA-A2 alpha-2 helix. Consistent with the common aspect of binding, the QCAI evaluation finds importance in the position of the notch-binding residue and in several N-terminal flanking positions of CDR3b. The distinct aspect of binding for the two TCRs arises in the longer and less-constrained CDR3a for the 5TEZ TCR, which may explain its 25-fold lower affinity than for the 1OGA TCR. We note that for both binding orientations, AttnLRP and TokenTM produce weaker importance scores overall.

The second case study considers a self-antigen in the autoimmune disease of rheumatoid arthritis. The HLA-DR4-bound citrullinated peptide, named vimentin-64cit59-71, has been analyzed in the complex with two different TCRs (Loh et al., 2024) (indicated with PDB codes, 8TRR and 8TRQ in Figure 4(b)). The QCAI evaluation finds an overall similar number of important positions in the two TCRs, including a concentration of importance along one edge of the hairpin formed by the CDR3s in both TCRs (highlighted with a dark outline in Figure 4(b)). The CDR3a contributes the largest direct contact with the peptide in both complexes. Interestingly, the CDR3b of the 5-fold-lower-affinity 8TRQ complex is longer and contains more positions of lower importance, again suggesting that the entropic cost of ordering this loop is responsible for the reduced affinity. For this case study,

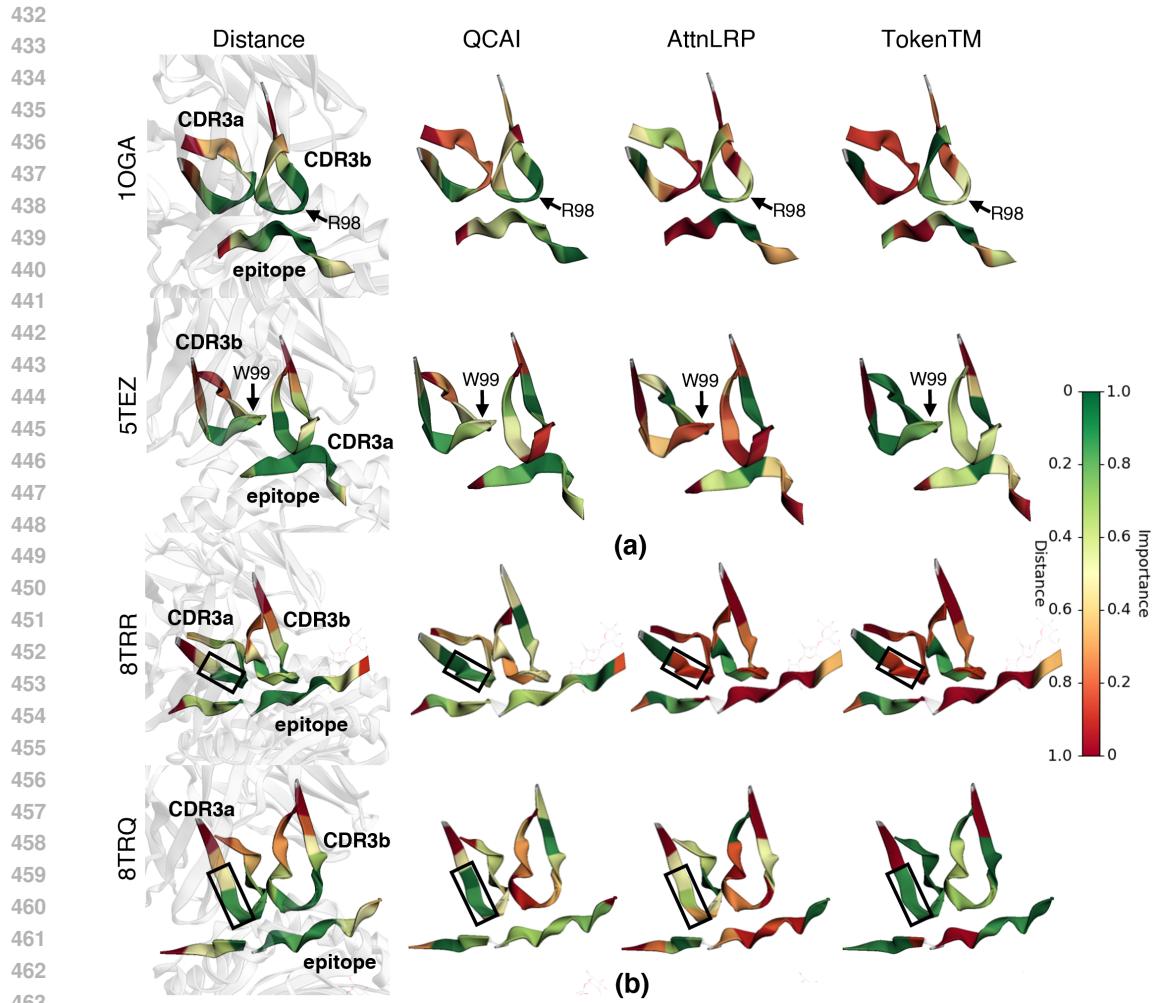
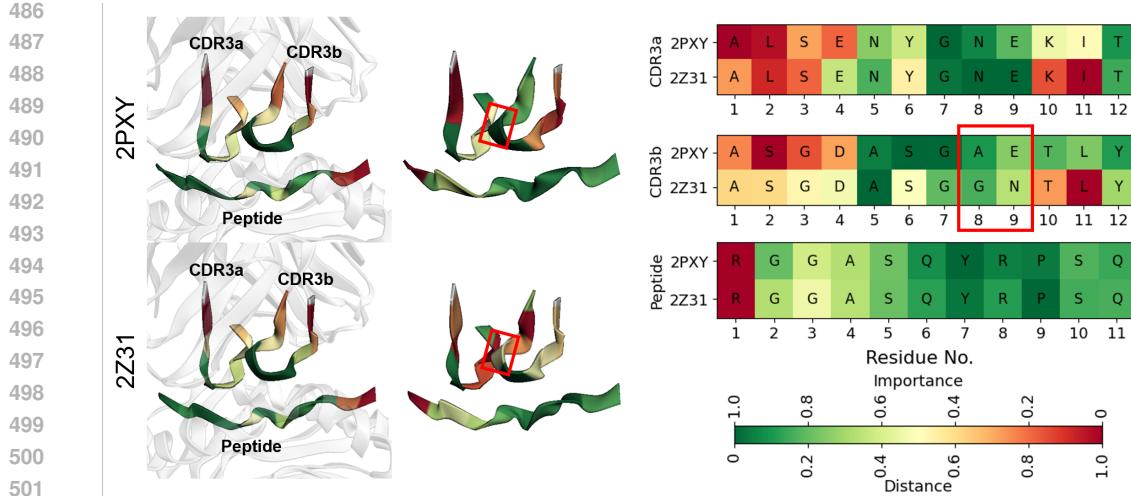


Figure 4: Case studies on systems from TCR-XAI. (a) We consider the same TCR-pMHC bound in two distinct binding orientations. For this system QCAI identifies key residues from both orientations. (b) We consider the same pMHC bound to two distinct TCRs. Here QCAI identifies the importance of the hairpin region of CDR3a in both cases.

AttnLRP does not produce meaningful results while TokenTM does not capture the importance of residues in the peptide proximal to the CDR3a hairpin.

To investigate how QCAI explanations differ for similar TCR-pMHC complexes, we conducted a case study on two TCR-MHCII-peptide structures, 2PXY and 2Z31, which investigates whether a germline-encoded motif structurally guides TCR recognition of MHC (Feng et al., 2007). They differ by two amino acids in the CDR3b loop (Feng et al., 2007). To convince chain alignment, amino acids were re-indexed starting from 1 for each chain. As shown in Figure 5, QCAI with TULIP assigns similar importance scores to the peptide in both complexes but produces different pattern of importance for the CDR3b loop. Both complexes correctly highlight A5 as an important contact region, and QCAI identifies additional contact sites in 2PXY. In 2PXY, residues S6 and G7 receive higher scores, whereas the corresponding region in 2Z31 receives lower scores, where are also the contact regions. These results indicate that QCAI can detect critical contact regions even with minor sequence changes. However, such changes can affect the overall explanation quality.



503 Figure 5: Case studies of two closely related TCR-pMHC complexes from TCR-XAI. These  
504 complexes differ by only two amino acids in the CDR3b, highlighted in the figure with red rectangles.  
505

## 5 CONCLUSIONS

510 In this paper, we present Quantifying Cross-Attention Interaction (QCAI) to interpret the cross-  
511 attention in the decoders of transformer models, aiming to better understand encoder-decoder TCR-  
512 pMHC binding prediction models. QCAI quantifies the importance of the cross-attention matrix  
513 into contributions from query and key inputs, revealing how they influence each other. To rigorously  
514 evaluate the explanations, we created a new structural explanation benchmark, TCR-XAI, along  
515 with a novel evaluation metric, the Binding Region Hit Rate (BRHR). On this benchmark, QCAI  
516 achieves state-of-the-art results across perturbation metrics (LOdds and AOPC), ROC-AUC, ROC  
517 curve analysis, and BRHR.

### 5.1 FUTURE WORK

522 In future work, we plan to pursue two primary directions: (1) extending the metrics used to evaluate  
523 explainability, and (2) applying QCAI to broader range of immunological and protein-protein interaction  
524 tasks. Beyond distance-based measures, energy functions (e.g., REF15 (Alford et al., 2017))  
525 offer a promising alternative for quantifying explanations in TCR-pMHC binding prediction. Investigating  
526 a range of energy-based models to better understand the relationship between explainability  
527 and protein energetics will be an important next step.

528 Given the emergence of several cross-attention models for protein-protein interactions and immunological  
529 tasks, such as PALM-H3 (He et al., 2024) for antigen generation, UniPMT (Zhao et al., 2025)  
530 for peptide-MHC prediction, ProtAttBA (Liu et al., 2025) for antibody-antigen prediction, and HB-  
531 Former (Zhang et al., 2024) for human-virus interaction identification, QCAI provides a method for  
532 opening the black box of cross-attentions in these models and revealing their underlying mecha-  
533 nisms. In addition, QCAI can be extended beyond these applications. For instance, we have already  
534 applied it to CLIP encoders with cross-attention, as discussed in Appendix A.10. Exploring broader  
535 applications of QCAI across these tasks and domains is an important direction for future work.

536 Transformers are widely used for TCR-pMHC binding prediction, but they remain black-box mod-  
537 els. While post-hoc methods like QCAI improve explainability, they cannot directly integrate these  
538 insights into prediction. Beyond post-hoc methods, an important future direction is to develop  
539 explain-by-design models that provide inherent explainability and utilize mechanistic TCR-pMHC  
insights to improve predictive performance.

---

540 REFERENCES  
541

542 Samira Abnar and Willem Zuidema. Quantifying attention flow in transformers. In *Proceedings of*  
543 *the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 4190–4197, 2020.

544 Reduan Achtibat, Sayed Mohammad Vakilzadeh Hatefi, Maximilian Dreyer, Aakriti Jain, Thomas  
545 Wiegand, Sebastian Lapuschkin, and Wojciech Samek. Attnlrp: Attention-aware layer-wise relevance  
546 propagation for transformers. In *Forty-first International Conference on Machine Learning*,  
547 2024.

548 Rebecca F Alford, Andrew Leaver-Fay, Jeliazko R Jeliazkov, Matthew J O’Meara, Frank P DiMaio,  
549 Hahnbeom Park, Maxim V Shapovalov, P Douglas Renfrew, Vikram K Mulligan, Kalli Kappel,  
550 et al. The rosetta all-atom energy function for macromolecular modeling and design. *Journal of*  
551 *chemical theory and computation*, 13(6):3031–3048, 2017.

552 Ameen Ali, Thomas Schnake, Oliver Eberle, Grégoire Montavon, Klaus-Robert Müller, and Lior  
553 Wolf. Xai for transformers: Better explanations through conservative propagation. In *International*  
554 *conference on machine learning*, pp. 435–451. PMLR, 2022.

555 Dmitry V Bagaev, Renske MA Vroomans, Jerome Samir, Ulrik Stervbo, Cristina Rius, Garry  
556 Dolton, Alexander Greenshields-Watson, Meriem Attaf, Evgeny S Egorov, Ivan V Zvyagin, et al.  
557 Vdjdb in 2019: database extension, new analysis infrastructure and a t-cell receptor motif com-  
558 pendium. *Nucleic acids research*, 48(D1):D1057–D1062, 2020.

559 Avik Bhattacharya, James O Wrabl, Samuel J Landry, and Ramgopal R Mettu. Parallel computa-  
560 tion of conformational stability for cd4+ t-cell epitope prediction. In *2023 IEEE International*  
561 *Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 88–93. IEEE, 2023.

562 Alexander Binder, Grégoire Montavon, Sebastian Lapuschkin, Klaus-Robert Müller, and Wojciech  
563 Samek. Layer-wise relevance propagation for neural networks with local renormalization layers.  
564 In *Artificial Neural Networks and Machine Learning–ICANN 2016: 25th International Confer-  
565 ence on Artificial Neural Networks, Barcelona, Spain, September 6-9, 2016, Proceedings, Part II*  
566 25, pp. 63–71. Springer, 2016.

567 Rémy Bosselut. T cell antigen recognition: Evolution-driven affinities. *Proceedings of the National*  
568 *Academy of Sciences*, 116(44):21969–21971, 2019.

569 Tysheena Charles, Daniel L Moss, Pawan Bhat, Peyton W Moore, Nicholas A Kummer, Avik Bhat-  
570 tacharya, Samuel J Landry, and Ramgopal R Mettu. Cd4+ t-cell epitope prediction by combined  
571 analysis of antigen conformational flexibility and peptide-mhcii binding affinity. *Biochemistry*,  
572 61(15):1585–1599, 2022.

573 Aditya Chattpadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N Balasubramanian. Grad-  
574 cam++: Generalized gradient-based visual explanations for deep convolutional networks. In *2018*  
575 *IEEE winter conference on applications of computer vision (WACV)*, pp. 839–847. IEEE, 2018.

576 Hila Chefer, Shir Gur, and Lior Wolf. Generic attention-model explainability for interpreting bi-  
577 modal and encoder-decoder transformers. In *Proceedings of the IEEE/CVF international confer-  
578 ence on computer vision*, pp. 397–406, 2021.

579 Junwei Chen, Bowen Zhao, Shenggeng Lin, Heqi Sun, Xueying Mao, Meng Wang, Yanyi Chu,  
580 Liang Hong, Dong-Qing Wei, Min Li, et al. Tepcam: Prediction of t-cell receptor–epitope binding  
581 specificity via interpretable deep learning. *Protein Science*, 33(1):e4841, 2024.

582 Lewis Cornwall, Grisha Szep, James Day, SR Gokul Krishnan, David Carter, Jamie Blundell, Lilly  
583 Wollman, Neil Dalchau, and Aaron Sim. Fine-tuned protein language models capture t cell re-  
584 ceptor stochasticity. In *NeurIPS 2023 Generative AI and Biology (GenBio) Workshop*, 2023.

585 Giancarlo Croce, Sara Bobisse, Dana Léa Moreno, Julien Schmidt, Philippe Guillame, Alexandre  
586 Harari, and David Gfeller. Deep learning predictions of tcr-epitope interactions reveal epitope-  
587 specific chains in dual alpha t cells. *Nature Communications*, 15(1):3211, 2024.

594 Pradyot Dash, Andrew J Fiore-Gartland, Tomer Hertz, George C Wang, Shalini Sharma, Aisha  
595 Souquette, Jeremy Chase Crawford, E Bridie Clemens, Thi HO Nguyen, Katherine Kedzierska,  
596 et al. Quantifiable predictive features define epitope-specific t cell receptor repertoires. *Nature*,  
597 547(7661):89–93, 2017.

598

599 Mark M Davis and Pamela J Bjorkman. T-cell antigen receptor genes and t-cell recognition. *Nature*,  
600 334(6181):395–402, 1988.

601

602 Jacob Devlin. Bert: Pre-training of deep bidirectional transformers for language understanding.  
603 *arXiv preprint arXiv:1810.04805*, 2018.

604

605 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep  
606 bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of  
607 the North American chapter of the association for computational linguistics: human language  
608 technologies, volume 1 (long and short papers)*, pp. 4171–4186, 2019.

609

610 Alice Driessen, Jannis Born, Rocío Castellanos Rueda, Sai T Reddy, and Marianna Rapsomaniki.  
611 Modeling car response at the single-cell level using conditional ot. In *NeurIPS 2024 Workshop  
612 on AI for New Drug Modalities*, 2024.

613

614 Rudresh Dwivedi, Devam Dave, Het Naik, Smiti Singhal, Rana Omer, Pankesh Patel, Bin Qian,  
615 Zhenyu Wen, Tejal Shah, Graham Morgan, et al. Explainable ai (xai): Core ideas, techniques,  
616 and solutions. *ACM Computing Surveys*, 55(9):1–33, 2023.

617

618 Dan Feng, Christopher J Bond, Lauren K Ely, Jennifer Maynard, and K Christopher Garcia. Struc-  
619 tural evidence for a germline-encoded t cell receptor–major histocompatibility complex interac-  
620 tion’codon’. *Nature immunology*, 8(9):975–983, 2007.

621

622 Ziyang Feng, Jingyang Chen, Youlong Hai, Xuelian Pang, Kun Zheng, Chenglong Xie, Xiujuan  
623 Zhang, Shengqing Li, Chengjuan Zhang, Kangdong Liu, et al. Sliding-attention transformer neu-  
624 ral architecture for predicting t cell receptor–antigen–human leucocyte antigen binding. *Nature  
625 Machine Intelligence*, 6(10):1216–1230, 2024.

626

627 Jacob Glanville, Huang Huang, Allison Nau, Olivia Hatton, Lisa E Wagar, Florian Rubelt, Xuhuai  
628 Ji, Arnold Han, Sheri M Krams, Christina Pettus, et al. Identifying specificity groups in the t cell  
629 receptor repertoire. *Nature*, 547(7661):94–98, 2017.

630

631 Haohuai He, Bing He, Lei Guan, Yu Zhao, Feng Jiang, Guanxing Chen, Qingge Zhu, Calvin Yu-  
632 Chian Chen, Ting Li, and Jianhua Yao. De novo generation of sars-cov-2 antibody cdrh3 with a  
633 pre-trained generative large language model. *Nature Communications*, 15(1):6867, 2024.

634

635 Ilka Hoof, Bjoern Peters, John Sidney, Lasse Eggers Pedersen, Alessandro Sette, Ole Lund, Søren  
636 Buus, and Morten Nielsen. Netmhcpn, a method for mhc class i binding prediction beyond  
637 humans. *Immunogenetics*, 61:1–13, 2009.

638

639 Huang Huang, Chunlin Wang, Florian Rubelt, Thomas J Scriba, and Mark M Davis. Analyzing  
640 the mycobacterium tuberculosis immune response by t-cell receptor clustering with gliph2 and  
641 genome-wide antigen screening. *Nature biotechnology*, 38(10):1194–1202, 2020.

642

643 Dan Hudson, Ricardo A Fernandes, Mark Basham, Graham Ogg, and Hashem Koohy. Can we  
644 predict t cell specificity with digital biology and machine learning? *Nature Reviews Immunology*,  
645 23(8):511–521, 2023.

646

647 Dan Hudson, Alex Lubbock, Mark Basham, and Hashem Koohy. A comparison of clustering models  
648 for inference of t cell receptor antigen specificity. *ImmunoInformatics*, 13:100033, 2024.

649

650 Mathias Fynbo Jensen and Morten Nielsen. Nettcr 2.2-improved tcr specificity predictions by com-  
651 bining pan-and peptide-specific training strategies, loss-scaling and integration of sequence simi-  
652 larity. *bioRxiv*, pp. 2023–10, 2023.

653

654 Alok V Joglekar and Guideng Li. T cell antigen discovery. *Nature methods*, 18(8):873–880, 2021.

---

648 Edita Karosiene, Claus Lundegaard, Ole Lund, and Morten Nielsen. Netmhccons: a consensus  
649 method for the major histocompatibility complex class i predictions. *Immunogenetics*, 64:177–  
650 186, 2012.

651

652 Dhavarakesh Karthikeyan, Colin Raffel, Benjamin Vincent, and Alex Rubinsteyn. Conditional gen-  
653 eration of antigen specific t-cell receptor sequences. In *NeurIPS 2023 Generative AI and Biology*  
654 (*GenBio*) Workshop, 2023.

655 Eoin M Kenny, Courtney Ford, Molly Quinn, and Mark T Keane. Explaining black-box classifiers  
656 using post-hoc explanations-by-example: The effect of explanations and error-rates in xai user  
657 studies. *Artificial Intelligence*, 294:103459, 2021.

658

659 Yohan Kim, John Sidney, Clemencia Pinilla, Alessandro Sette, and Bjoern Peters. Derivation of  
660 an amino acid similarity matrix for peptide: Mhc binding and its application as a bayesian prior.  
661 *BMC bioinformatics*, 10:1–11, 2009.

662 Kyohei Koyama, Kosuke Hashimoto, Chioko Nagao, and Kenji Mizuguchi. Attention network  
663 for predicting t-cell receptor-peptide binding can associate attention with interpretable protein  
664 structural properties. *Frontiers in Bioinformatics*, 3:1274599, 2023.

665

666 Bjørn PY Kwee, Marius Messemaek, Eric Marcus, Giacomo Oliveira, Wouter Schepers, Catherine J  
667 Wu, Jonas Teuwen, and Ton N Schumacher. Stapler: efficient learning of tcr-peptide specificity  
668 prediction from full-length tcr-peptide data. *bioRxiv*, pp. 2023–04, 2023.

669

670 Jinwoo Leem, Saulo H P de Oliveira, Konrad Krawczyk, and Charlotte M Deane. Stcrdab: the  
671 structural t-cell receptor database. *Nucleic acids research*, 46(D1):D406–D412, 2018.

672

673 Jiarui Li, Samuel J Landry, and Ramgopal R Mettu. Gpu acceleration of conformational stabil-  
674 ity computation for cd4+ t-cell epitope prediction. In *2024 IEEE International Conference on*  
*Bioinformatics and Biomedicine (BIBM)*, pp. 191–196. IEEE, 2024a.

675

676 Jiarui Li, Samuel J Landry, and Ramgopal R Mettu. Gpu acceleration for markov chain monte carlo  
677 sampling. In *4th International Conference on AIML Systems (AIMLSystems 2024)*. ACM, 2024b.

678

679 Tianxiao Li, Hongyu Guo, Filippo Grazioli, Mark Gerstein, and Martin Renqiang Min. Disentan-  
680 gled wasserstein autoencoder for t-cell receptor engineering. *Advances in Neural Information*  
*Processing Systems*, 36:73604–73632, 2023.

681

682 Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr  
683 Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European*  
*conference on computer vision*, pp. 740–755. Springer, 2014.

684

685 Valerie Lin, Melyssa Cheung, Ragul Gowthaman, Maya Eisenberg, Brian M Baker, and Brian G  
686 Pierce. Tcr3d 2.0: expanding the t cell receptor structure database with new structures, tools and  
687 interactions. *Nucleic Acids Research*, 53(D1):D604–D608, 2025.

688

689 Chen Liu, Mingchen Li, Yang Tan, Wenrui Gou, Guisheng Fan, and Bingxin Zhou. Sequence-only  
690 prediction of binding affinity changes: a robust and interpretable model for antibody engineering.  
691 *Bioinformatics*, 41(8):btaf446, 2025.

692

693 Tiing Jen Loh, Jia Jia Lim, Claerwen M Jones, Hien Thy Dao, Mai T Tran, Daniel G Baker, Nicole L  
694 La Gruta, Hugh H Reid, and Jamie Rossjohn. The molecular basis underlying t cell specificity  
695 towards citrullinated epitopes presented by hla-dr4. *Nature Communications*, 15(1):6201, 2024.

696

697 Claus Lundegaard, Kasper Lamberth, Mikkel Harndahl, Søren Buus, Ole Lund, and Morten Nielsen.  
698 Netmhcc-3.0: accurate web accessible predictions of human, mouse and monkey mhc class i affini-  
699 ties for peptides of length 8–11. *Nucleic acids research*, 36(suppl\_2):W509–W512, 2008.

700

701 Koshlan Mayer-Blackwell, Stefan Schattgen, Liel Cohen-Lavi, Jeremy C Crawford, Aisha Sou-  
702 quette, Jessica A Gaevert, Tomer Hertz, Paul G Thomas, Philip Bradley, and Andrew Fiore-  
703 Gartland. Tcr meta-clonotypes for biomarker discovery with tcrdist3 enabled identification of  
704 public, hla-restricted clusters of sars-cov-2 tcrs. *Elife*, 10:e68605, 2021.

---

702 Ramgopal R Mettu, Tysheena Charles, and Samuel J Landry. Cd4+ t-cell epitope prediction using  
703 antigen processing constraints. *Journal of immunological methods*, 432:72–81, 2016.  
704

705 Barthelemy Meynard-Piganeau, Christoph Feinauer, Martin Weigt, Aleksandra M Walczak, and  
706 Thierry Mora. Tulip: A transformer-based unsupervised language model for interacting peptides  
707 and t cell receptors that generalizes to unseen epitopes. *Proceedings of the National Academy of  
708 Sciences*, 121(24):e2316401121, 2024.

709 Alexander Myronov, Giovanni Mazzocco, Paulina Król, and Dariusz Plewczynski.  
710 Bertrand—peptide: Tcr binding prediction using bidirectional encoder representations from  
711 transformers augmented with random tcr pairing. *Bioinformatics*, 39(8):btad468, 2023.  
712

713 Jacques Neefjes, Marlieke LM Jongsma, Petra Paul, and Oddmund Bakke. Towards a systems  
714 understanding of mhc class i and mhc class ii antigen presentation. *Nature reviews immunology*,  
715 11(12):823–836, 2011.

716 Morten Nielsen, Claus Lundegaard, Peder Worning, Sanne Lise Lauemøller, Kasper Lamberth,  
717 Søren Buus, Søren Brunak, and Ole Lund. Reliable prediction of t-cell epitopes using neural  
718 networks with novel sequence representations. *Protein Science*, 12(5):1007–1017, 2003.  
719

720 Morten Nielsen, Claus Lundegaard, Thomas Blicher, Kasper Lamberth, Mikkel Harndahl, Sune  
721 Justesen, Gustav Røder, Bjoern Peters, Alessandro Sette, Ole Lund, et al. Netmhcpn, a method  
722 for quantitative predictions of peptide binding to any hla-a and-b locus protein of known sequence.  
723 *PloS one*, 2(8):e796, 2007.

724 Morten Nielsen, Massimo Andreatta, Bjoern Peters, and Søren Buus. Immunoinformatics: predict-  
725 ing peptide–mhc binding. *Annual review of biomedical data science*, 3(1):191–215, 2020.  
726

727 Bjoern Peters and Alessandro Sette. Generating quantitative models describing the sequence speci-  
728 ficity of biological processes with the stabilized matrix method. *BMC bioinformatics*, 6(1):1–9,  
729 2005.

730 Bjoern Peters, Morten Nielsen, and Alessandro Sette. T cell epitope predictions. *Annual Review of  
731 Immunology*, 38(1):123–145, 2020.  
732

733 Mansour Poorebrahim, Niloufar Mohammadkhani, Reza Mahmoudi, Monireh Gholizadeh, Elham  
734 Fakhr, and Angel Cid-Arregui. Tcr-like cars and tcr-cars targeting neoepitopes: an emerging  
735 potential. *Cancer gene therapy*, 28(6):581–589, 2021.

736 Yao Qiang, Deng Pan, Chengyin Li, Xin Li, Rhongho Jang, and Dongxiao Zhu. Attcat: Explaining  
737 transformers via attentive class activation tokens. *Advances in neural information processing  
738 systems*, 35:5052–5064, 2022.  
739

740 Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language under-  
741 standing by generative pre-training. *OpenAI*, 2018.

742 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,  
743 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual  
744 models from natural language supervision. In *International conference on machine learning*, pp.  
745 8748–8763. PmLR, 2021.

746

747 Luis A Rojas, Zachary Sethna, Kevin C Soares, Cristina Olcese, Nan Pang, Erin Patterson, Jayon  
748 Lihm, Nicholas Ceglia, Pablo Guasp, Alexander Chu, et al. Personalized rna neoantigen vaccines  
749 stimulate t cells in pancreatic cancer. *Nature*, 618(7963):144–150, 2023.

750 A Saranya and R Subhashini. A systematic review of explainable artificial intelligence models and  
751 applications: Recent developments and future trends. *Decision analytics journal*, 7:100230, 2023.  
752

753 Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh,  
754 and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based local-  
755 ization. In *Proceedings of the IEEE international conference on computer vision*, pp. 618–626,  
2017.

756 Ido Springer, Nili Tickotsky, and Yoram Louzoun. Contribution of t cell receptor alpha and beta  
757 cdr3, mhc typing, v and j genes to peptide binding prediction. *Frontiers in immunology*, 12:  
758 664514, 2021.

759 Guillaume BE Stewart-Jones, Andrew J McMichael, John I Bell, David I Stuart, and E Yvonne  
760 Jones. A structural basis for immunodominant human t cell receptor recognition. *Nature im-*  
761 *munology*, 4(7):657–663, 2003.

762 Nili Tickotsky, Tal Sagiv, Jaime Prilusky, Eric Shifrut, and Nir Friedman. Mcpas-tcr: a manu-  
763 ally curated catalogue of pathology-associated t cell receptor sequences. *Bioinformatics*, 33(18):  
764 2924–2929, 2017.

765 Sebastiaan Valkiers, Max Van Houcke, Kris Laukens, and Pieter Meysman. Cluster: a python  
766 interface for rapid clustering of large sets of cdr3 sequences with unknown antigen specificity.  
767 *Bioinformatics*, 37(24):4865–4867, 2021.

768 A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.

769 Randi Vita, Swapnil Mahajan, James A Overton, Sandeep Kumar Dhanda, Sheridan Martini, Jason R  
770 Cantrell, Daniel K Wheeler, Alessandro Sette, and Bjoern Peters. The immune epitope database  
771 (iedb): 2018 update. *Nucleic acids research*, 47(D1):D339–D343, 2019.

772 Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. Analyzing multi-head  
773 self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In *Proceedings of  
774 the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 5797–5808, 2019.

775 Anna Weber, Jannis Born, and María Rodriguez Martínez. Titan: T-cell receptor specificity predic-  
776 tion with bimodal attention networks. *Bioinformatics*, 37(Supplement\_1):i237–i244, 2021.

777 Sarah Wiegreffe and Yuval Pinter. Attention is not explanation. In *Proceedings of the 2019  
778 Conference on Empirical Methods in Natural Language Processing and the 9th International  
779 Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 11–20, 2019.

780 Junyi Wu, Bin Duan, Weitai Kang, Hao Tang, and Yan Yan. Token transformation matters: Towards  
781 faithful post-hoc explanation for vision transformer. In *Proceedings of the IEEE/CVF Conference  
782 on Computer Vision and Pattern Recognition*, pp. 10926–10935, 2024a.

783 Kevin E Wu, Kathryn Yost, Bence Daniel, Julia Belk, Yu Xia, Takeshi Egawa, Ansuman Satpathy,  
784 Howard Chang, and James Zou. Tcr-bert: learning the grammar of t-cell receptors for flexible  
785 antigen-binding analyses. In *Machine Learning in Computational Biology*, pp. 194–229. PMLR,  
786 2024b.

787 Xinbo Yang, Guobing Chen, Nan-ping Weng, and Roy A Mariuzza. Structural basis for clonal  
788 diversity of the human t-cell response to a dominant influenza virus epitope. *Journal of Biological  
789 Chemistry*, 292(45):18618–18627, 2017.

790 Hongyi Zhang, Longchao Liu, Jian Zhang, Jiahui Chen, Jianfeng Ye, Sachet Shukla, Jian Qiao,  
791 Xiaowei Zhan, Hao Chen, Catherine J Wu, et al. Investigation of antigen-specific t-cell receptor  
792 clusters in human cancers. *Clinical Cancer Research*, 26(6):1359–1371, 2020.

793 Hongyi Zhang, Xiaowei Zhan, and Bo Li. Giana allows computationally-efficient tcr clustering and  
794 multi-disease repertoire classification by isometric transformation. *Nature communications*, 12  
795 (1):4699, 2021.

796 Liyuan Zhang, Sicong Wang, Yadong Wang, and Tianyi Zhao. Hbformer: a single-stream frame-  
797 work based on hybrid attention mechanism for identification of human-virus protein–protein in-  
798 teractions. *Bioinformatics*, 40(12):btae724, 2024.

799 Yunxiang Zhao, Jijun Yu, Yixin Su, You Shu, Enhao Ma, Jing Wang, Shuyang Jiang, Congwen Wei,  
800 Dongsheng Li, Zhen Huang, et al. A unified deep framework for peptide–major histocompatibility  
801 complex–t cell receptor binding prediction. *Nature Machine Intelligence*, pp. 1–11, 2025.

802 Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep  
803 features for discriminative localization. In *Proceedings of the IEEE conference on computer  
804 vision and pattern recognition*, pp. 2921–2929, 2016.

---

810           A SUPPLEMENTARY MATERIAL  
811

812           A.1 POST-HOC EXPLANATION METHODS  
813

814           A variety of explainable AI (XAI) methods have been developed to interpret deep learning mod-  
815           els (Saranya & Subhashini, 2023). These methods fall into two broad categories: explain-by-design,  
816           which integrates interpretability into the model architecture (Dwivedi et al., 2023), and post-hoc,  
817           which analyzes model behavior after training (Kenny et al., 2021). Post-hoc approaches offer a  
818           promising avenue for interpreting TCR-pMHC models and uncovering the underlying factors driv-  
819           ing binding predictions. Several families of post-hoc methods have been proposed, including:

820           • The Class Activation Map (CAM) (e.g., CAM (Zhou et al., 2016), GradCAM (Selvaraju  
821           et al., 2017), GradCAM++ (Chattopadhyay et al., 2018))  
822           • Layer-wise Relevance Propagation (LRP) (e.g., LRP (Binder et al., 2016), Partial  
823           LRP (Voita et al., 2019), Conservative LRP (Ali et al., 2022), AttnLRP (Achtibat et al.,  
824           2024))  
825           • Attention-based methods (e.g., Raw Attention (Wiegreffe & Pinter, 2019), Attention Roll-  
826           out (Abnar & Zuidema, 2020), AttCAT (Qiang et al., 2022))  
827           • Model-specific hybrid methods (e.g., TokenTM (Wu et al., 2024a), GAE (Chefer et al.,  
828           2021))  
829

830           These techniques have been successfully applied to TCR-pMHC models. For example, TEPCAM  
831           uses CAM to interpret a CNN-based predictor (Chen et al., 2024), while TCR-BERT relies on at-  
832           tention weight analysis for interpretability (Wu et al., 2024b). These efforts have revealed struc-  
833           tural determinants of TCR-pMHC binding. However, existing post-hoc methods primarily support  
834           encoder-only or co-attention mechanisms (Chefer et al., 2021), limiting their applicability to modern  
835           encoder-decoder models, which consists of cross-attention. This poses a major barrier to understand-  
836           ing how such models capture TCR-pMHC interactions.  
837

838           A.2 CLASS ACTIVATION MAPS  
839

840           Class Activation Map (CAM)-based methods have achieved significant success in explaining Con-  
841           volutional Neural Networks (CNNs) by generating class-discriminative localization maps. Grad-  
842           CAM (Selvaraju et al., 2017), one of the most effective CAM methods, leverages the gradient of the  
843           class score  $L^c$  with respect to the feature maps  $F_d$  from the last convolutional layer. These gradients  
844           are used to compute importance weights for each feature map channel, enabling spatial localiza-  
845           tion of the regions most relevant for class  $c$ . The importance weight  $w_d^c$  for feature map  $F_d$  is computed  
846           as:

847

$$w_d^c = \mathbb{E} \left( \frac{\partial L^c}{\partial F_d} \right),$$

848           where  $\mathbb{E}$  denotes global average and  $w_d^c$  represents the global average pooled gradient for feature  
849           map  $F_d$ . The final CAM is then computed as a weighted sum over channels, followed by a ReLU  
850           activation:

851

$$\text{GradCAM}^c = \text{ReLU} \left( \sum_d w_d^c F_d \right).$$

852           The resulting heatmap is upsampled to the input resolution to highlight input regions most relevant  
853           to the prediction for class  $c$ .  
854

855           A.2.1 ATTENTION ROLLOUT  
856

857           CAM-based approaches are primarily designed for CNNs. To interpret transformer-based models,  
858           Attention Rollout was proposed by [Abnar & Zuidema \(2020\)](#), which estimates the flow of attention  
859           across layers. This method computes how information propagates through the self-attention mecha-  
860           nism across layers. Given the raw attention weights  $W_l^A$  for layer  $l$ , the augmented attention matrix  
861           is defined as  
862

863

$$A_l = \frac{1}{2}(W_l^A + I),$$

864 where  $I$  is the identity matrix, modeling the residual connection. The cumulative attention, or  
865 rollout, is then computed recursively:  
866

$$867 \quad R_l = \begin{cases} A_l R_{l-1}, & \text{if } l > 0 \\ A_l, & \text{if } l = 0 \end{cases},$$

869 capturing the total attention contribution from input tokens through layer  $l$ .  
870

### 871 A.3 TCR-pMHC BINDING PREDICTION

873 T cells are important component of our immune system, which can be mainly categorized in two  
874 CD8+ and CD4+ T cells. CD8+ T cells are initiated through the Major Histocompatibility Complex I (MHC I)  
875 pathway, while CD4+ T cells are initiated through the MHC II pathway. Epitope prediction for CD8+ T cells has had remarkable success, while the mechanisms of CD4+ T cell  
876 response are less understood. T cell immune response can be viewed as consisting of two stages  
877 of recognition. In the first stage, an antigen is taken up by antigen-presenting cells (APCs), where  
878 it undergoes joint processing (i.e., cleavage) and binding to Major Histocompatibility Complex II (MHC II)  
879 molecules. Peptide-MHC complexes are then presented on the APC cell surface (Davis & Bjorkman, 1988; Neefjes et al., 2011). In the second stage, T cell receptors (TCRs) on T cells  
880 “recognize” pMHC complexes and a T cell response is initiated. TCR recognition is mediated by  
881 its  $\alpha$  and  $\beta$  domains, which consist of variable (V), joining (J), constant (C), and, in the  $\beta$  chain,  
882 diversity (D) regions (Bosselut, 2019).  
883

885 Accurate prediction of T cell responses requires a comprehensive understanding of both of these  
886 stages (Peters et al., 2020; Nielsen et al., 2020). Early efforts in the area of computational epitope  
887 prediction focused on characterizing peptide-MHC II binding using allele-specific machine learning  
888 models (Nielsen et al., 2020) with tools such as SMM (Peters & Sette, 2005; Kim et al., 2009),  
889 NetMHC (Lundsgaard et al., 2008; Nielsen et al., 2003), NetMHCpan (Hoof et al., 2009; Nielsen  
890 et al., 2007), and NetMHCcons (Karosiene et al., 2012). More recent work has focused on modeling  
891 antigen processing computationally with the Antigen Processing Likelihood (APL) algorithm (Mettu  
892 et al., 2016; Bhattacharya et al., 2023; Li et al., 2024a;b; Charles et al., 2022), which seeks to model  
893 the contributions of antigen structure on which peptides are made available for MHC II binding.  
894

895 Accurately predicting TCR-pMHC binding remains critical for advancing quantitative immunology and adaptive immunity research (Hudson et al., 2023). For this stage of prediction, both  
896 unsupervised and supervised methods have been developed (Hudson et al., 2023; 2024). Unsupervised methods process cluster TCR sequencing datasets through dimensionality reduction and  
897 clustering (Dash et al., 2017; Glanville et al., 2017) through a carefully chosen similarity metric  
898 (e.g., TCRdist3 (Mayer-Blackwell et al., 2021)). These methods cluster TCRs by analyzing their  
899 complementarity-determining regions (CDRs) using only TCR sequence data, without requiring  
900 binding labels or epitope information (e.g., GIANA (Zhang et al., 2021), ClusTCR (Valkiers et al.,  
901 2021), GLIPH2 (Huang et al., 2020) iSMART (Zhang et al., 2020)). The resulting cluster labels  
902 serve as the output for each input TCR sequence (Hudson et al., 2024) and are typically analyzed by  
903 practitioners to guide and supplement experimental methods. In contrast, supervised machine learning  
904 techniques make use of large amounts of TCR-pMHC data for training (Hudson et al., 2023)  
905 from databases such as VDJdb (Bagaev et al., 2020), McPAS-TCR (Tickotsky et al., 2017) and the  
906 IEDB (Vita et al., 2019). Supervised approaches (e.g. TITAN (Weber et al., 2021), STAPLER (Kwee  
907 et al., 2023), ERGO2 (Springer et al., 2021), MixTCRpred (Croce et al., 2024), NetTCR2.2 (Jensen  
908 & Nielsen, 2023), TULIP (Meynard-Piganeau et al., 2024)) use a variety of deep learning models  
909 providing reasonable performance and generalization capability.  
910

#### 911 A.3.1 TCR-pMHC BINDING PROBLEM FORMATION

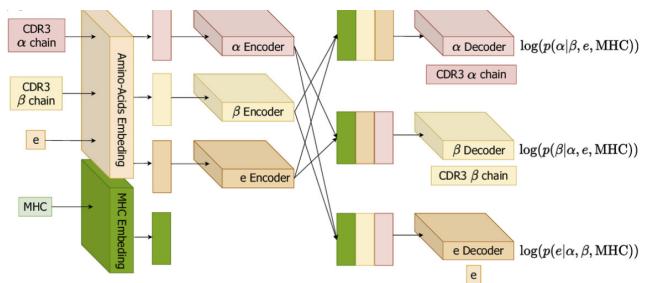
912 The TCR-pMHC binding prediction problem can be formulated as a classification task: given the  
913 TCR alpha ( $\alpha$ ) and beta ( $\beta$ ) chains, an epitope  $e$ , and an MHC molecule  $m$ , the model predicts  
914 whether the pair binds (binder) or does not bind (non-binder). The TCR chains and the epitope are  
915 proteins or peptides, typically represented as amino acid sequences. Formally, we define amino acid  
916 units as  $a \in \mathbb{A}$ , where  $\mathbb{A}$  is the set of amino acid characters. For a single TCR-pMHC binding case,  
917  $\alpha = [a_i^\alpha]_{i=1}^{N_\alpha}$ ,  $\beta = [a_i^\beta]_{i=1}^{N_\beta}$ , and  $e = [a_i^e]_{i=1}^{N_e}$ , with  $N_\alpha, N_\beta, N_e \in \mathbb{Z}^+$  representing the sequence  
lengths. The MHC allele type is denoted by  $m \in M$ , where  $M$  is the set of all MHC alleles. The

918 pMHC-TCR binding classification is formulated as a conditional probability:  $p_{\text{bind}} = P(e|\alpha, \beta, m)$ .  
 919 If  $p_{\text{bind}} > t$ , where  $t \in [0, 1]$ , the case is classified as positive, otherwise negative.  
 920

#### 921 A.4 TCR-PMHC PREDICTION TRANSFORMER MODELS

923 Transformers (Vaswani, 2017), as a successful deep learning models in different areas, have a series  
 924 of variants such as Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al.,  
 925 2019) and Generative Pre-training Transformers (GPT) (Radford et al., 2018). These models support  
 926 multi-sequence inputs and excel in modeling interactions, are well-suited for this task. Because  
 927 TCR-pMHC interactions are determined by interactions among the TCR  $\alpha$  and  $\beta$  chains, epitope,  
 928 and MHC, several state-of-the-art models, such as TULIP (Meynard-Piganeau et al., 2024) and  
 929 cross-TCR-interpreter (Koyama et al., 2023), adopt encoder-decoder transformer architectures to  
 930 learn these complex relationships.

931 **TULIP:** TULIP is a transformer-  
 932 based model with an encoder-  
 933 decoder architecture designed for  
 934 TCR-pMHC binding prediction.  
 935 It operates through three parallel  
 936 modality processing pipelines, pro-  
 937 cessing CDR3a, CDR3b, and epitope  
 938 sequences separately (Meynard-  
 939 Piganeau et al., 2024). The encoders  
 940 transform the input sequences  
 941 into feature representations, while  
 942 the decoders model interactions  
 943 across different sequences (Devlin,  
 944 2018; Vaswani, 2017). As an auto-  
 945 regressive generative model, TULIP

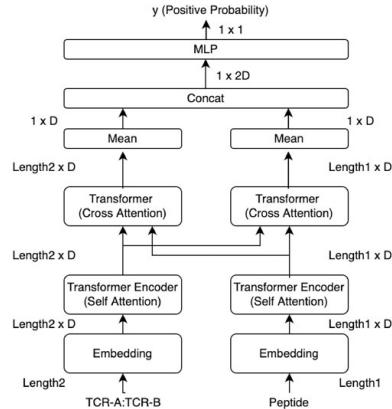


946 Figure 6: The architecture figure of TULIP  
 947 model (Meynard-Piganeau et al., 2024).

948 computes the conditional probability distribution of sequences (e.g., epitope) given others (e.g.,  
 949 CDR3a, CDR3b, and MHC) during training (Meynard-Piganeau et al., 2024). **For evaluation,**  
 950 **TULIP retains only the epitope stream to produce the binding score. In this setting, the peptide**  
 951 **features serve as the query, while the CDR3a and CDR3b features are used as the keys and values**  
 952 **in the cross-attention module.** To compute gradients for TULIP, we design an amino-acid-wise loss  
 953 function. The ground truth is derived from the TCR alpha, TCR beta, and epitope sequences. These  
 954 sequences are first one-hot encoded, and the model’s predicted probabilities are compared against  
 955 them using a negative log-likelihood (NLL) loss. This formulation allows us to attribute importance  
 956 scores at the amino acid level based on how well the model reconstructs each residue.

#### 957 **CrossTCRInterpreter:**

958 CrossTCRInterpreter is an encoder-  
 959 decoder transformer for TCR-pMHC  
 960 binding prediction (Koyama et al.,  
 961 2023). It takes the CDR regions of  
 962 the alpha and beta chains, along with  
 963 the peptide sequence, as inputs. The  
 964 CDR alpha and beta chains are  
 965 concatenated using a colon (:) to form  
 966 the TCR input. The TCR and peptide  
 967 sequences are then independently  
 968 encoded by an encoder module.  
 969 Subsequently, cross-attention is  
 970 employed to model the interaction  
 971 between the two inputs and predict  
 972 whether the pair represents a binder  
 973 or a non-binder. We apply a binary  
 974 classification loss to extract the  
 975 model gradients.



976 Figure 7: The architecture of CrossTCRInterpreter  
 977 model (Koyama et al., 2023).

972                    **BERTrand:** BERTrand is  
 973                    an encoder-only transformer  
 974                    model (Myronov et al., 2023).  
 975                    It takes the TCR beta chain and pep-  
 976                    tide sequence as inputs. These two  
 977                    sequences are concatenated using  
 978                    a <SEP> token and are processed  
 979                    jointly by a transformer encoder  
 980                    as an integrated input. Similar to  
 981                    CrossTCRInterpreter, BERTrand is  
 982                    a classification model designed to  
 983                    predict whether the TCR-pMHC pair  
 984                    is a binder or a non-binder. We apply  
 985                    a binary classification loss to obtain  
 986                    the model gradients.

## 988 A.5 PERTURBATION EXPERIMENTS

990                    We evaluated the robustness of interpretability methods using perturbation-based metrics across  
 991                    varying values of  $k$ . Figure 9 presents the comparison results for both AOPC and LOdds across all  
 992                    chains.

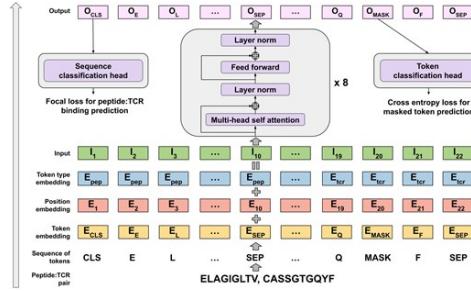


Figure 8: The architecture figure of BERTrand model (Myronov et al., 2023).

## 988 A.5 PERTURBATION EXPERIMENTS

990                    We evaluated the robustness of interpretability methods using perturbation-based metrics across  
 991                    varying values of  $k$ . Figure 9 presents the comparison results for both AOPC and LOdds across all  
 992                    chains.

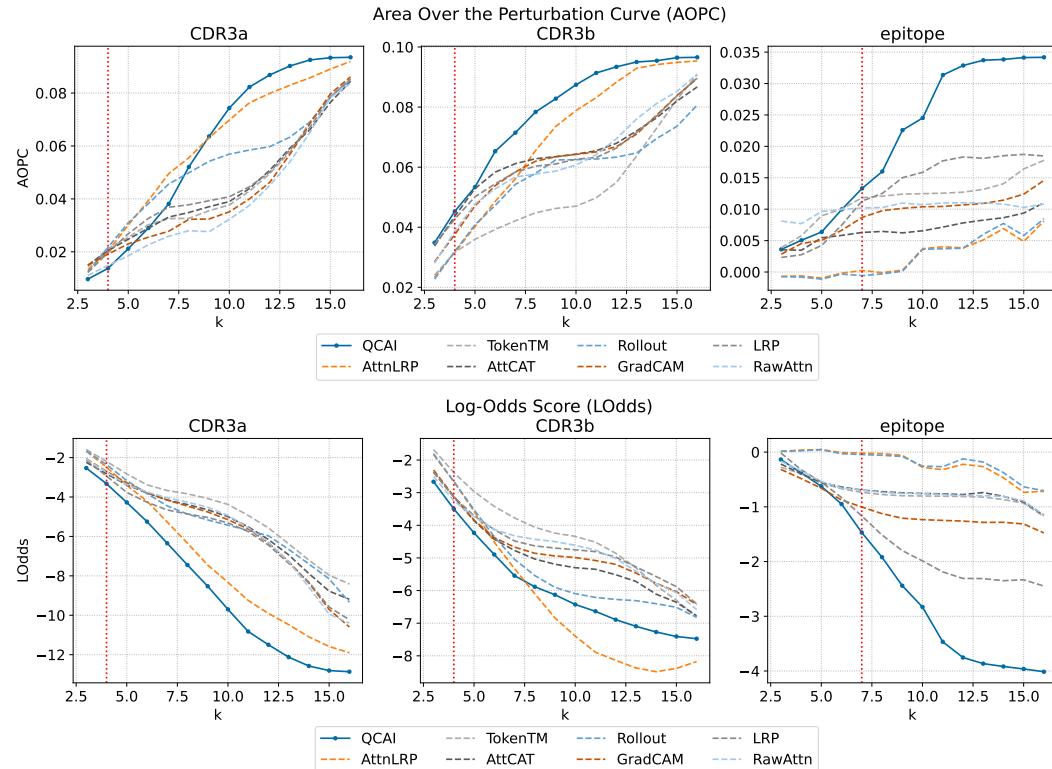


Figure 9: Comparison of Area Over the Perturbation Curve (AOPC) and Comparison of Log-Odds Score (LOdds) across different values of  $k$  for all chains.

1024                    We also conducted with an integrated dataset that includes data from VDJdb, IEDB and McPAS-  
 1025                    TCR. For both AOPCs and LOdds, the thresholds for peptide, CDR3a, and CDR3b are 7, 4, and 4  
 respectively.

Chain	Method	AOPCs		LOdds	
		TCR-XAI	Integrated	TCR-XAI	Integrated
peptide	QCAI	0.014	<b>0.036</b>	-1.62	<b>-0.84</b>
peptide	TokenTM	0.013	0.026	-0.77	0.09
peptide	AttnLRP	0.012	0.026	-0.42	-0.50
CDR3a	QCAI	0.014	0.020	-3.50	<b>-2.63</b>
CDR3a	TokenTM	0.021	0.020	-2.43	-2.35
CDR3a	AttnLRP	0.020	0.020	-2.72	-2.44
CDR3b	QCAI	0.048	<b>0.027</b>	-3.61	<b>-3.08</b>
CDR3b	TokenTM	0.033	0.025	-2.53	-2.78
CDR3b	AttnLRP	0.034	0.024	-2.82	-2.88

Table 2: AOPCs and LOdds comparison on TCR-XAI and Integrated datasets.

#### A.6 MAXIMUM VS. AVERAGE FOR AGGREGATION

High attention weights indicate meaningful interactions and so we used maximum across different cross-attention layers to retain all activated signals. Ablation studies in the table below show that max generally outperforms average, with small exceptions on peptide BRHR and CDR LOdds.

Chain	Mix	ROC-AUC(3.4)	BRHR.25	AOPCs	LOdds
peptide	Max.	0.60	74.3	<b>0.014</b>	<b>-1.52</b>
peptide	Avg.	0.60	<b>76.7</b>	0.013	-1.51
CDR3a	Max.	<b>0.55</b>	<b>79.1</b>	<b>0.014</b>	-3.37
CDR3a	Avg.	0.50	72.6	0.013	<b>-3.51</b>
CDR3b	Max.	<b>0.55</b>	<b>79.3</b>	<b>0.046</b>	-3.54
CDR3b	Avg.	0.54	75.3	0.045	<b>-3.63</b>

Table 3: Maximum vs. Average for aggregation comparison across chains for ROC-AUC (3.4), BRHR, AOPCs, and LOdds.

#### A.7 ROC CURVES OF RESIDUE-LEVEL IMPORTANCE SCORES FOR BINDING REGION IDENTIFICATION

We compared QCAI with other methods using ROC curves. We mainly set the distance threshold at 3, 3.4, 4, and 5 Å. 3.4 Å is chosen because it corresponds to the van der Waals diameter between two carbon atoms, implying that residues within this distance are considered to be in contact.

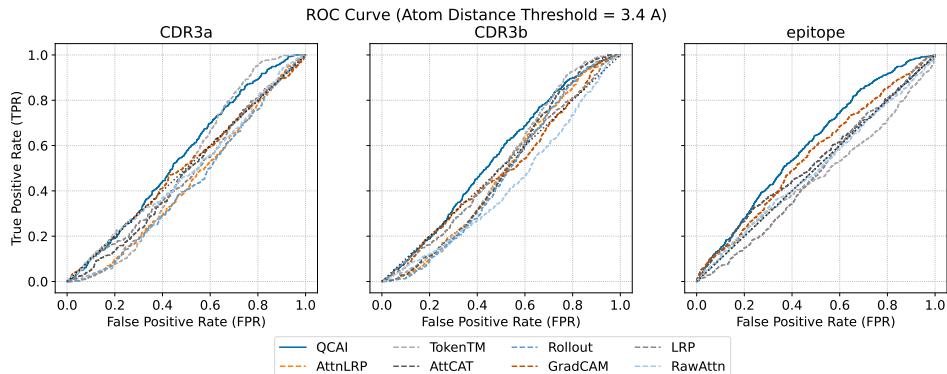


Figure 10: ROC curve comparison of the alpha, beta, and epitope chains between QCAI and other post-hoc methods. The distance threshold is set to 3.4 Å.

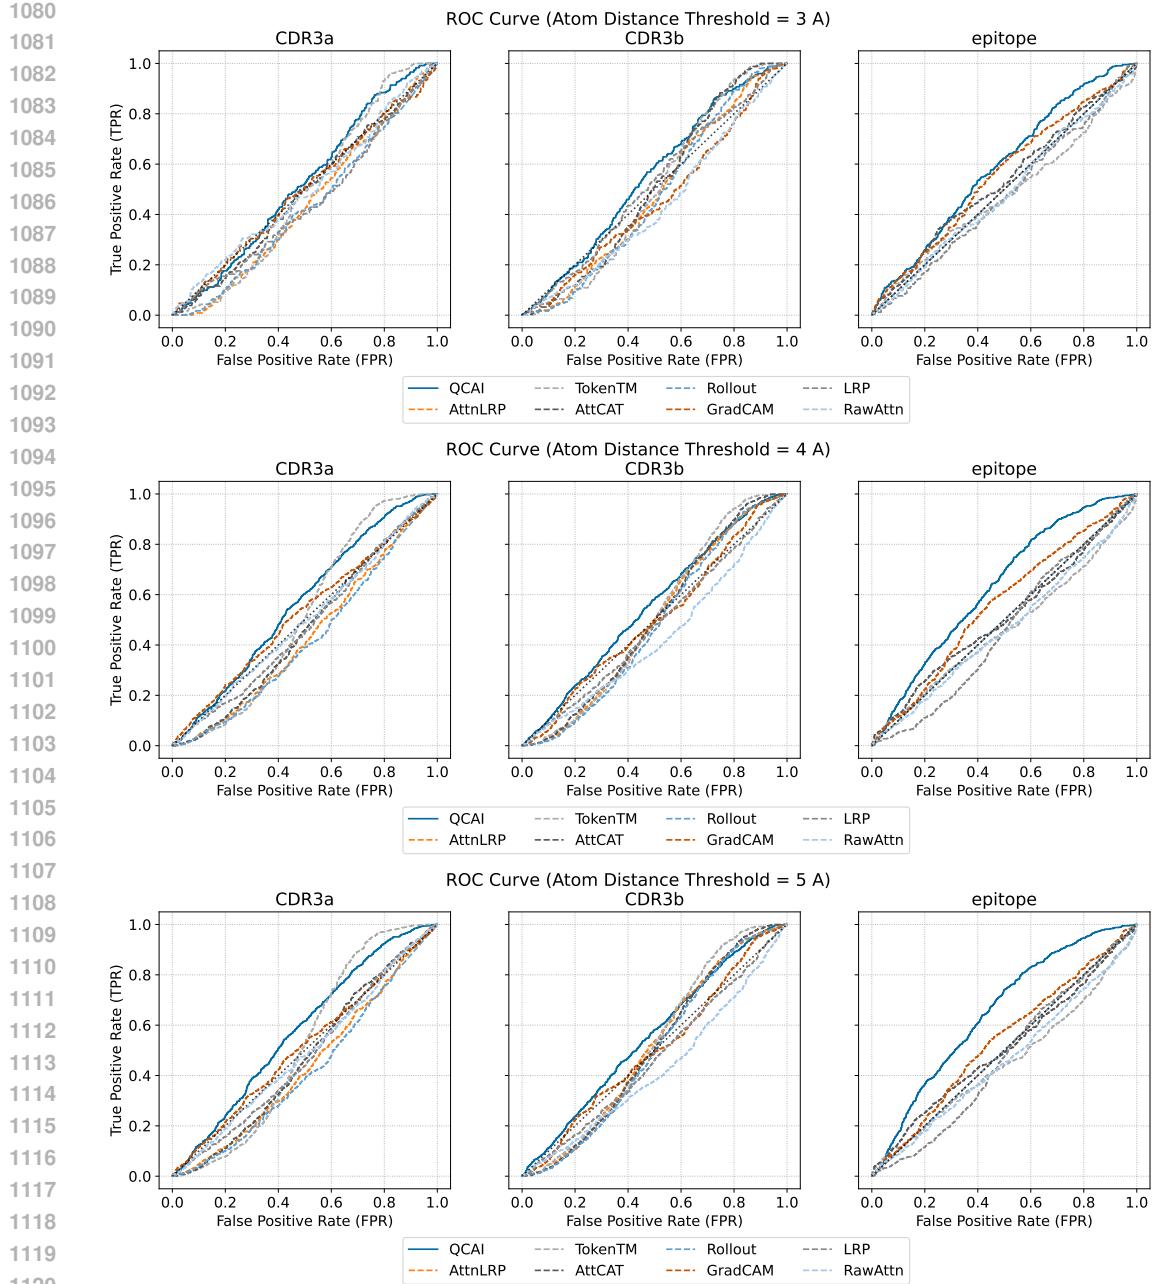


Figure 11: ROC curve comparison of the alpha, beta, and epitope chains between QCAI and other post-hoc methods. The distance thresholds are set to 3, 4, and 5 Å.

#### A.8 BINDING REGION HIT RATE

We compare the Binding Region Hit Rate (BRHR) across the TCR  $\alpha$  and  $\beta$  chains as well as the epitope region for different explanation methods. Here,  $HR.t$  denotes the hit rate calculated based on the top  $t$  percentile of importance scores.

1134  
1135  
1136  
1137  
1138  
1139  
1140  
1141  
1142  
1143  
1144  
1145  
1146  
1147  
1148  
1149  
1150  
1151  
1152  
1153  
1154  
1155  
1156  
1157  
1158  
1159  
1160  
1161  
1162  
1163  
1164  
1165  
1166  
1167  
1168  
1169  
1170  
1171  
1172  
1173  
1174  
1175  
1176  
1177  
1178  
1179  
1180  
1181  
1182  
1183  
1184  
1185  
1186  
1187

Chain	Method	HR.25	HR.30	HR.40	HR.50
epitope	<b>QCAI (Ours)</b>	<b>74.3(±24.5)%</b>	<b>72.7(±24.6)%</b>	<b>66.4(±19.8)%</b>	<b>55.3(±15.7)%</b>
	AttLRP (Achitibat et al., 2024)	58.4(±29.2)%	60.6(±25.9)%	60.4(±19.8)%	53.7(±15.8)%
	TokenTM (Wu et al., 2024a)	68.5(±29.6)%	66.4(±29.4)%	56.3(±25.9)%	44.4(±20.7)%
	AttCAT (Qiang et al., 2022)	69.1(±30.8)%	66.8(±30.0)%	57.9(±23.9)%	46.5(±18.4)%
	Rollout (Abnar & Zuidema, 2020)	58.4(±29.2)%	60.6(±25.9)%	60.4(±19.8)%	53.7(±15.8)%
	GradCAM (Selvaraju et al., 2017)	72.2(±29.5)%	70.9(±29.3)%	61.9(±24.9)%	50.6(±19.9)%
	LRP (Binder et al., 2016)	60.8(±31.4)%	58.0(±29.8)%	51.7(±21.4)%	42.2(±17.3)%
	RawAttn (Wiegreffe & Pinter, 2019)	68.6(±27.7)%	65.2(±26.0)%	53.7(±23.5)%	42.8(±20.0)%
CDR3a	<b>QCAI (Ours)</b>	<b>79.1(±20.1)%</b>	<b>74.9(±19.4)%</b>	<b>66.7(±16.7)%</b>	<b>55.2(±16.6)%</b>
	AttLRP (Achitibat et al., 2024)	68.5(±24.5)%	63.6(±25.1)%	54.6(±20.9)%	43.0(±19.4)%
	TokenTM (Wu et al., 2024a)	64.4(±30.8)%	60.8(±30.2)%	57.1(±27.1)%	50.1(±20.9)%
	AttCAT (Qiang et al., 2022)	62.7(±25.0)%	60.4(±24.9)%	54.9(±23.4)%	45.0(±19.9)%
	Rollout (Abnar & Zuidema, 2020)	66.6(±25.2)%	61.5(±24.7)%	51.5(±20.7)%	41.1(±18.8)%
	GradCAM (Selvaraju et al., 2017)	66.7(±26.7)%	62.7(±25.5)%	56.1(±20.1)%	46.5(±17.0)%
	LRP (Binder et al., 2016)	66.3(±27.1)%	61.7(±26.6)%	54.9(±21.8)%	46.0(±19.0)%
	RawAttn (Wiegreffe & Pinter, 2019)	65.8(±27.2)%	60.8(±25.3)%	53.1(±22.5)%	44.0(±17.1)%
CDR3b	<b>QCAI (Ours)</b>	<b>79.3(±19.0)%</b>	<b>76.7(±18.9)%</b>	<b>67.7(±16.2)%</b>	<b>56.5(±14.9)%</b>
	AttLRP (Achitibat et al., 2024)	72.6(±25.3)%	66.1(±23.4)%	57.1(±21.8)%	49.5(±17.8)%
	TokenTM (Wu et al., 2024a)	69.5(±27.5)%	66.0(±26.7)%	60.8(±22.0)%	51.6(±18.3)%
	AttCAT (Qiang et al., 2022)	66.9(±25.9)%	64.9(±24.4)%	57.9(±23.2)%	49.3(±19.4)%
	Rollout (Abnar & Zuidema, 2020)	70.4(±25.6)%	64.4(±23.3)%	55.5(±21.5)%	48.3(±18.0)%
	GradCAM (Selvaraju et al., 2017)	71.7(±26.8)%	67.1(±27.1)%	61.0(±24.3)%	48.6(±19.1)%
	LRP (Binder et al., 2016)	61.8(±26.3)%	58.8(±23.1)%	54.6(±20.6)%	45.1(±18.5)%
	RawAttn (Wiegreffe & Pinter, 2019)	69.3(±24.3)%	65.0(±22.2)%	55.9(±19.8)%	44.0(±17.8)%

Table 4: The Binding Region Hit Rate comparison among TCR alpha and beta chains and epitope between various methods. The  $HR.t$  denotes the hit rate computed based on the top  $t$  percentile.

1188 To consider performance relative to training set similarity we consider the change in BRHR of  
 1189 samples indexed by the Levenshtein distance each input modality to the TULIP training dataset (this  
 1190 approach is also used in the original TULIP paper). In the Table A.8 each cell represents the BRHR  
 1191 of all samples with the minimum Levenshtein distance to the sequences of TULIP training dataset  
 1192 smaller than the given threshold. As a sample’s distance between the sequences of TCR-XAI and the  
 1193 sequences of training dataset increases we find that the BRHR score of QCAI on TULIP decreases  
 1194 slightly but remains reliable with the BRHR drop within 0.05, which is small relative to the BRHR  
 1195 difference between QCAI and other methods. This shows that QCAI’s performance is preserved  
 1196 even as samples differ from the training set.

Levenshtein Distance ( $d$ )	$1 > d$	$2 > d$	$3 > d$	$4 > d$	$5 > d$	$6 > d$	$7 \leq d$
Peptide	.77(.23)	.77(.24)	.80(.23)	.77(.24)	.75(.25)	.74(.25)	.76(.25)
CDR3a	.83(.17)	.81(.18)	.84(.19)	.82(.19)	.79(.20)	.79(.20)	.79(.20)
CDR3b	.83(.20)	.80(.21)	.77(.21)	.78(.21)	.78(.20)	.78(.20)	.78(.20)

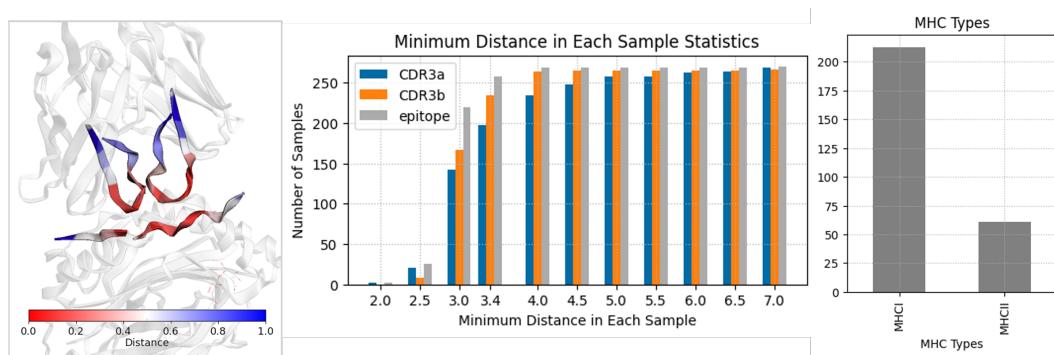
1203 Table 5: The change in BRHR for samples indexed by their Levenshtein distance.

1204 To examine how model confidence and prediction outcomes affect BRHR, we further compare  
 1205 BRHR on positive and negative samples for both TULIP and Cross-TCR-Interpreter. Because  
 1206 TULIP provides only relative binding likelihood scores, we perform QCAI analysis on Cross-TCR-  
 1207 Interpreter separately for its predicted positive and negative samples. For TULIP, we additionally  
 1208 set a manual threshold by treating the top 50% scoring pairs as positive and the remaining 50% as  
 1209 negative. The BRHR results shown in Table A.8 indicate that negative samples decreased the quality  
 1210 of explanation comparing to the positive samples.

Samples	BRHR	Cross-TCR-Interpreter	Cross-TCR-Interpreter	TULIP	TULIP
		Positive	Negative	Positive	Negative
Peptide	.61		.55	.75	.74
CDR3a	.41		.46	.79	.79
CDR3b	.87		.87	.81	.79

1219 Table 6: The BRHR of predicted positive and negative samples.

## A.9 TCR-XAI BENCHMARK



1225 Figure 12: In the example (8TRQ) from the TCR-XAI benchmark, the peptide, CDR3a, and CDR3b  
 1226 regions are highlighted based on their residue-level distances to the nearest interacting residues.  
 1227 Additionally, we report statistics for the minimum distance in each sample and MHC distribution.

1228  
 1229  
 1230  
 1231  
 1232  
 1233  
 1234  
 1235  
 1236  
 1237  
 1238  
 1239  
 1240  
 1241 We have compiled 274 samples from the STCRDab (Leem et al., 2018) and TCR3d 2.0 (Lin et al.,  
 1242 2025) datasets. Only samples with fully provided CDR3 regions and peptide sequences were se-  
 1243 lected. Among them, 213 (77.7%) are MHC-I and 61 (22.3%) are MHC-II complexes. For each

---

1242 sample, we computed the distance from each residue in the CDR3 regions to the nearest atom in  
 1243 the peptide, and vice versa from the peptide residues to the CDR3 regions. The resulting dataset in-  
 1244 cludes both the CDR3 and peptide sequences along with their corresponding residue-level distances.  
 1245 Since the model lacks structural input, we allow a one-residue positional tolerance to account for  
 1246 minor attention shifts. To this end, we smooth each method’s output importance scores by convolv-  
 1247 ing them with the kernel [1/3, 1/3, 1/3] prior to evaluation. The detailed information can be found  
 1248 in Table 9. **Compared with the TULIP training dataset, there are 176 distinct epitopes, and none**  
 1249 **appears in more than 3.3% (9) of the samples.**

1250 **A.10 COMPUTATIONAL EFFICIENCY OF QCAI**

1251 We evaluate QCAI efficiency based on datasets including VDJdb, IEDB, McPAS-TCR, and TCR-  
 1252 XAI. All evaluations are conducted on CPU (32 E5 cores). QCAI involves pseudo-inverse operations  
 1253 making it more computationally expensive than alternative approaches, but it is still relatively ef-  
 1254 ficient on a per sample basis. For example benchmark sets with thousands of test samples would  
 1255 need on the order of seconds for QCAI evaluation - this is far smaller than what would be needed by  
 1256 practitioners.

Method	TCR-XAI	McPAS-TCR	VDJdb	IEDB
QCAI	2.19 ms	2.18 ms	1.30 ms	1.90 ms
TokenTM	0.11 ms	0.15 ms	0.05 ms	0.11 ms
AttnLRP	0.04 ms	0.02 ms	0.02 ms	0.04 ms

1257 Table 7: Milliseconds per sample for each method across different datasets.

1258 **A.11 ABLATION STUDY: QCAI ON CROSS- VS. SELF-ATTENTION**

1259 To investigate whether QCAI applied to cross-attention or self-attention contributes more to the  
 1260 final explanation, we compare QCAI applied only to cross-attention, only to self-attention, and  
 1261 to both, using perturbation experiments. Applying QCAI solely to self-attention is equivalent to  
 1262 Rollout. As shown in Figure A.10 and Table A.10, the performance of QCAI on cross-attention  
 1263 alone is comparable to applying it to both cross- and self-attention, and both outperform Rollout.  
 1264 These results indicate that cross-attention is the main contributor to the final explanation and plays  
 1265 a significant role in cross-attention incorporated transformers.

	CDR3a <sub>k=4</sub>		CDR3b <sub>k=4</sub>		Peptide <sub>k=7</sub>	
	LOdds	AOPC	LOdds	AOPC	LOdds	AOPC
QCAI	-3.328	0.014	-3.498	0.045	-1.470	0.013
QCAI (Cross-Attention)	-3.728	0.017	-3.511	0.048	-1.417	0.012
Rollout (Self-Attention)	-2.356	0.022	-2.653	0.032	-0.044	-0.001
AttnLRP	-2.481	0.020	-2.662	0.032	-0.017	0.000
TokenTM	-2.195	0.021	-2.383	0.032	-0.736	0.012

1266 Table 8: **Comparison of AOPC and LOdds for QCAI applied to cross-attention only, self-attention**  
 1267 **only (Rollout), and both.**

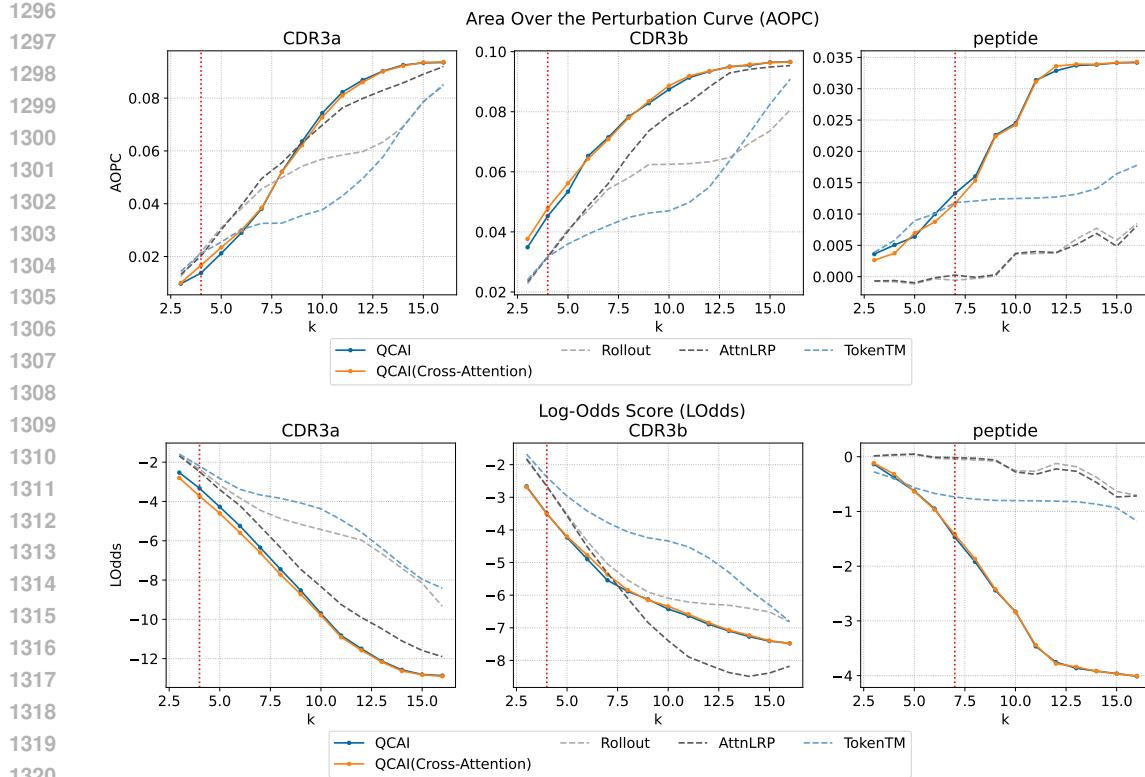


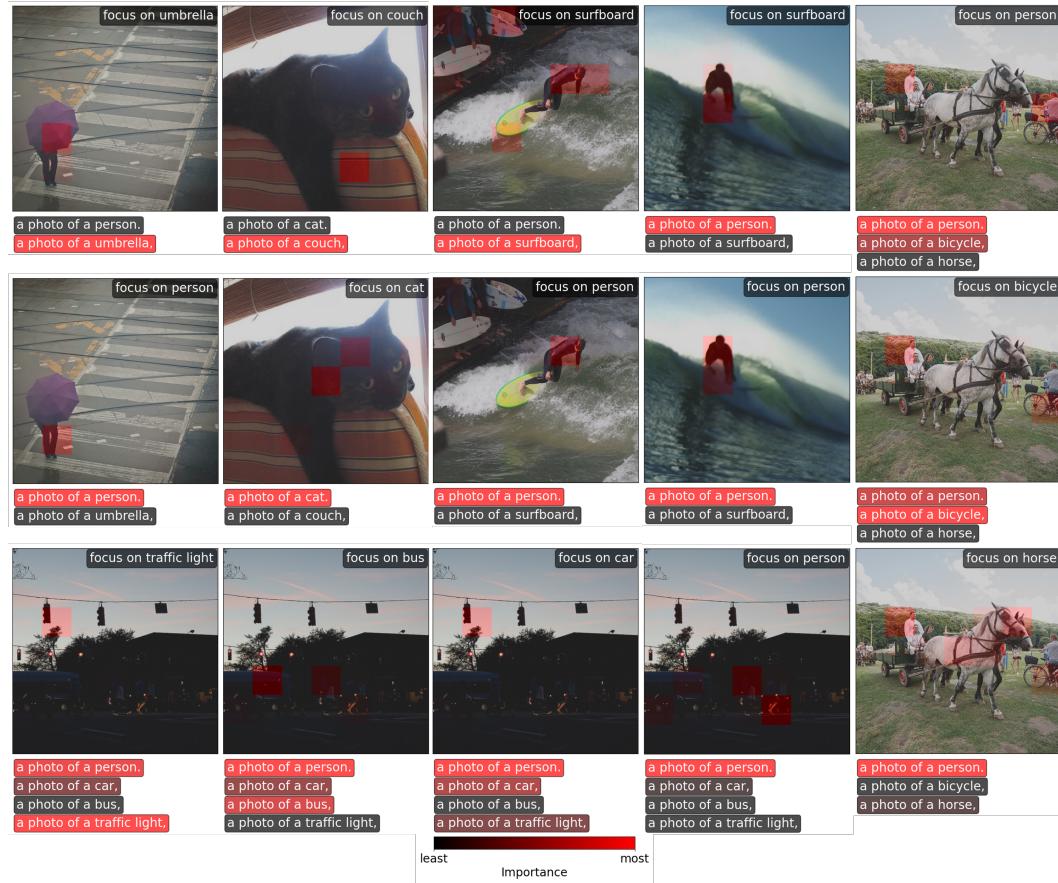
Figure 13: Comparison of AOPC and LOdds for QCAI applied to cross-attention only, self-attention only (Rollout), and both.

## A.12 APPLICATION OF VISION-LANGUAGE MODELS

Since QCAI can be applied broadly to cross-attention modules, we illustrate its use in a vision-language model (VLM). For this case study, we employ CLIP Radford et al. (2021), a widely used vision foundation model. CLIP provides separate vision and text encoders with aligned features, so we added a cross-attention layer to fuse image features (as key and value) with text features (as query). We use a subset of the MS-COCO dataset Lin et al. (2014), containing 73,000 images for multi-label classification. The dataset is split 9:1, with 65,700 training samples and 7,300 test samples. The input consists of an image–text pair, where the text is generated following the CLIP paper’s recommendation as “a photo of a ...” with the corresponding labels (e.g., “a photo of a cat”, “a photo of a couch”) Radford et al. (2018). Features extracted via cross-attention are used for label prediction. After 100 epochs of training, the model achieves 94.08% accuracy and 0.9997 ROC-AUC on the test set.

QCAI is then applied to analyze the model. Since each image-text pair has multiple labels, gradients and QCAI are computed for one label at a time. Figure A.10 presents case studies on both training and test samples, demonstrating that QCAI can identify interactions between the two input modalities in cross-attention and highlight the relative importance of the image and text for a given classification label.

1350  
1351  
1352  
1353  
1354  
1355  
1356  
1357  
1358  
1359  
1360  
1361



1390  
1391  
1392 Figure 14: Example of QCAI explaining CLIP with cross-attention.  
1393  
1394  
1395  
1396  
1397  
1398  
1399  
1400  
1401  
1402  
1403

---

1404 A.13 TCR-XAI BENCHMARK SAMPLES

1405

PDB	MHC	Peptide	CDRA3	CDRB3
8TRQ	MHCII	GVYATSSAVRLR	ALGDHSGSWQLI	ASSLRGANSDYT
4OZI	MHCII	QPFPQPELPYP	LVGDGGSFSGGYNKLI	SAGVGGQETQY
2AK4	MHCI	LPEPLPQGQLTAY	ALSGFYNTDKLI	ASPGLAGEYEQY
5EU6	MHCI	YLEPGPVTV	AVLSSGGSNYKLT	ASSFIGGTDTQY
7PBE	MHCI	YLQPRTFLL	VVNINTDKLI	ASSSANSGEFL
4Z7W	MHCII	PSGEGSFQPSQENPQ	AVGETGANNLF	ASSEARRYNEQF
6V18	MHCII	GGYRAPAKAAAT	ALSDSGSFNKL	ASSLDWGGQNTLY
8E08	MHCI	LPFDKATIM	AADGGAGSYQLT	SAGPTSGRTDTQY
5WKH	MHCI	GTSGSPINR	GLGDAGNMLT	ASSLGQGLLYGYT
7T2B	MHCII	ATGLAWEWWRTVYE	ATDKGGATNKLI	ASSQGGGEQY
7SG1	MHCII	QPFPQPELPYGGGS	LVGGLARDMR	SVALGSDTGEFL
3W0W	MHCI	RFPLTFGWCF	GTYNQGGKLI	ASSGASHEQY
5W1V	MHCI	VMAPRTLIL	AGQPLGGSNYKLT	ASSANPGDSSNEKLF
6AVF	MHCI	APRGPHGGAASGL	LVGEILDNFNKFY	ASSQRQEGLDTQY
5NHT	MHCI	ELAGIGILTV	AVGGGADGLT	ASSQGLAGAGELF
3TPU	MHCI	FLSPFWFDI	AVSAKGTGSKLS	ASSDAPGQLY
2P5E	MHCI	SLLLMWIITQC	AVRPLLDGTYIPT	ASSYLGNTGELLF
4P2O	MHCII	PADPLAFFSSAIKGGGGSLV	AALRATGGNNKLT	ASSLNWSQDTQY
7NME	MHCI	QLPRLFPPLL	AEPSGNTGKLI	ASSLHHEQY
5JZI	MHCI	KLVALGINAV	AYGEDDKII	ASRRGPYEQY
8I5C	MHCI	VVGAVGVGK	AARDNSYQLI	ASGDTGGYEQY
4E41	MHCII	GELIGILNAAKVPAD	AVDRGSTLGRLY	ASSQIRETQY
6AM5	MHCI	SMLGIGIVPV	AVNFGGGKLI	ASSLSFGTEAF
3E2H	MHCI	QLSPFPFDL	AVSLERPYLT	ASGGGGTLY
3MV8	MHCI	HPVGEADYFEY	AVQDLGTSGSRLT	ASSARSGEFL
3KPR	MHCI	EEYLKAWTF	ILPLAGGTSYGKLT	ASSLGQAYEQY
5HYJ	MHCI	AQWGPDPAAA	AMRGDSSYKLI	ASSLWEKLAKNIQY
5KS9	MHCII	APSGEGSFQPSQENPQ	AVALNNNAGNMLT	ASSVAPGSDTQY
7T2D	MHCII	ATGLAWEWWRTVYE	ALSGSARQLT	ASSHREGETQY
1KJ2	MHCI	KVITFIDL	AARYQGGRALI	TCSAAPDWGASAETLY
6RPB	MHCI	SLLMWITQV	AVKSGGSYIPT	ASSYLNRDSDL
6UON	MHCI	GADGVGKSAL	AAAMDSSYKLI	ASSDPGTEAF
1ZGL	MHCII	VHFFKNIVTPRTPG	ALSGGDSSYKLI	ASSLADRVNTEAF
3PQY	MHCI	SSLENFRAYV	ILSGGSNYKLT	ASSFGREQY
4Y19	MHCII	QPLALEGSLQKRG	AASVYAGGTSYGKLT	ASRPRRDNEQF
6AVG	MHCI	APRGPHGGAASGL	LVVDQKLV	ASSGGHTGSNEQF
6V15	MHCII	GGYAPAKAAAT	ALSPSNTNKVV	ASSLDWGVNTLY
4Z7U	MHCII	APSGEGSFQPSQENPQ	ILRDRSNQFY	ASSTTPGTGTETQY
7RM4	MHCI	HMTTEVVRHC	ALDIYPHDMR	ASSLDPGDTGEFL
4QOK	MHCI	EAAGIGILTV	AVNVAGKST	AWSETGLGTGEFL
3PWP	MHCI	LGYGFVNYI	AVTTDSWGKLQ	ASRPGLAGGRPEQY
7N6E	MHCI	YLQPRTFLL	VVNRNNNDMR	AGQVTNTGEFL
8WUL	MHCI	VVGAVGVGK	AARSSGSWQLI	ASSQDRGDSAHTLY
3DXA	MHCI	EENLLDFVRF	IVWGGYQKV	ASRYRDDSNEQF
6RP9	MHCI	SLLMWITQV	ALTRGPGNQFY	ASSSPGGVSTEAF
4MJI	MHCI	TAFTIPSI	ATDDDSARQLT	ASSLTGGGEFL
6V19	MHCII	GGYAPAKAAAT	ALSDSSSFSLV	ASSLDWASQNTLY
7RDV	MHCII	EGRVRVNSAYQS	AASDDNNNRIF	ASGGQSNERLF
3E3Q	MHCI	QLSPFPFDL	AVSDPPPPLT	ASGGGGTLY
4MS8	MHCI	SPAEEAGFFL	AVSAKGTGSKLS	ASSDAPGQLY
2F53	MHCI	SLLMWITQC	AVRPTSGGSYIPT	ASSYVGNTGEFL
3QDJ	MHCI	AAGIGILTV	AVNFGGGKLI	ASSLSFGTEAF
6BJ2	MHCI	IPLTEEEAEL	ALSHNSGGSNYKLT	ASSFRGGKTQY
3QDM	MHCI	ELAGIGILTV	AGGTGNQFY	AISEVGVGQPQH
5TIL	MHCI	KAPYNFATM	AALYGNEKIT	ASSDAGGRNRTLY
6VMX	MHCI	RPPIFIRRL	AFGSSNTGKLI	ASSQDLFTGGYT

1448 Table 9: The samples contained in TCRxAI benchmarks

1449

1450

1451

1452

1453

1454

1458

1459

PDB	MHC	Peptide	CDRA3	CDRB3	
7RTR	MHCI	YLQPRTFLL	AVNRDDKII	ASSPDIEQY	
8EN8	MHCI	LPFDKSTIM	AADGGAGGSYQLT	SAGPTSGRTDTQY	
1YMM	MHCII	ENPVVHFFFKNIVTP	ATDTTSGTYKYL	SARDLTSGANNEQF	
5C07	MHCI	YQFGPDFPIA	AMRGDSSSYKLI	ASSLWEKLAKNIQY	
3VXU	MHCI	RFPLTFGWCF	GTYNQGGKLI	ASSGASHEQY	
6VQO	MHCI	HMTEVVRHC	AMSGLKEDSSYKLI	ASSIQQGADTQY	
1J8H	MHCII	PKYVKQNTLKLAT	AVSESPFGNEKLT	ASSSTGLPYGYT	
8ENH	MHCI	LPFEKSTIM	AADGGAGGSYQLT	SAGPTSGRTDTQY	
2P5W	MHCI	SLLMWITQC	AVRPLLDGTYIPT	ASSYLGNTGELF	
3UTT	MHCI	ALWGPDPAAA	AMRGDSSYKLI	ASSLWEKLAKNIQY	
7Q99	MHCI	NLSALGIFST	AVNVAGKST	AWSETGLGTGELF	
6ZKZ	MHCI	RLPAKAPL	AVTNQAGTALI	ASSYSIRGSRGEQF	
4GG6	MHCII	SGEGSFQPSQENP	ILRDGRGGADGLT	ASSVAVSAGTYEQY	
6TMO	MHCI	EAAGIGILTV	AVNDGGRLT	AWSETGLGMGGWQ	
3QIU	MHCII	ADLIAYLKQATKG	AAEPSSGQKLV	ASSLNANSDYT	
5WKF	MHCI	GTSGSPIVNR	GLGDAGNMLT	ASSLGQGLLYGYT	
8SHI	MHCI	VRSRRLRL	ATDALYSGGGADGLT	ASSYSEGEDEAF	
5D2N	MHCI	NLVPVMVATV	ILDNNNDMR	ASSLAPGTTNEKLF	
5KSA	MHCII	QPQQSFPEQEA	AVQFMDNSYQLI	ASSVAGTPSYEQY	
6MTM	MHCI	FEDLRLVLSF	GTERSGGYQKV	ASSMSAMGTEAF	
2BNQ	MHCI	SLLMWITQC	AVRPTSGGGYIPT	ASSYVGNTGELF	
4Z7V	MHCII	SGEGSFQPSQENP	ILRDSRAQKLV	ASSAGTSGEYEQY	
2F54	MHCI	SLLMWITQC	AVRPTSGGGYIPT	ASSYVGNTGELF	
5BS0	MHCI	ESDPPIVAQY	AVRPGGAGPFFVV	ASSFNMMATGQY	
6CQR	MHCII	RFYKTLRAEQASQ	AFKAAGNKL	ASSRLAGGMDEQF	
5M00	MHCI	KA VANFATM	AALYGNEKIT	ASSDDAAGGGGRNTLY	
7N2Q	MHCI	LRVMMLAPF	AVSNFNKFY	ASSVATYSTDQY	
6EQB	MHCI	AAAAGGIIGGIILTV	AVNDGGRLT	AWSETGLGMGGWQ	
4P2R	MHCII	ANGVAFFLTPFKA	AAEASNTNKVV	ASSLNANSDYT	
4P2Q	MHCI	ADGLAYFRSSFKGG	AAEASNTNKVV	ASSLNANSDYT	
8DNT	MHCI	LLLDRLNQL	AVREGAQKLV	ASSLDLGADEQF	
5E6I	MHCI	GILGFVFTL	AGPGGSSNTGKL	ASSLIYPGELF	
5TJE	MHCI	KAVYNFATM	AALYGNEKIT	ASSDAGGRNTLY	
2J8U	MHCI	ALWGFFFVL	ALFLASSFSKL	ASSDWVSYEQY	
1LP9	MHCI	ALWGFFFVL	ALFLASSFSKL	ASSDWVSYEQY	
3KPS	MHCI	EEYLQAFTY	ILPLAGGTSYGKLT	ASSLGQAYEQY	
2BNR	MHCI	SLLMWITQC	AVRPTSGGGYIPT	ASSYVGNTGELF	
5W1W	MHCI	VMAPRTLVL	AGQPLGGSNYKLT	ASSANPGDSSNEKLF	
1493	6CQL	MHCII	RFYKTLRAEQASQ	AFKAAGNKL	ASSRLAGGMDEQF
5C09	MHCI	YLGGPDFPTI	AMRGDSSYKLI	ASSLWEKLAKNIQY	
4MXQ	MHCI	SPAPRPLDL	AVSAKGTGSKLS	ASSDAPGQLY	
3SJV	MHCI	FLRGRAYGL	VVRAGKLI	ASGQGNFDIQQY	
1496	1QRN	MHCI	LLFGYAVYV	AVTTDSWGKLQ	ASRPGLAGGRPEQY
1497	3KXF	MHCI	LPEPLPQGQLTAY	ALSGFYNTDKL	ASPGLAGEYEQY
1498	5C0A	MHCI	MVWGPDPLYV	AMRGDSSYKLI	ASSLWEKLAKNIQY
1499	7N2N	MHCI	TRLALIAPK	AVLSPVQETSGSRLT	ASSVGLFSTDQY
1500	8ES9	MHCI	GVYDGREHTV	AVQPLNAGNNRKL	SAREWGGTEAF
1501	2NX5	MHCI	EPLPQGQLTAY	AVQASGGSYIPT	ATGTGDSNQPQH
1502	1G6R	MHCI	SIYRYYGL	AVSGFASALT	ASGGGGTL
1503	8GVB	MHCI	RYPLTFGW	AVGFTGGGNKLT	ASSDRDRVPETQY
1504	8TRL	MHCII	EIFDSGNPTGEV	IVNPANTGNQFY	ASRRDYFSYEQY
1505	5D2L	MHCI	NLVPVMVATV	AFITGNQFY	ASSQTQLWETQY
1506	5WLG	MHCI	SQLLNAKYL	ATVYAQGLT	ASSDWGDTGQLY
1507	5NMG	MHCI	SLFNTIAVL	AVRTNSGYALN	ASSDTVSYEQY
1508	7DZM	MHCI	TPQDLNTML	IVRGLNNAGNM	ASSLGIDAIY
1509	7BYD	MHCI	GGAI	LVGGGGYVLT	ASSQDLGAGEVYEQY
1510	5HHO	MHCI	GILEFVFTL	AGAGSQGNLI	ASSIRSSYEQY

1509

1510

1511

Table 10: The samples contained in TCRxAI benchmarks (continue table 1)

1512

1513

1514

1515

1516

1517

1518

1519

1520

1521

1522

1523

1524

1525

1526

1527

1528

1529

1530

1531

1532

1533

1534

1535

1536

1537

1538

1539

1540

1541

1542

1543

1544

1545

1546

1547

1548

1549

1550

1551

1552

1553

1554

1555

1556

1557

1558

1559

1560

1561

1562

1563

1564

1565

PDB	MHC	Peptide	CDRA3	CDRB3
1QSE	MHCI	LLFGYPRYV	AVTTDSWGKLQ	ASRPGLAGGRPEQY
3RGV	MHCI	WIYVYRPMGCGGS	AANSGTYQR	ASGDFWGDTLY
2E7L	MHCI	QLSPFPFDL	AVSHQGRYLT	ASGGGGTLY
3MBE	MHCII	GAMKRHGLDNYRGYSLG	AAEDGGSGNKLI	ASSWDRAGNLY
5M01	MHCI	KAPANFATM	AALYGNEKIT	ASSDDAAGGGRRNNTLY
5SWZ	MHCI	ASNENMETM	AASETSGSWQLI	ASSRDLGRDTQY
5NMF	MHCI	SLYNTIATL	AVRTNSGYALN	ASSDTVSYEQY
8GVI	MHCI	RYPLTFGW	AVVFTGGGNKLT	ASSLRDRVPETQY
7N5C	MHCI	SSLCNFRAYV	ILSGGCCNYKLT	ASSFGREQY
8TRR	MHCII	GYVATSSAVRLR	ALGDTGNYKV	ASSAVNSGNLY
4OZG	MHCII	APQPELPPQPG	IVLGGADGLT	ASSFRFTDTQY
4OZH	MHCII	APQPELPPQPGS	IVWGGATNKLI	ASSVRSTDTQY
2OL3	MHCI	SQYYYNSL	AMRGDYGGSGNKLI	TCSADRVGNTLY
7QPJ	MHCI	GLYDGMEHL	AVRGTGRRALT	ASSFATEAF
3VXM	MHCI	RFPLTFGWCF	AVGAPSGAGSYQLT	ASSPTSGIYEQY
6EQA	MHCI	AAAAGGIIGGIILTV	AVNVAGKST	AWSETGLGTGELF
2YPL	MHCI	KAFSPEVIMF	AVSGGYQKV	ASTGSYGYT
7RK7	MHCI	YMDGTMSSQV	LVALNYGGSQGNLI	AISPTEEGLIFPGNTIY
8WTE	MHCI	VVGAVGVGK	AARSSGSWQLI	ASSQDRGDSAETLY
3UTS	MHCI	ALWGPDPAAA	AMRGDSSYKLI	ASSLWEKLAKNIQY
1QSF	MHCI	LLFGYPVAV	AVTTDSWGKLQ	ASRPGLAGGRPEQY
1OGA	MHCI	GILGFVFTL	AGAGSQGNLI	ASSSRSSYEQY
2GJ6	MHCI	LLFGKPVYV	AVTTDSWGKLQ	ASRPGLAGGRPEQY
3QDG	MHCI	ELAGIGILTV	AVNFGGGKLI	ASSLSFGTEAF
2VLR	MHCI	GILGFVFTL	AGAGSQGNLI	ASSSRASYEQY
7NA5	MHCI	YGFRNVVHI	AVSNYNVLY	ASSQEPGGYAEQF
8CX4	MHCI	LRVMMMLAPF	AVNSPMSGAGSYQLT	ASSVGTYSTDQY
4PRI	MHCI	HPVGEADYFYEY	AVQDLGTSGSRLT	ASSARSGEFL
8YE4	MHCI	NYNYLYRLF	VVNAHSGAGSYQLT	ASSETGGYEQY
5M02	MHCI	KAPFNFMATM	AALYGNEKIT	ASSDAGGRNNTLY
2CKB	MHCI	EQYKFYSV	AVSGFASALT	ASGGGGTLY
3TFK	MHCI	QLSDVPMDL	AVSAKGTGSKLS	ASSDAPGQLY
7N2S	MHCI	TRLALIAPK	AVSLGTGAGSYQLT	ASSVGLYSTDQY
5KSB	MHCII	GPQQSFPEQEAA	AVQASGGSYIPT	ASSNRGLGDTQY
2UWE	MHCI	ALWGFPPVL	ALFLASSFSKLV	ASSDWVSYEQY
7Q9A	MHCI	LLLIGILVL	AVNVAGKST	AWSETGLGTGELF
5C08	MHCI	RQWGPDPAAV	AMRGDSSYKLI	ASSLWEKLAKNIQY
3HG1	MHCI	ELAGIGILTV	AVNVAGKST	AWSETGLGTGELF
8I5D	MHCI	VVGAVGVGK	AASSGSWQLI	ASSLEGTVEETLY
5JHD	MHCI	GILGFVFTL	AWGVNAGGTSYGKLT	ASSIGVYGYT
7JWJ	MHCI	ASNENMETM	AAVTGNTGKLI	ASSRGTIHSNTEVF
4MNQ	MHCI	ILAKFLHWL	AVDSATALPYGYI	ASSYQGTEAF
6PY2	MHCII	APFSEQECPVVLG	ASPQGGSEKLV	ASSSGGWGGGTEAF
7DZN	MHCI	TPQDLNTML	IVRGLNNAGNMLT	ASSLGIDAIY
4EUP	MHCI	ALGIGILTV	AVSGGGADGLT	ASSFLGTGVEQY
7N1E	MHCI	RLQLSQTYYV	ALSGFNNAGNMLT	ASSLGGAGGADTQY
3SEQ	MHCI	AAGIGILTV	AGGTGNQFY	AISEVGVGQPQH
2IAN	MHCII	GELIGTLNAAKVPAD	AALIQGAQKLV	ASTYHGTGY
2VLJ	MHCI	GILGFVFTL	AGAGSQGNLI	ASSSRSSYEQY
6CQN	MHCII	RFYKTLRAEQASQ	AFKAAGNKLT	ASSGLAGGMDEQF
3VXR	MHCI	RYPLTFGWCF	AVRMDSSYKLI	ASSSWDTGELF
7NMG	MHCI	LWMRLLPLL	AEPSGNTGKLI	ASSLHHEQY
3D3V	MHCI	LLFGPVYV	AVTTDSWGKLQ	ASRPGLAGGRPEQY
5ISZ	MHCI	GILGFVFTL	AFDTNAGKST	ASSIFGQREQY
6U3N	MHCII	APMPMPELPYP	AVGAGSNYQLI	ASSLEGQGASEQF
6RSY	MHCI	RMFPNAPYL	IGGGTTSGTYKYI	ASSLGFRDVMR
4MVB	MHCI	QPAEGGFQL	AVSAKGTGSKLS	ASSDAPGQLY

Table 11: The samples contained in TCRxAI benchmarks (continue table 2)

1566

1567

1568

1569

PDB	MHC	Peptide	CDRA3	CDRB3
1M15	MHCI	FLRGRAYGL	ILPLAGCTSYGKLT	ASSLGQAYEQY
3VXS	MHCI	RYPLTLGWCF	AVRMDSYYKLI	ASSSWDTGELF
7OW5	MHCI	VVVGAGGVGK	AMSVPSGDGSYQFT	ASKVPGQHNSPLH
8GVG	MHCI	RFPLTFW	AVGFTGGGNKLT	ASSDRDRVPETQY
2VLK	MHCI	GILGFVFTL	AGAGSQGNLI	ASSSRSSYEQY
8GOM	MHCI	RLQSLQTYV	ASSGNTPLV	ASTWGRASTDTQY
1D9K	MHCII	GNSHRGAIIEWEGIESG	AATGSFNKLT	ASGGQGRAEQQF
6CQQ	MHCII	RFYKTLRAEQASQ	AFKAAGNKLT	ASSRLAGGMDEQF
5HHM	MHCI	GILGLVFTL	AGAGSQGNLI	ASSSSRSSYEQY
4N0C	MHCI	MPAGRWPDL	AVSAKGTGSKLS	ASSDAPGQLY
5C0B	MHCI	RQFGPDFPTI	AMRGDSSYKLI	ASSLWEKLAKNIQY
7PB2	MHCI	VVVGADGVGK	ALSGPSGAGSYQLT	ASSYGPQHNSPLH
1MWA	MHCI	EQYKFYSV	AVSGFASALT	ASGGGGTLY
4QRP	MHCI	HSKKKCDEL	ALSDPVNDMR	ASSLRGRGDQPQH
6RPA	MHCI	SLLMWITQV	AVRDINSGAGSYQLT	SVGGSGGADTQY
7N5P	MHCI	SSLCNFRAYV	ILSGGSNYKLT	ASSFFGREQY
7N1F	MHCI	YLQPRTFLL	AVNRDDKII	ASSPDIEQY
5C0C	MHCI	RQFGPDWIVA	AMRGDSSYKLI	ASSLWEKLAKNIQY
6D78	MHCI	AAGIGILTV	AVNFGGGKLI	ASSWSFGTEAF
4JFF	MHCI	ELAGIGILTV	AVNDGGRLT	AWSETGLGMGGWQ
4N5E	MHCI	VPYMAEFGM	AVSAKGTGSKLS	ASSDAPGQLY
4JRX	MHCI	LPEPLPQGQLTAY	ALSGFYNTDKLII	ASPGETEAF
7NMF	MHCI	QLPRLFPLL	AEPSGNTGKLI	ASSLHHEQY
3QIW	MHCII	ADLIAYLEQATKG	AAEPSSGQKLV	ASSLNNANSNDYT
6ZKX	MHCI	RLPAKAPLLGCG	AVTNQAGTALI	ASSYSIRGSRGEQF
1NAM	MHCI	RGYVYQGL	AMRGDYGGSGNKLII	TCSADRVGNTLY
8PJG	MHCII	PKYVKQNTLKLAR	AVSEQDDKII	ATSDESYGYT
8VCX	MHCII	GQVELGGGPGAESQ	IVSHNAGNMLT	ASSLERETQY
5YXU	MHCI	KLVALGINAV	AYGEDDKII	ASRRGSAELY
3O4L	MHCI	GLCTLVAML	AEDNNARLM	SARDGTGNGYT
7SG2	MHCII	QPPFPQEQQFPFGS	LVGGLARDMR	SVALGSDTGEELF
8GON	MHCI	RLQSLQIYV	ASSGNTPLV	ASTWGRASTDTQY
2JCC	MHCI	ALWGFFPVL	ALFLASSFSKLV	ASSDWVSYEQY
6G9Q	MHCI	KAPYDYAPI	AALYGNEKIT	ASSDAGGRNTLY
6DKP	MHCI	ELAGIGILTV	AVNFGGGKLI	ASSWSFGTEAF
5NQK	MHCI	ELAGIGILTV	AGGGGADGLT	ASSQGLAGAGELF
2PYE	MHCI	SLLMWITQC	AVRPLLDGTYIPT	ASSYLGNTGEELF
6R2L	MHCI	SLSKILDTV	AVGGNDWNTDKLII	ASSPLDVSISYYNEQF
2O19	MHCI	QLSPFPFDL	AVSGFASALT	ASGGGGTLY
8F5A	MHCI	TSTLQEIQIGW	AVTLNNNAGNMLT	ASSVGGTEAF
7Z50	MHCII	LQTIALEVEDDPC	AASVRNYKYV	ASSRQGQNTLY
6BGA	MHCII	YVVVPD	AALRATGGNNKLT	ASSLNWSQDTQY
3MV7	MHCI	HPVGEADYFEY	AVVQDLGTSGSRLT	ASSARSGELF
8VD2	MHCII	GQVELGGGPIESC	IVRVAIEGSQGNLI	ASSLRRGDTIY
8VCY	MHCII	GQVELGGGSSPETCI	IVSHNAGNMLT	ASSLERETQY
5YXN	MHCI	KLVALGINAV	AYGEDDKII	ASRRGPYEQY
5E9D	MHCI	ELAGIGILTV	AVTKYSWGKLQ	ASRPWMAGGVELY
6AMU	MHCI	MMWDRGLGMM	AVNFGGGKLI	ASSLSFGTEAF
5BRZ	MHCI	EVDPIGHLY	AVRPGGAGPFFVV	ASSFNMATGQY
3TJH	MHCI	SPLDSLWVI	AVSAKGTGSKLS	ASSDAPGQLY
3H9S	MHCI	MLWGYLQYY	AVTTDSWGKLQ	ASRPGLAGGRPEQY
4PRP	MHCI	HPVGQADYFEY	AVQDLGTSGSRLT	ASSARSGELF
5IVX	MHCI	RGPGRAFVTI	AASASFQDNSKLI	ASSLGHTEVF
4Y1A	MHCII	LQPLALEGSLQKRG	AASSSSAGGTSYGKLT	ASRPRDPVTQY
2IAM	MHCII	GELIGILNAAKVPAD	AALIQGAQKLV	ASTYHGTGY
6U3O	MHCII	AVVQSELPYPEGS	IAFQGAQKLV	ASSFRALAADTQY

1617

1618

1619

Table 12: The samples contained in TCRxAI benchmarks (continue table 3)

1620  
1621  
1622  
1623  
1624  
1625

	PDB	MHC	Peptide	CDRA3	CDRB3
1627	6V0Y	MHCII	GGYAPAKAAAT	ALSDSGSFNKLT	ASSLDWGGQNTLY
1628	5NME	MHCI	SLYNTVATL	AVRTNSGYALN	ASSDTVSYEQY
1629	7T2C	MHCII	TGLAWEWWRTVY	LVGDTGFQKLV	SARDPGGGSSYEQY
1630	2PXY	MHCII	RGGASQYRPSQ	ALSENYGNEKIT	ASGDASGAETLY
1631	4G9F	MHCI	KRWIIMGLNK	AMRDLRDNFNKFY	ASREGLGGTEAF
1632	3QIB	MHCII	ADLIAYLKQATKG	AALRATGGNNKLT	ASSLNWSQDTQY
1633	4G8G	MHCI	KRWIILGLNK	AMRDLRDNFNKFY	ASREGLGGTEAF
1634	7N2O	MHCI	LRVMMMLAPF	AVLSPVQETSGSRLT	ASSVGLFSTDTQY
1635	5TEZ	MHCI	GILGFVFTL	AASFIIQGAQKLV	ASSLLGGWSEAF
1636	3D39	MHCI	LLFGPVYV	AVTTDSWGKLQ	ASRPGLLAGGRPEQY
1637	6ZKW	MHCI	RLPAKAPLL	AVTNQAGTALI	ASSYSIRGSRGEQF
1638	4FTV	MHCI	LLFGYPVYV	AVTTDSWGKLQ	ASRPGLMSAQPEQY
1639	6PX6	MHCII	APFSEQEQPVLG	AVHTGARLM	ASSHGASTDTQY
1640	6V1A	MHCII	GGYRAPAKAAAT	ALSDSSSFSKLV	ASSLDWASQNTLY
1641	1U3H	MHCII	SRGGASQYRPSQ	AASANSGYQQR	ASGDAGGGYEQY
1642	7N4K	MHCI	SSLENFRRAYV	ILSGGSNYKLT	ASSFFGREQY
1643	2Z31	MHCII	RGGASQYRPSQ	ALSENYGNEKIT	ASGDASGGNTLY
1644	2ESV	MHCI	VMAPRTLIL	IVVRSSNTGKLI	ASSQDRDTQY
1645	5EUO	MHCI	GILGFVFTL	AGAIGPSNTGKLI	ASSIRSSYEQY
1646	4JFD	MHCI	ELAAIGILT	AVNDGGRLT	AWSETGLGMGGWQ
1647	6ZKY	MHCI	RLPAKAPL	AVTNQAGTALI	ASSYSIRGSRGEQF
1648	6TRO	MHCI	GVYDGREHTV	VVNHSGGSYIPT	ASSFLMTSGDPYEQY
1649	7N2P	MHCI	GQVMVVAPR	AVSNFNKFY	ASSVATYSTDTQY
1650	7R80	MHCI	QASQEVKNW	AQLNQAGTALI	ASSYGTGINYGYT
1651	1BD2	MHCI	LLFGYPVYV	AAMEGAQKLV	ASSYPGGGFYEQY
1652	4L3E	MHCI	ELAGIGILT	AVNFGGGKLI	ASSWSFGTEAF
1653	7PHR	MHCI	YLEPGPVTV	ATDGSTPMQ	ASSWGAPYEQY
1654	3FFC	MHCI	FLRGRAYGL	AMREDTGNQFY	ASSFTWTSGGATDTQY
1655	4JRY	MHCI	LPEPLPQGQLTAY	AVGGGSNYQLI	ASSRTGSTYEQY
1656	5SWS	MHCI	ASNENMETM	AASEGSGSWQLI	ASSAGLDAEQY
1657	6UZ1	MHCI	LLFGYPVYV	AVTTDRSGKLQ	ASRPAGAAGGRPELY
1658	1FO0	MHCI	INFDFNTI	AMRGDYGGSGNKLI	TCSADRVGNTLY
1659	7JWI	MHCI	ASNENMETM	AASETSGSWQLI	ASSRDLGRDTQY
1660	8D5Q	MHCI	HPGSVNEFDF	ALGDPTGANTGKLT	TCSAGRGGYAEQF
1661	6VRM	MHCI	HMTEVVVRHC	VVQPGGYQKVT	ASSEGGLWQVGDEQY
1662	7N2R	MHCI	TRLALIAPK	AVSNFNKFY	ASSVATYSTDTQY
1663	1FYT	MHCII	PKYVKQNTLKLAT	AVSESPFGNEKLT	ASSSTGLPYGYT
1664	3QFJ	MHCI	LLFGFPVYV	AVTTDSWGKLQ	ASRPGLLAGGRPEQY
1665	3GSN	MHCI	NLVPMVATV	ARNTGNQFY	ASSPVTGGIYGYT
1666	6V13	MHCII	GGYRAPAKAAAT	ALSPSNTNKVV	ASSLDWGVNTLY
1667	7OW6	MHCI	VVVGADGVGK	AMSVPSPGDGSYQFT	ASKVGPQGHNSPLH
1668	4OZF	MHCII	APQPELPYPQPGS	IAFQGAQKLV	ASSFRALAADTQY
1669	4JFE	MHCI	ELAGIGALT	AVNDGGRLT	AWSETGLGMGGWQ
1670	3MV9	MHCI	HPVGEADYFEY	AVQDLGTSGSRLT	ASSARSGEFL
1671	6Q3S	MHCI	SLLMWITQV	AVRPTSGGSYIPT	ASSYVGNTGEFL
1672	5MEN	MHCI	ILAKFLHWL	AVDSATSGTYKYI	ASSYQGTEAF
1673	1AO7	MHCI	LLFGYPVYV	AVTTDSWGKLQ	ASRPGLLAGGRPEQY
1674	4H1L	MHCII	QHIRCNIPKRISA	AVGASGNTGKLI	ASSLRDGYTGEFL

Table 13: The samples contained in TCRxAI benchmarks (continue table 4)

---

1674  
1675

## B REPRODUCIBILITY STATEMENT

1676  
1677  
1678  
1679  
1680  
1681

To comply with the double-blind review policy while ensuring reproducibility, we provide a self-contained code package with detailed training instructions. Model weights are not included due to their large size; however, the protein language model weights can be obtained from publicly available repositories referenced in our documentation. After publication, we will release a public repository containing the full code and trained weights.

1682  
1683

## C LARGE LANGUAGE MODEL USAGE STATEMENT

1684  
1685

We employed large language models (LLMs), primarily ChatGPT, in two limited ways:

1686  
1687  
1688

- as a coding assistant, and
- for polishing written text.

1689  
1690  
1691  
1692

**Coding Assistant** LLMs were consulted to clarify documentation, organize API references, and suggest debugging strategies. All code, documentation, and fixes obtained were manually reviewed and verified by the authors.

1693  
1694  
1695  
1696  
1697  
1698  
1699  
1700  
1701  
1702  
1703  
1704  
1705  
1706  
1707  
1708  
1709  
1710  
1711  
1712  
1713  
1714  
1715  
1716  
1717  
1718  
1719  
1720  
1721  
1722  
1723  
1724  
1725  
1726  
1727

**Polishing Article** LLMs were used only to refine the clarity and style of sentences written by the authors and to format tables from raw data. No raw text or substantive content was generated by LLMs. All refined content was manually checked and further revised by the authors.