HIFT-TCR: INTERFACE-STRUCTURED TCR LIBRARY SCREENING UNDER EPITOPE-FAMILY SHIFT

Anonymous authors
Paper under double-blind review

ABSTRACT

Emerging pathogen variants and tumor neoantigens motivate screening characterized T-cell receptor (TCR) libraries against new peptide families. We introduce the Hierarchical Interface Field Transformer for T-cell Receptor screening (HIFT-TCR), a structure-based graph Transformer for modeled interfaces between TCRs and peptide–major histocompatibility complexes (pMHCs). It combines typed local interactions, molecular hierarchy, and separate TCR–peptide and TCR–MHC compatibility. Five-fold cross-validation on the STAG pan-peptide benchmark holds out complete peptide components defined at 75% sequence identity. Across 25 matched seed–fold runs, HIFT-TCR achieves 0.5874 AUROC, 0.2679 AUPRC, and 0.5112 standardized partial AUROC at a false-positive rate of 0.1, leading six same-split sequence, language-model, and structure baselines across all three mean metrics. Performance is higher without an exact paired-CDR3 training match. Retrained ablations support the combined objective, interface topology, and hierarchy, while interventions identify TCR residues and TCR–peptide edges as consistent support for positive predictions. These findings support interface-structured receptor prioritization for experimental binding validation under peptide-family shift.

1 INTRODUCTION

Pathogen evolution and tumor mutations continually change the peptides presented to T cells. DNA-barcoded pMHC multimer assays characterize antigen-specific T cells at scale (Bentzen et al., 2016), pre-existing memory T cells can recognize emerging-virus epitopes (Mateus et al., 2020), and T cells bearing engineered TCRs can sustain antigen-specific responses in adoptive therapy (Rapoport et al., 2015). These advances motivate screening characterized receptor libraries against new peptide families to prioritize experimental binding validation.

New peptides challenge screening because they lie outside the training distribution. TCR specificity is relational because in a class-I complex, the receptor contacts both peptide and MHC (Rossjohn et al., 2015). Cross-reactivity, recurring receptor motifs (Dash et al., 2017), and convergent specificity groups (Glanville et al., 2017) also preclude reducing specificity to a single sequence match. Models must combine sequence and interface evidence while preserving partner identity.

Evaluation can obscure this challenge. Data source, negative construction, and held-out entity influence reported TCR–epitope performance (Drost et al., 2025; Lu et al., 2026). Random example splits may place close peptide relatives across partitions, obscuring peptide similarity's contribution to apparent generalization (Bernett et al., 2024).

We therefore hold out complete peptide sequence-identity components, termed "peptide families," while allowing previously measured receptors to recur against other peptides. We quantify exact paired-CDR3 recurrence explicitly. This tests transfer to unseen peptide families within a characterized receptor library.

The Hierarchical Interface Field Transformer for T-cell Receptor screening (HIFT-TCR) addresses this task with a graph Transformer on role-annotated residue graphs of modeled TCR–peptide–MHC class I complexes. Each encoder block combines typed edge-gated local messages, role-aware interface-shell attention (RISA) with one-hop and two-hop heads, and molecular-role and whole-complex context pooling. By preserving molecular-partner identity from local interactions to

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

binding prediction, HIFT-TCR unifies interface evidence and molecular context for characterized-library screening under peptide-family shift.

2 Related Work

Sequence-based TCR–peptide prediction. Sequence predictors differ in representation and peptide generalization. ERGO-II (Springer et al., 2021) uses recurrent or pretrained autoencoder receptor representations and a recurrent peptide encoder. TCR-ESM (Yadav et al., 2024) classifies binding from pretrained protein-language-model embeddings. PanPep (Gao et al., 2023) combines meta-learning and external memory for peptides with limited or no known binders. TULIP (Meynard-Piganeau et al., 2024) learns conditional sequence probabilities from interacting pairs, scoring unseen epitopes without synthetic nonbinding training pairs. EPACT (Zhang et al., 2024) combines pretrained TCR/pMHC representations with epitope-anchored contrastive transfer learning.

Complementary binding evidence offers another approach. NetTCR-2.2 (Jensen & Nielsen, 2024) combines pan-specific and peptide-specific training with sequence-similarity predictions. UniPMT (Zhao et al., 2025) jointly models peptide–MHC, peptide–TCR, and peptide–MHC–TCR binding through a heterogeneous molecular-entity graph. Its graph links binding evidence across examples, whereas HIFT-TCR's residue graphs describe interfaces within complexes. Sequence prediction can also use structural supervision such as EPACT (Zhang et al., 2024) which includes residue-interaction fine-tuning. HIFT-TCR instead explicitly organizes modeled residue interfaces throughout prediction.

Structure-aware TCR–pMHC prediction. STAG (Slone et al., 2025a) uses spatial and physicochemical features on modeled residue graphs whereas STAG-LLM (Slone et al., 2025b) adds protein-language-model information. TCRen (Karnaukhov et al., 2024) ranks unseen candidate epitopes using TCR–peptide contact maps and a crystal-structure-derived statistical potential. SageTCR (Li et al., 2025) exchanges residue-level and atom-level information before graph aggregation and evaluates unseen-epitope transfer. TCRLens (Siriarchawatana et al., 2026) uses equivariant residue-interface message passing for peptide–MHC, peptide–TCR, and full-complex prediction. Structural and multiscale representations thus have established precedents. HIFT-TCR integrates typed local interactions, molecular hierarchy, and separate TCR–peptide and TCR–MHC compatibility.

Graph Transformers and receptive fields. Message passing (Gilmer et al., 2017) aggregates graph neighborhoods, while Transformer attention (Vaswani et al., 2017) exchanges broader context. GraphGPS (Rampášek et al., 2022) combines both with positional or structural encodings. Graphormer (Ying et al., 2021) adds shortest-path attention information whereas Retroformer (Wan et al., 2022) combines local molecular-graph and global sequence attention for retrosynthesis. HopFormer (Yun et al., 2026) explicitly assigns hop limits to individual attention heads. STAG (Slone et al., 2025a) and SageTCR (Li et al., 2025) expand graph receptive fields through successive graph-convolution layers. HIFT-TCR adapts this head-specific principle to typed TCR–pMHC interfaces, integrating local messages, molecular-role context, and partner-factorized prediction.

Evaluation under biological distribution shift. The held-out entity defines the task. IMMREP22 (Meysman et al., 2023) compared predictors and emphasized independent follow-up evaluation. ePytope-TCR (Drost et al., 2025) finds pronounced differences between frequently and sparsely observed epitopes. Lu et al. (2026) show effects of negative construction and independent testing. Related train/test examples can inflate generalization estimates (Bernett et al., 2024), and DOME (Walsh et al., 2021) emphasizes application-aligned validation. Random splits can therefore obscure peptide-family transfer. We hold out complete peptide components while conditioning on a characterized receptor library and measuring exact receptor recurrence.

All in all, these studies motivate peptide-family exclusion and explicit interface organization. HIFT-TCR evaluates typed local evidence, molecular-role context, controlled receptive fields, and separate TCR–peptide and TCR–MHC reasoning together in characterized-library screening under peptide-family shift.

3 Problem Formulation and Evaluation Setting

3.1 Characterized-Receptor-Library Screening under Peptide-Family Shift

Let $\mathcal{D} = \{(G_n, y_n, p_n, r_n)\}_{n=1}^{N_{\mathrm{ex}}}$, where G_n is the modeled TCR–peptide–MHC class I residue graph, $y_n \in \{0, 1\}$ is the released interaction label, p_n is the peptide, r_n is the paired receptor, and N_{ex} is the number of examples. For outer fold k, let $\mathcal{I}_s^{(k)}$ contain the indices assigned to split $s \in \mathcal{S}_{\mathrm{split}} = \{\mathrm{train}, \mathrm{val}, \mathrm{test}\}$.

Let $\mathcal{P} = \{p_n : 1 \le n \le N_{\mathrm{ex}}\}$ be the set of observed peptides. Peptides are assigned to sequence-identity components by $C : \mathcal{P} \longrightarrow \mathcal{C}$, where $C(p_n)$ is the component containing peptide p_n. The partitions are peptide-component disjoint:

$$\{C(p_n) : n \in \mathcal{I}_s^{(k)}\} \cap \{C(p_n) : n \in \mathcal{I}_t^{(k)}\} = \varnothing, \qquad s, t \in \mathcal{S}_{\mathrm{split}}, \quad s \neq t. \tag{1}$$

No peptide component crosses partitions. Receptor recurrence is permitted to evaluate unseen-peptide transfer within a characterized receptor library. Component disjointness also prevents identical TCR–peptide–MHC sequence complexes from crossing partitions.

3.2 Cohort and Component-Disjoint Splits

The pinned STAG (Slone et al., 2025a) pan-epitope cohort contains 22,410 rows, 4,890 positives (21.82%), 344 peptides, and 19,282 unique sample identifiers. STAG integrates McPAS-TCR (Tickotsky et al., 2017), VDJdb (Bagaev et al., 2020), IEDB (Vita et al., 2015), and 10x resources, with observed and paired negatives. Structures follow its inherited reconstruction and homology-modeling pipeline: Stitchr (Heather et al., 2022), TCRpMHCmodels (Jensen et al., 2019), and ImmuneScape (Lusiany et al., 2022).

To prevent close peptide relatives from crossing partitions, we normalize peptide strings and define identity as $1 - d_{\mathrm{Lev}}(a, b)/\max(|a|, |b|)$, using unit-cost Levenshtein distance. Transitive single-link clustering at identity 75% produces 249 peptide components. Whole components are assigned to five groups while balancing rows and positives. In each outer fold, one group is test, the next is validation, and the other three are training, yielding exactly 13,446/4,482/4,482 released rows. Among 47,096 cross-fold peptide comparisons, none reaches 75% identity and the maximum is 0.7273. As peptide components are indivisible, exact TCR–peptide–MHC sequence complexes remain confined to one partition.

The same frozen manifest governs every method. CDR3 strings and α/β role labels are inherited from the dataset. An independent check with ANARCI (Dunbar & Deane, 2016) and IMGT numbering (Lefranc et al., 2003) confirms the expected chain roles and CDR3 anchors.

4 HIFT-TCR Architecture

Figure 1 follows HIFT-TCR from interface construction to prediction. For B complexes, $b(i)$ identifies residue i's complex and $N = \sum_{b=1}^{B} |V_b|$ counts retained residues. Blocks use indices $\ell \in \{0, 1, 2\}$ and Drop denotes dropout with probability 0.20.

4.1 Interface Graph Construction and Physicochemical Encoding

Complex b has residue graph $G_b = (V_b, \mathcal{E}_b)$ (Figure 1, Part A). The map $\rho : V_b \to \{P, M, T_\alpha, T_\beta, O\}$, assigns each residue i its molecular role $\rho(i)$.

Cropping uses Cα coordinates, or available-atom centroids when Cα is absent. All peptide residues remain. TCR residues remain within 8 Å of peptide or MHC; MHC residues remain within 8 Å of peptide or TCR. If fewer than eight residues remain, the uncropped parsed set is restored. Directed edges connect retained residues within 8 Å minimum heavy-atom distance or consecutive in retained same-chain order, with both directions stored.

A node vector $x_i \in \mathbb{R}^{38}$ concatenates a 21-state one-hot amino-acid code, four continuous descriptors (charge, normalized hydropathy (Kyte & Doolittle, 1982), normalized volume (Zamyatnin,

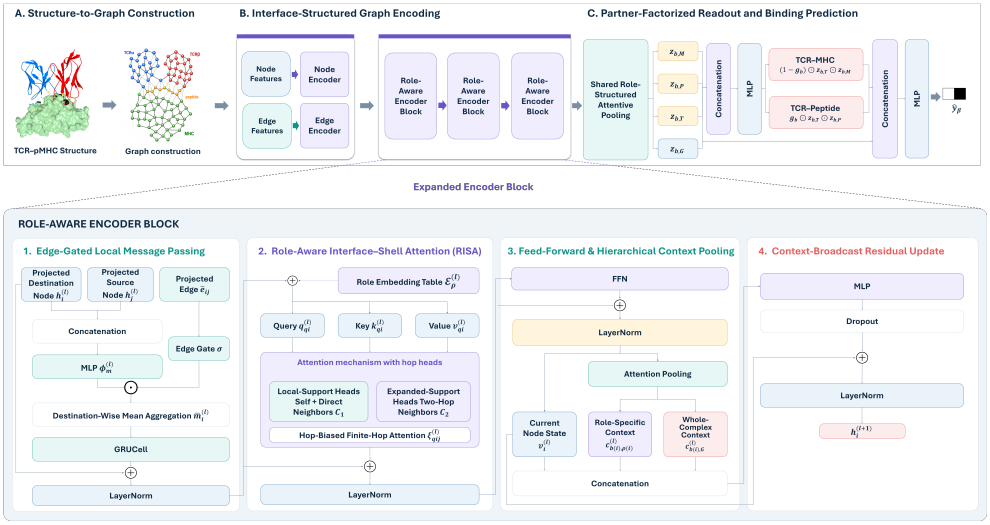


Figure 1: HIFT-TCR architecture. Modeled TCR–pMHC graphs are encoded with edge-gated local messages, RISA, and molecular hierarchy, then pooled by a partner-factorized readout. Training combines binding BCE, edge-proximity self-reconstruction, and pairwise ranking.

1972), and normalized polarity (Grantham, 1974)), eight binary physicochemical indicators, and a five-state one-hot molecular-role code for peptide, MHC, TCRα, TCRβ, or other. For each directed edge $j \to i$, $e_{ji} \in \mathbb{R}^{22}$ contains scaled and inverse distance, proximity and interaction-proxy indicators, distance-scaled charge and hydropathy products, polarity and volume differences, and six chain, sequential, and molecular-partner relations.

Node and edge features share a projected latent space (Figure 1, Part B):

$$h_i^{(0)} = \mathrm{GELU}\big(\mathrm{LN}(W_x x_i + b_x)\big), \qquad \bar{e}_{ji} = \mathrm{GELU}\big(\mathrm{LN}(W_e e_{ji} + b_e)\big), \qquad h_i^{(0)}, \bar{e}_{ji} \in \mathbb{R}^{128}. \tag{2}$$

Here, LN denotes layer normalization. Projected edge states are computed once and reused across all three encoder blocks.

4.2 Edge-Gated Local Interface Field

Each block first propagates typed local evidence (Figure 1, Part B1). On each edge $j \to i$, a block-specific message MLP and vector gate compute

$$m_{ji}^{(\ell)} = \phi_m^{(\ell)}([h_j^{(\ell)}, h_i^{(\ell)}, \bar{e}_{ji}]) \odot \sigma(W_g^{(\ell)} \bar{e}_{ji} + b_g^{(\ell)}), \tag{3}$$

where $\phi_m^{(\ell)}$ has dimensions $384 \to 256 \to 256 \to 128$. Incoming messages are mean-aggregated at their destination:

$$\bar{m}_i^{(\ell)} = \frac{\sum_{j \in \mathcal{N}_{\mathrm{in}}(i)} m_{ji}^{(\ell)}}{\max(1, |\mathcal{N}_{\mathrm{in}}(i)|)}, \qquad \mathcal{N}_{\mathrm{in}}(i) = \{j : (j,i) \in \mathcal{E}_b\} \tag{4}$$

A gated recurrent unit cell (GRUCell) receives aggregate $\bar{m}_i^{(\ell)}$ as input and destination state $h_i^{(\ell)}$ as its hidden state. A dropout–residual–LayerNorm update yields $\hat{h}_i^{(\ell)} \in \mathbb{R}^{128}$. Typed edges thus control latent message propagation before attention.

4.3 Role-Aware Interface-Shell Attention

RISA assigns four one-hop and four two-hop heads per block (Figure 1, Part B2), adapting the head-specific hop budgets of HopFormer (Yun et al., 2026) to typed residue interfaces. Let $A_b \in \{0,1\}^{|V_b| \times |V_b|}$ be complex b's symmetrized loop-free adjacency and I_b the identity. Multiplication

is ordinary; inequalities, $\mathbf{1}[\cdot]$, and Boolean AND/OR/NOT ($\wedge, \vee, \neg$) act elementwise. Exact graph-distance classes are:

$$\begin{aligned} D_{b,0} &= I_b, & D_{b,1} &= A_b, \\ D_{b,2} &= \mathbf{1}[A_b A_b > 0] \wedge \neg(D_{b,0} \vee D_{b,1}), & & \\ R_{b,1} &= D_{b,0} \vee D_{b,1}, & R_{b,2} &= R_{b,1} \vee D_{b,2}. \end{aligned} \tag{5}$$

Here, $D_{b,0}$, $D_{b,1}$, and $D_{b,2}$ represent self, direct neighbors, and exact two-hop neighbors; $R_{b,1}$ and $R_{b,2}$ give cumulative supports.

With zero-based heads $q \in \{0, \ldots, 7\}$, let

$$k_q = (q \bmod 2) + 1, \qquad \mathcal{N}_q(i) = \{j : (R_{b(i),k_q})_{ij} = 1\}. \tag{6}$$

Heads $0, 2, 4, 6$ use self-plus-one-hop support; heads $1, 3, 5, 7$ additionally admit exact two-hop neighbors. The two bias classes are:

$$C_{b,1} = D_{b,0} \vee D_{b,1}, \qquad C_{b,2} = D_{b,2}. \tag{7}$$

Self and direct neighbors share the first bias; newly reached two-hop neighbors use the second. The learned block-specific role table $\mathcal{E}_\rho^{(\ell)} \in \mathbb{R}^{5\times 128}$ supplies row $\rho(i)$ to attention:

$$a_i^{(\ell)} = \hat{h}_i^{(\ell)} + \mathcal{E}_\rho^{(\ell)}[\rho(i)] \in \mathbb{R}^{128}. \tag{8}$$

For each head, learned affine projections map $a_i^{(\ell)}$ into d_h-dimensional query, key, and value vectors:

$$\mathbf{q}_{qi}^{(\ell)} = W_{Qq}^{(\ell)} a_i^{(\ell)} + b_{Qq}^{(\ell)}, \quad \mathbf{k}_{qi}^{(\ell)} = W_{Kq}^{(\ell)} a_i^{(\ell)} + b_{Kq}^{(\ell)}, \quad \mathbf{v}_{qi}^{(\ell)} = W_{Vq}^{(\ell)} a_i^{(\ell)} + b_{Vq}^{(\ell)}. \tag{9}$$

For query residue i and an admissible key residue $j \in \mathcal{N}_q(i)$, the attention logit is

$$\xi_{qij}^{(\ell)} = \frac{\mathbf{q}_{qi}^{(\ell)\top} \mathbf{k}_{qj}^{(\ell)}}{\sqrt{d_h}} + \sum_{c=1}^{2} \beta_{qc}^{(\ell)} \mathbf{1}[(C_{b(i),c})_{ij} = 1]. \tag{10}$$

The head dimension d_h is 16; $\beta^{(\ell)} \in \mathbb{R}^{8\times 2}$ contains head-specific biases $\beta_{qc}^{(\ell)} \in \mathbb{R}$ for the two classes. Pairs outside $\mathcal{N}_q(i)$ receive $-\infty$ logits. Permitted logits are normalized to aggregate values:

$$\begin{aligned} \alpha_{qij}^{(\ell)} &= \frac{\exp \xi_{qij}^{(\ell)}}{\sum_{m\in\mathcal{N}_q(i)} \exp \xi_{qim}^{(\ell)}}, \qquad j \in \mathcal{N}_q(i), \\ o_{qi}^{(\ell)} &= \sum_{j\in\mathcal{N}_q(i)} \mathrm{Drop}(\alpha_{qij}^{(\ell)}) \mathbf{v}_{qj}^{(\ell)} \in \mathbb{R}^{16}. \end{aligned} \tag{11}$$

The head outputs $o_{qi}^{(\ell)}$ are concatenated and projected as

$$o_i^{(\ell),\mathrm{att}} = W_O^{(\ell)} [o_{0i}^{(\ell)} \| \cdots \| o_{7i}^{(\ell)}] + b_O^{(\ell)} \in \mathbb{R}^{128}. \tag{12}$$

A dropout–residual–LayerNorm update combines this output with $\hat{h}_i^{(\ell)}$. A position-wise $128 \to 256 \to 128$ FFN and a second residual update then produce $v_i^{(\ell)} \in \mathbb{R}^{128}$.

4.4 Role Context and Partner-Factorized Readout

After the FFN as shown in Figure 1 Part B3, $v_i^{(\ell)}$ is shared by the whole-complex and role-specific pools. For any nonempty set $S \subseteq V_b$,

$$\begin{aligned} \gamma_i^{(\ell)} &= w_h^{(\ell)\top} v_i^{(\ell)} + b_h^{(\ell)}, \\ \omega_i^{(\ell)}(S) &= \frac{\exp \gamma_i^{(\ell)}}{\sum_{j\in S} \exp \gamma_j^{(\ell)}}, \\ c^{(\ell)}(S) &= \sum_{i\in S} \omega_i^{(\ell)}(S) v_i^{(\ell)} \in \mathbb{R}^{128}. \end{aligned} \tag{13}$$

For $V_{b,r} = \{i \in V_b : \rho(i) = r\}$, set $c_{b,G}^{(\ell)} = c^{(\ell)}(V_b)$ and $c_{b,r}^{(\ell)} = c^{(\ell)}(V_{b,r})$ for nonempty roles; absent roles use $c_{b,G}^{(\ell)}$. Each residue concatenates its state, role context, and whole-complex context:

$$\psi_i^{(\ell)} = [v_i^{(\ell)} \| c_{b(i),\rho(i)}^{(\ell)} \| c_{b(i),G}^{(\ell)}] \in \mathbb{R}^{384}. \tag{14}$$

A $384 \to 256 \to 128$ hierarchy MLP and dropout–residual–LayerNorm update yield $h_i^{(\ell+1)} \in \mathbb{R}^{128}$ (Figure 1, Part B4). Three independently parameterized blocks apply the local field, RISA, FFN, and context broadcast, yielding $H^{(3)} \in \mathbb{R}^{N \times 128}$.

After the third block, set $S_{b,G} = V_b$; $S_{b,T}$ contains combined TCRα/TCRβ residues, $S_{b,P}$ peptide, and $S_{b,M}$ MHC. A shared score and value projection pool each subset (Figure 1, Part C). For $R \in \{G, T, P, M\}$,

$$\begin{aligned} \nu_i &= w_p^\top h_i^{(3)} + b_p, \\ \eta_i(S_{b,R}) &= \frac{\exp \nu_i}{\sum_{j \in S_{b,R}} \exp \nu_j}, \\ z_{b,R} &= \sum_{i \in S_{b,R}} \eta_i(S_{b,R})(W_v h_i^{(3)} + b_v) \in \mathbb{R}^{128}. \end{aligned} \tag{15}$$

TCRα and TCRβ retain separate encoder roles but pool jointly at readout. A feature-wise gate combines TCR, peptide, and MHC summaries:

$$g_b = \sigma(\phi_g[z_{b,T} \| z_{b,P} \| z_{b,M}]) \in (0,1)^{128}, \tag{16}$$

where ϕ_g has dimensions $384 \to 128 \to 128$ with GELU between affine maps and no LayerNorm or dropout. TCR–peptide and TCR–MHC paths combine as

$$\begin{aligned} z_{b,I} &= g_b \odot z_{b,T} \odot z_{b,P} + (1 - g_b) \odot z_{b,T} \odot z_{b,M} \in \mathbb{R}^{128}, \\ z_b &= [z_{b,G} \| z_{b,T} \| z_{b,P} \| z_{b,I}] \in \mathbb{R}^{512}. \end{aligned} \tag{17}$$

Each latent dimension can emphasize peptide or MHC compatibility. MHC enters through the gate and TCR–MHC product, without independent concatenation.

A three-layer binding MLP ϕ_b with dimensions $512 \to 128 \to 128 \to 1$ maps z_b to the binding logit s_b:

$$s_b = \phi_b(z_b), \qquad \hat{y}_b = \Pr(y_b = 1 \mid G_b) = \sigma(s_b). \tag{18}$$

Hidden layers use GELU, LayerNorm, and dropout. Losses use s_b directly; evaluation and inference use its sigmoid.

4.5 Training Objective and Optimization

Training aligns three complementary signals: pointwise binding classification, retention of local geometric evidence, and within-batch ranking:

$$\mathcal{L} = \mathcal{L}_{\text{BCE}} + \lambda_{contact}\mathcal{L}_{\text{contact}} + \lambda_{rank}\mathcal{L}_{\text{rank}}, \tag{19}$$

The auxiliary decoder minimizes $\mathcal{L}_{\text{contact}}$ by reconstructing the 5-Å edge-proximity indicator from endpoint states and the projected edge representation. Following the pairwise-margin ranking principle (Joachims, 2002), $\mathcal{L}_{\text{rank}}$ applies a hinge margin to all available positive–negative logit pairs. We set $\lambda_{contact} = 0.05$ and $\lambda_{rank} = 0.20$. Appendix A specifies the decoder and losses.

Five folds with five random seeds yielded 25 independently trained models using identical fold assignments and minimum-atom-distance residue graphs. Means and standard deviations summarize the 25 test evaluations. AdamW used an initial learning rate of 1×10^{-4}, weight decay 1×10^{-4}, momentum coefficients $(0.9, 0.999)$, and numerical stability constant 10^{-8}. Training used global gradient-norm clipping at 5.0, batch size 64, and at most 10 epochs. After one epoch without validation-AUPRC improvement, the learning rate was multiplied by 0.5, to a minimum of 10^{-7}. Early stopping followed four consecutive epochs without improvement.

5 RESULTS

The experiments connect predictive performance to HIFT-TCR's modeling of the molecular interface. We first evaluate screening across unseen peptide families and examine receptor familiarity. Retrained ablations and attention-family comparisons then assess the model components, while frozen-checkpoint interventions identify the inputs supporting its predictions.

5.1 BENCHMARK PERFORMANCE UNDER PEPTIDE-FAMILY SHIFT

Table 1: Characterized-library screening across completely held-out sequence-defined peptide components. Values are mean ± sample standard deviation over 25 matched seed–fold runs. Best values are bold, second-best values are underlined.

Method	AUROC	AUPRC	pAUC@0.1
HIFT-TCR	**0.5874 ± 0.0555**	**0.2679 ± 0.0684**	**0.5112 ± 0.0190**
STAG	<u>0.5122 ± 0.0546</u>	<u>0.2251 ± 0.0542</u>	<u>0.4996 ± 0.0108</u>
ERGO-II LSTM	0.5021 ± 0.0637	0.2166 ± 0.0618	0.4939 ± 0.0085
TCR-ESM	0.4975 ± 0.0491	0.2200 ± 0.0523	0.4975 ± 0.0100
ERGO-II AE	0.4726 ± 0.0743	0.2068 ± 0.0716	0.4916 ± 0.0072
STAG-LLM	0.4674 ± 0.0435	0.2004 ± 0.0504	0.4915 ± 0.0089
NetTCR-2.2	0.4404 ± 0.0829	0.2071 ± 0.0619	0.4974 ± 0.0084

We first compare HIFT-TCR with sequence, language-model, and structure baselines on the same frozen peptide-component partitions as shown in Table 1. Each entry aggregates metrics computed separately from the saved test predictions of each matched seed–fold run. AUROC evaluates overall ranking, AUPRC summarizes precision as positive examples are recovered, and standardized pAUC evaluates the low-FPR region. They assess candidate ordering across the score range and when false positives are restricted.

HIFT-TCR leads all three mean metrics, with STAG providing the strongest baseline throughout. Its AUROC is 0.5874, and its AUPRC of 0.2679 exceeds the 21.8% prevalence reference. The higher means span overall discrimination and low-FPR ranking. Excluding close peptide relatives across partitions makes this a test of prioritization against unseen peptide families within a characterized receptor library.

5.2 RECEPTOR-FAMILIARITY ANALYSIS

Table 2: Receptor familiarity after seed ensembling. Fold mean ± sample SD. Nonidentical denotes no exact normalized paired-CDR3 training match. Counts are unique samples, with positives in parentheses.

Stratum	n (positive)	AUROC	AUPRC	pAUC@0.1
Nonidentical receptor	7,728 (1,612)	0.6367 ± 0.0555	0.3140 ± 0.0540	0.5351 ± 0.0348
Exact-pair familiar	11,554 (2,535)	0.5596 ± 0.0583	0.2487 ± 0.0650	0.5019 ± 0.0118

We then separate test samples by whether their exact normalized paired-CDR3 sequence occurs in the corresponding training fold as shown in Table 2. Predictions are first averaged over repeated sample identifiers within each seed, then across seeds for each unique sample. Metrics are computed for each stratum within each fold and summarized across folds. Thus, repeated identifiers and optimization seeds do not count as additional observations.

Nonidentical receptors achieve 0.6367 AUROC compared with 0.5596 for exact-pair-familiar receptors, with the same ordering in AUPRC and pAUC. Predictive information therefore persists without an exact receptor-pair training match, supporting screening beyond exact-pair recurrence.

5.3 ARCHITECTURAL COMPONENT ABLATIONS

To investigate the contributions of each component of the model, we conduct an ablation study to verify the effect of removing components and its impact to the model performance. Figure 2

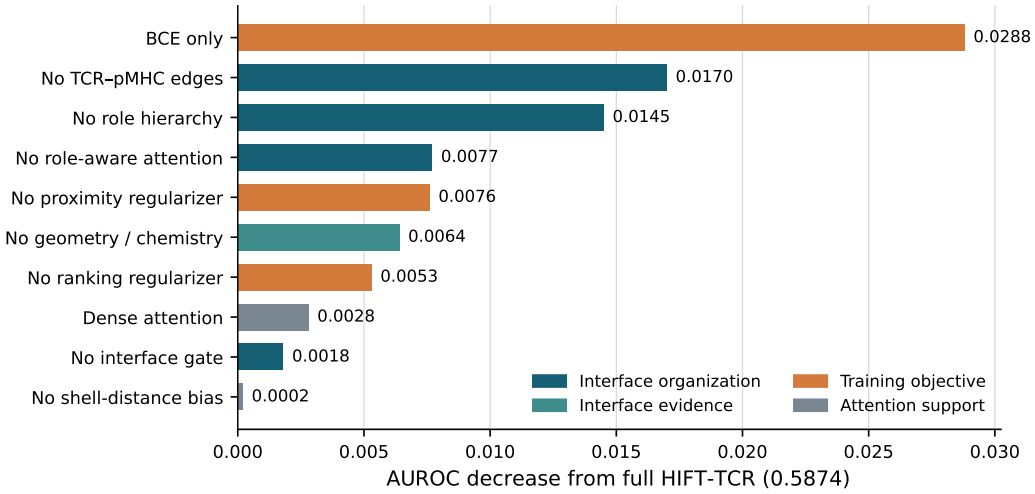


Figure 2: Architectural and objective ablations of HIFT-TCR. Bars show the decrease in mean AUROC from the full model (0.5874) to each retrained variant across 25 matched seed–fold runs. Positive differences favor HIFT-TCR.

summarizes component contributions through retraining, with complete metrics in Table 4. Each bar subtracts a variant's mean AUROC from the full-model mean under the matched seed–fold protocol. The models can adapt during retraining, so the bars measure performance after a component is removed rather than the immediate response of a fixed checkpoint.

The largest decreases follow BCE-only training (0.0288), removal of TCR–pMHC edges (0.0170), and removal of molecular hierarchy (0.0145). The latter changes remove cross-partner graph connections and molecular-role and whole-complex context aggregation, respectively. Their effects extend across metrics. Removing TCR–pMHC edges lowers AUPRC from 0.2679 to 0.2524 and pAUC@0.1 from 0.5112 to 0.5043. Removing hierarchy yields 0.2547 AUPRC and 0.5062 pAUC@0.1. BCE-only training has the lowest AUROC and AUPRC among these component ablations.

Removing shell-distance bias changes AUROC by 0.0002 while retaining the finite-hop attention support. Interface-gate, individual auxiliary-term, and role-aware-attention effects are smaller or metric dependent. These results support the combined objective, interface topology, and hierarchy as contributors to the integrated model. These ablations distinguish the effect of the learned shell-distance bias from the effect of changing the attention receptive field.

5.4 Transformer-Family Controls

Table 3: Attention-support variants under matched training and evaluation. Values are mean ± sample SD over 25 seed–fold runs. Bold and underlined entries mark the highest and second-highest means.

Model family	AUROC	AUPRC	pAUC@0.1
HIFT-TCR (RISA-2)	**0.5874±0.0555**	**0.2679±0.0684**	<u>0.5112±0.0190</u>
RISA-1	<u>0.5853±0.0551</u>	0.2679±0.0671	**0.5124±0.0174**
RISA-3	0.5826±0.0578	0.2640±0.0680	0.5095±0.0192
Dense transformer	0.5762±0.0511	0.2588±0.0675	0.5079±0.0146
Shared-hop sparse	0.5727±0.0507	0.2569±0.0632	0.5079±0.0142

Beyond removing shell-distance bias, we compare attention support while fixing data splits, optimization, local interface field, hierarchy, readout, and objective. RISA-K cycles head radii from one through K hops: RISA-1 uses one-hop heads, RISA-2 alternates one- and two-hop heads, and RISA-3 adds three-hop heads. Dense attention removes the finite-hop restriction whereas shared-hop attention gives every head two-hop support.

RISA-2 leads mean AUROC and AUPRC whereas RISA-1 leads pAUC. AUROC and AUPRC ordering favors complementary immediate and second-shell support over uniformly local or broader attention. RISA-2 exceeds dense attention by 0.0112 AUROC, 0.0091 AUPRC, and 0.0033 pAUC@0.1.

5.5 Interventional Sensitivity to Molecular Inputs

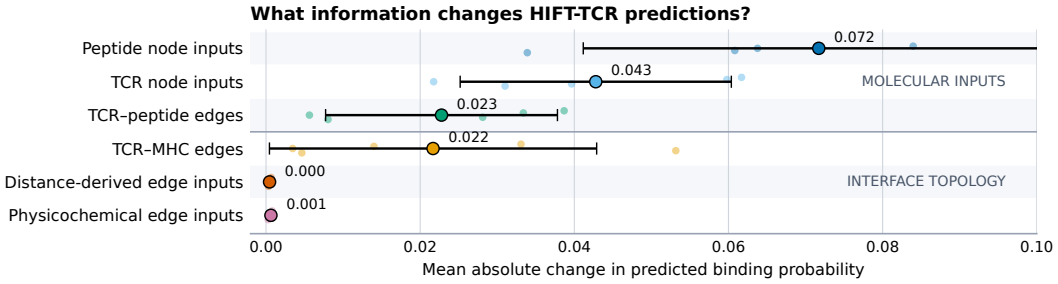


Figure 3: Frozen-checkpoint input sensitivity. Mean absolute change in binding probability. Small points average across 5 seeds within each outer fold. Large circles and horizontal bars show the across-fold mean ± sample SD.

We complement the retrained ablations with post-hoc perturbation analysis to identify which molecular inputs influence predictions, without implying biological causality. Figure 3 examines post-training sensitivity by comparing each frozen checkpoint's predictions on matched test samples before and after an intervention. Interventions zero selected node or edge features, or remove specified bidirectional partner edges. Changes are averaged across seeds within each fold and reported as the mean ± sample standard deviation across folds.

Peptide and TCR inputs produce the largest absolute changes, 0.0717 ± 0.0306 and 0.0428 ± 0.0176, respectively. Absolute change captures sensitivity regardless of direction. To determine which inputs support positive predictions, we also examine the unperturbed probability minus the perturbed probability on positive examples. Zeroing TCR nodes and removing TCR–peptide edges lower these scores by 0.0364 ± 0.0246 and 0.0195 ± 0.0169, respectively. Together with the retrained edge ablation, this consistent direction connects predictive performance to the modeled receptor–peptide interface.

6 Discussion and Conclusion

HIFT-TCR's leading mean AUROC, AUPRC, and low-FPR pAUC demonstrate gains in both overall discrimination and ranking at low false-positive rates. Holding out entire peptide families makes these gains directly relevant to selecting characterized-library receptors for new peptide families. The result supports learning from modeled interfaces with molecular-partner identity preserved throughout.

The ablations connect performance to how interface evidence is organized. Cross-partner edges carry local interaction evidence, while hierarchy places residue states in molecular-role and whole-complex context. Their AUROC effects, together with that of the combined objective, exceed the effect of shell-distance bias alone. This pattern supports coupling interface connectivity and molecular context with training that combines binding classification, proximity self-reconstruction, and ranking. RISA organizes attention over complementary local neighborhoods, while the readout keeps TCR–peptide and TCR–MHC compatibility distinct before fusion.

Higher performance without exact paired-CDR3 training matches shows that predictive signal extends beyond exact recurrence. Agreement between retrained edge ablations and frozen-checkpoint interventions links the performance contribution of interface topology to its use in positive-example predictions. Peptide and TCR inputs also produce the largest absolute sensitivities. HIFT-TCR thus integrates residue interactions, molecular context, and partner compatibility into a framework for prioritizing receptor candidates for experimental binding validation against new peptide families.

AI USE STATEMENT

Generative AI tools assisted with manuscript drafting and editing. The authors have reviewed this manuscript and take responsibility for the final content.

ACKNOWLEDGMENTS

This work was supported by AcRF Tier-2 grant MOE-T2EP20224-0004 of Ministry of Education, Singapore.

REFERENCES

Dmitry V. Bagaev, Renske M. A. Vroomans, Jerome Samir, Ulrik Stervbo, Cristina Rius, Garry Dolton, Alexander Greenshields-Watson, Meriem Attaf, Evgeny S. Egorov, Ivan V. Zvyagin, Nina Babel, David K. Cole, Andrew J. Godkin, Andrew K. Sewell, Can Kesmir, Dmitriy M. Chudakov, Fabio Luciani, and Mikhail Shugay. VDJdb in 2019: Database extension, new analysis infrastructure and a T-cell receptor motif compendium. *Nucleic Acids Research*, 48(D1):D1057–D1062, 2020. doi: 10.1093/nar/gkz874.

Amalie Kai Bentzen, Andrea Marion Marquard, Rikke Lyngaa, Sunil Kumar Saini, Sofie Ramskov, Marco Donia, Lina Such, Andrew J. S. Furness, Nicholas McGranahan, Rachel Rosenthal, Per thor Straten, Zoltan Szallasi, Inge Marie Svane, Charles Swanton, Sergio A. Quezada, Søren Nyboe Jakobsen, Aron Charles Eklund, and Sine Reker Hadrup. Large-scale detection of antigen-specific T cells using peptide–MHC-I multimers labeled with DNA barcodes. *Nature Biotechnology*, 34(10):1037–1045, 2016. doi: 10.1038/nbt.3662.

Judith Bernett, David B. Blumenthal, Dominik G. Grimm, Florian Haselbeck, Roman Joeres, Olga V. Kalinina, and Markus List. Guiding questions to avoid data leakage in biological machine learning applications. *Nature Methods*, 21:1444–1453, 2024. doi: 10.1038/s41592-024-02362-y.

Pradyot Dash, Andrew J. Fiore-Gartland, Tomer Hertz, George C. Wang, Shalini Sharma, Aisha Souquette, Jeremy Chase Crawford, E. Birdie Clemens, Thi H. O. Nguyen, Katherine Kedzierska, Nicole L. La Gruta, Philip Bradley, and Paul G. Thomas. Quantifiable predictive features define epitope-specific T cell receptor repertoires. *Nature*, 547:89–93, 2017. doi: 10.1038/nature22383.

Felix Drost, Anna Chernysheva, Mahmoud Albahah, Katharina Kocher, Kilian Schober, and Benjamin Schubert. Benchmarking of T cell receptor-epitope predictors with ePytope-TCR. *Cell Genomics*, 5(8):100946, 2025. doi: 10.1016/j.xgen.2025.100946.

James Dunbar and Charlotte M. Deane. ANARCI: Antigen receptor numbering and receptor classification. *Bioinformatics*, 32(2):298–300, 2016. doi: 10.1093/bioinformatics/btv552.

Yicheng Gao, Yuli Gao, Yuxiao Fan, Chengyu Zhu, Zhiting Wei, Chi Zhou, Guohui Chuai, Qinchang Chen, He Zhang, and Qi Liu. Pan-peptide meta learning for T-cell receptor–antigen binding recognition. *Nature Machine Intelligence*, 5:236–249, 2023. doi: 10.1038/s42256-023-00619-3.

Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1263–1272, 2017. URL https://proceedings.mlr.press/v70/gilmer17a.html.

Jacob Glanville, Huang Huang, Allison Nau, Olivia Hatton, Lisa E. Wagar, Florian Rubelt, Xuhuai Ji, Arnold Han, Sheri M. Krams, Christina Pettus, Nikhil Haas, Cecilia S. Lindestam Arlehamn, Alessandro Sette, and Scott D. Boyd. Identifying specificity groups in the T cell receptor repertoire. *Nature*, 547:94–98, 2017. doi: 10.1038/nature22976.

R. Grantham. Amino acid difference formula to help explain protein evolution. *Science*, 185(4154): 862–864, 1974. doi: 10.1126/science.185.4154.862.

James M. Heather, Matthew J. Spindler, Marta Herrero Alonso, Yifang Ivana Shui, David G. Millar, David S. Johnson, Mark Cobbold, and Aaron N. Hata. Stitchr: Stitching coding TCR nucleotide sequences from V/J/CDR3 information. *Nucleic Acids Research*, 50(12):e68, 2022. doi: 10.1093/nar/gkac190.

Kamilla Kjærgaard Jensen, Vasileios Rantos, Emma Christine Jappe, Tobias Hegelund Olsen, Martin Closter Jespersen, Vanessa Jurtz, Leon Eyrich Jessen, Esteban Lanzarotti, Swapnil Mahajan, Bjoern Peters, Morten Nielsen, and Paolo Marcatili. TCRpMHCmodels: Structural modelling of TCR-pMHC class I complexes. *Scientific Reports*, 9:14530, 2019. doi: 10.1038/s41598-019-50932-4.

Mathias Fynbo Jensen and Morten Nielsen. Enhancing TCR specificity predictions by combined pan- and peptide-specific training, loss-scaling, and sequence similarity integration. *eLife*, 12, 2024. doi: 10.7554/eLife.93934.3.

Thorsten Joachims. Optimizing search engines using clickthrough data. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 133–142, 2002. doi: 10.1145/775047.775067.

Vadim K. Karnaukhov, Dmitrii S. Shcherbinin, Anton O. Chugunov, Dmitriy M. Chudakov, Roman G. Efremov, Ivan V. Zvyagin, and Mikhail Shugay. Structure-based prediction of T cell receptor recognition of unseen epitopes using TCRen. *Nature Computational Science*, 4:510–521, 2024. doi: 10.1038/s43588-024-00653-0.

Jack Kyte and Russell F. Doolittle. A simple method for displaying the hydropathic character of a protein. *Journal of Molecular Biology*, 157:105–132, 1982. doi: 10.1016/0022-2836(82)90515-0.

Marie-Paule Lefranc, Christelle Pommié, Manuel Ruiz, Véronique Giudicelli, Elodie Foulquier, Lisa Truong, Valérie Thouvenin-Contet, and Gérard Lefranc. Imgt unique numbering for immunoglobulin and t cell receptor variable domains and ig superfamily v-like domains. *Developmental and Comparative Immunology*, 27(1):55–77, 2003. ISSN 0145-305X. doi: https://doi.org/10.1016/S0145-305X(02)00039-3.

Xiangyi Li, Chuance Sun, Weiran Huang, Yanjing Wang, and Buyong Ma. SageTCR: A structure-based model integrating residue- and atom-level representations for enhanced TCR-pMHC binding prediction. *Briefings in Bioinformatics*, 26(5), 2025. doi: 10.1093/bib/bbaf496.

Yanping Lu, Yuyan Wang, Meng Xu, Bingbing Xie, Yumeng Yang, Haodong Xu, and Shengbao Suo. Assessment of computational methods in predicting TCR–epitope binding recognition. *Nature Methods*, 23:248–259, 2026. doi: 10.1038/s41592-025-02910-0.

Tina Lusiany, Zichang Xu, Dianita S. Saputri, Hendra S. Ismanto, Sedat Aybars Nazlica, and Daron M. Standley. Structural modeling of adaptive immune responses to infection. In *Computer-Aided Antibody Design*, volume 2552 of *Methods in Molecular Biology*, pp. 283–294. Springer US, 2022. doi: 10.1007/978-1-0716-2609-2_15.

Jose Mateus, Alba Grifoni, Alison Tarke, John Sidney, Sydney I. Ramirez, Jennifer M. Dan, Zoe C. Burger, Stephen A. Rawlings, Davey M. Smith, Elizabeth Phillips, Simon Mallal, Marshall Lammers, Paul Rubiro, Lorenzo Quiambao, Aaron Sutherland, Esther Dawen Yu, Ricardo da Silva Antunes, Jason Greenbaum, April Frazier, Alena J. Markmann, Lakshmanane Premkumar, Aravinda de Silva, Bjoern Peters, Shane Crotty, Alessandro Sette, and Daniela Weiskopf. Selective and cross-reactive SARS-CoV-2 T cell epitopes in unexposed humans. *Science*, 370 (6512):89–94, 2020. doi: 10.1126/science.abd3871.

Barthelemy Meynard-Piganeau, Christoph Feinauer, Martin Weigt, Aleksandra M. Walczak, and Thierry Mora. TULIP: A transformer-based unsupervised language model for interacting peptides and T cell receptors that generalizes to unseen epitopes. *Proceedings of the National Academy of Sciences*, 121(24), 2024. doi: 10.1073/pnas.2316401121.

Pieter Meysman, Justin Barton, Barbara Bravi, Liel Cohen-Lavi, Vadim Karnaukhov, Elias Lilleskov, Alessandro Montemurro, Morten Nielsen, Thierry Mora, Paul Pereira, Anna Postovskaya, María Rodríguez Martínez, Jorge Fernandez-de Cossio-Diaz, Alexandra Vujkovic, Aleksandra M. Walczak, Anna Weber, Rose Yin, Anne Eugster, and Virag Sharma. Benchmarking solutions to the T-cell receptor epitope prediction problem: IMMREP22 workshop report. *ImmunoInformatics*, 9, 2023. doi: 10.1016/j.immuno.2023.100024.

Ladislav Rampášek, Michael Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Dominique Beaini. Recipe for a general, powerful, scalable graph transformer. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 14501–14515. Curran Associates, Inc., 2022. doi: 10.52202/068431-1054.

Aaron P. Rapoport, Edward A. Stadtmauer, Gwendolyn K. Binder-Scholl, Olga Goloubeva, Dan T. Vogl, Simon F. Lacey, Ashraf Z. Badros, Alfred Garfall, Brendan Weiss, Jeffrey Finklestein, Irina Kulikovskaya, Sanjoy K. Sinha, Shari Kronsberg, Minnal Gupta, Sarah Bond, Luca Melchiori, Joanna E. Brewer, Alan D. Bennett, Andrew B. Gerry, Nicholas J. Pumphrey, Daniel Williams, Helen K. Tayton-Martin, Lilliam Ribeiro, Tom Holdich, Saul Yanovich, Nancy Hardy, Jean Yared, Naseem Kerr, Sunita Philip, Sandra Westphal, Don L. Siegel, Bruce L. Levine, Bent K. Jakobsen, Michael Kalos, and Carl H. June. NY-ESO-1-specific TCR-engineered T cells mediate sustained antigen-specific antitumor effects in myeloma. *Nature Medicine*, 21:914–921, 2015. doi: 10.1038/nm.3910.

Jamie Rossjohn, Stephanie Gras, John J. Miles, Stephen J. Turner, Dale I. Godfrey, and James McCluskey. T cell antigen receptor recognition of antigen-presenting molecules. *Annual Review of Immunology*, 33:169–200, 2015. doi: 10.1146/annurev-immunol-032414-112334.

Paopit Siriarchawatana, Supawadee Ingsriswang, Challika Kaewborisuth, and Anan Jongkaewwattana. TCRLens: Structure-aware equivariant graph learning for TCR-pMHC-I recognition and immunogenic epitope discovery. *Bioinformatics Advances*, 6(1), 2026. doi: 10.1093/bioadv/vbag066.

Jared K. Slone, Anja Conev, Mauricio M. Rigo, Alexandre Reuben, and Lydia E. Kavraki. TCR-pMHC binding specificity prediction from structure using graph neural networks. *IEEE Transactions on Computational Biology and Bioinformatics*, 22(1):171–179, 2025a. doi: 10.1109/TCBBIO.2024.3504235.

Jared K. Slone, Minying Zhang, Peixin Jiang, Amanda Montoya, Emily Bontekoe, Barbara Nassif Rausseo, Alexandre Reuben, and Lydia E. Kavraki. STAG-LLM: Predicting TCR-pHLA binding with protein language models and computationally generated 3d structures. *Computational and Structural Biotechnology Journal*, 27:3885–3896, 2025b. doi: 10.1016/j.csbj.2025.09.004.

Ido Springer, Nili Tickotsky, and Yoram Louzoun. Contribution of T cell receptor alpha and beta CDR3, MHC typing, V and J genes to peptide binding prediction. *Frontiers in Immunology*, 12, 2021. doi: 10.3389/fimmu.2021.664514.

Nili Tickotsky, Tal Sagiv, Jaime Prilusky, Eric Shifrut, and Nir Friedman. McPAS-TCR: A manually curated catalogue of pathology-associated T cell receptor sequences. *Bioinformatics*, 33(18):2924–2929, 2017. doi: 10.1093/bioinformatics/btx286.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.

Randi Vita, James A. Overton, Jason A. Greenbaum, Julia Ponomarenko, Jason D. Clark, Jason R. Cantrell, Daniel K. Wheeler, Joseph L. Gabbard, Deborah Hix, Alessandro Sette, and Bjoern Peters. The immune epitope database (IEDB) 3.0. *Nucleic Acids Research*, 43(D1):D405–D412, 2015. doi: 10.1093/nar/gku938.

Ian Walsh, Dmytro Fishman, Dario Garcia-Gasulla, Tiina Titma, Gianluca Pollastri, Jennifer ELIXIR Machine Learning Focus, Harrow, Fotis E. Psomopoulos, and Silvio C. E. Tosatto. DOME: Recommendations for supervised machine learning validation in biology. *Nature Methods*, 18:1122–1127, 2021. doi: 10.1038/s41592-021-01205-4.

Yue Wan, Chang-Yu Hsieh, Ben Liao, and Shengyu Zhang. Retroformer: Pushing the limits of end-to-end retrosynthesis transformer. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 22475–22490, 2022.

Shashank Yadav, Dhvani Sandip Vora, Durai Sundar, and Jaspreet Kaur Dhanjal. TCR-ESM: Employing protein language embeddings to predict TCR-peptide-MHC binding. *Computational and Structural Biotechnology Journal*, 23:165–173, 2024. doi: 10.1016/j.csbj.2023.11.037.

Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. Do transformers really perform badly for graph representation? In *Advances in Neural Information Processing Systems*, volume 34, pp. 28877–28888, 2021.

Sanggeon Yun, Raheeb Hassan, Ryozo Masukawa, Sungheon Jeong, and Mohsen Imani. HopFormer: Sparse graph transformers with explicit receptive field control. *arXiv preprint arXiv:2602.02268*, 2026. doi: 10.48550/arXiv.2602.02268.

A. A. Zamyatnin. Protein volume in solution. *Progress in Biophysics and Molecular Biology*, 24: 107–123, 1972. doi: 10.1016/0079-6107(72)90005-3.

Yumeng Zhang, Zhikang Wang, Yunzhe Jiang, Dene R. Littler, Mark Gerstein, Anthony W. Purcell, Jamie Rossjohn, Hong-Yu Ou, and Jiangning Song. Epitope-anchored contrastive transfer learning for paired CD8+ T cell receptor–antigen recognition. *Nature Machine Intelligence*, 6: 1344–1358, 2024. doi: 10.1038/s42256-024-00913-8.

Yunxiang Zhao, Jijun Yu, Yixin Su, You Shu, Enhao Ma, Jing Wang, Shuyang Jiang, Congwen Wei, Dongsheng Li, Zhen Huang, Gong Cheng, Hongguang Ren, and Jiannan Feng. A unified deep framework for peptide–major histocompatibility complex–T cell receptor binding prediction. *Nature Machine Intelligence*, 7:650–660, 2025. doi: 10.1038/s42256-025-01002-0.

A COMPLETE HIFT-TCR MATHEMATICAL SPECIFICATION

This appendix specifies the architecture that generated the reported predictions. In particular, it follows the archived checkpoint tensor layout: two RISA bias bins per head and a 384-dimensional auxiliary edge decoder.

A.1 TRAINING-ONLY EDGE-PROXIMITY SELF-RECONSTRUCTION

For each retained directed edge $j \to i$, where j is the source and i is the destination, the auxiliary decoder receives the final endpoint states and the projected edge representation:

$$\kappa_{ji} = \phi_c\Big(h_j^{(3)} \| h_i^{(3)} \| \bar{e}_{ji}\Big), \qquad \hat{t}_{ji} = \sigma(\kappa_{ji}), \tag{20}$$

where $\kappa_{ji} \in \mathbb{R}$ is the edge-proximity logit and $\hat{t}_{ji}$ is its corresponding probability. The decoder has dimensions

$$\phi_c : \mathbb{R}^{384} \longrightarrow \mathbb{R}^{128} \longrightarrow \mathbb{R}. \tag{21}$$

Its hidden stage applies an affine transformation, GELU, LayerNorm, and dropout with probability 0.20; the output stage is affine. The supervision target is

$$t_{ji} = \mathbf{1}[d_{ji} \le 5 \text{ Å}], \tag{22}$$

where d_{ji} is the minimum heavy-atom distance between the two residues.

The edge representation $\bar{e}_{ji}$ is projected from all 22 raw edge channels, including the target-bearing 5-Å proximity indicator. Consequently, this branch is an edge self-reconstruction regularizer rather than independent geometric supervision. No edge channel is masked, and no second encoder pass is performed. The decoder reuses the endpoint states from the primary binding encoder.

The auxiliary loss propagates gradients through the contact decoder, the shared edge projection, and the shared graph encoder. The decoder is executed only during training and never contributes directly to the binding logit.

A.2 COMPLETE TRAINING OBJECTIVE

For a batch containing B complexes, the binding-classification loss is

$$\mathcal{L}_{\text{BCE}} = -\frac{1}{B}\sum_{b=1}^{B}\left[y_b \log \sigma(s_b) + (1-y_b)\log(1-\sigma(s_b))\right]. \tag{23}$$

This term is implemented using binary cross-entropy with logits for numerical stability. No positive-class weighting is applied to the binding labels.

The auxiliary contact loss is evaluated over the retained directed-edge set $\mathcal{E}$. Positive 5-Å contacts receive weight $w^{+}_{\text{contact}} = 2$:

$$\mathcal{L}_{\text{contact}} = -\frac{1}{|\mathcal{E}|}\sum_{(j,i)\in\mathcal{E}}\left[2t_{ji}\log\sigma(\kappa_{ji}) + (1-t_{ji})\log(1-\sigma(\kappa_{ji}))\right]. \tag{24}$$

This term is likewise implemented using binary cross-entropy with logits. If the retained candidate edge set is empty, the implementation sets $\mathcal{L}_{\text{contact}} = 0$.

Let

$$\mathcal{I}_+ = \{b : y_b = 1\}, \qquad \mathcal{I}_- = \{b : y_b = 0\}. \tag{25}$$

The ranking objective is computed over the Cartesian product of all positive and negative examples in the batch:

$$\mathcal{L}_{\text{rank}} = \frac{1}{|\mathcal{I}_+|\,|\mathcal{I}_-|}\sum_{p\in\mathcal{I}_+}\sum_{n\in\mathcal{I}_-}\max(0,\, m_{\text{rank}} - s_p + s_n), \qquad m_{\text{rank}} = 0.20. \tag{26}$$

Where $m_{\text{rank}} = 0.20$. A positive–negative pair incurs no penalty when $s_p \ge s_n + 0.20$. If either class is absent from the minibatch, no valid pair can be formed and the implementation sets $\mathcal{L}_{\text{rank}} = 0$. The ranking loss is global within each minibatch and is not peptide- or query-conditioned.

During validation and testing, the auxiliary edge-proximity decoder is disabled and neither the contact nor ranking loss is evaluated. Operational inference requires only the binding logit and probability $\hat{y}_b = \sigma(s_b)$.

B Additional Architecture and Evidence Results

B.1 Final Model Ablations

Table 4: Matched component ablations of HIFT-TCR, grouped by function. Each variant is retrained under the same protocol. Values are mean ± sample standard deviation over 25 seed–fold runs.

Model variant	AUROC	ΔAUROC	AUPRC	pAUC@0.1
HIFT-TCR (full)	**0.5874±0.0555**	—	**0.2679±0.0684**	**0.5112±.0190**
RISA component				
w/o shell-distance bias	0.5872±0.0554	0.0002	0.2676±0.0678	0.5111±0.0190
Interface organization and evidence				
w/o role-aware attention	0.5797±0.0665	0.0077	0.2675±0.0742	0.5109±0.0266
w/o interface gate	0.5856±0.0563	0.0018	0.2661±0.0658	0.5107±0.0185
w/o geometry/chemistry	0.5810±0.0589	0.0064	0.2631±0.0699	0.5096±0.0168
w/o hierarchy	0.5729±0.0551	0.0145	0.2547±0.0679	0.5062±0.0140
w/o TCR–pMHC edges	0.5704±0.0473	0.0170	0.2524±0.0708	0.5043±0.0098
Training objective				
w/o ranking regularizer	0.5821±0.0541	0.0053	0.2648±0.0624	0.5111±0.0175
w/o proximity regularizer	0.5798±0.0576	0.0076	0.2614±0.0665	0.5083±0.0187
BCE only	0.5586±0.0650	0.0288	0.2494±0.0615	0.5064±0.0124

Among these component ablations, the complete model has the highest mean on all three metrics under the matched protocol. Removing shell-distance bias changes little; clearer losses occur when molecular-role context or the modeled TCR–pMHC interface-edge system is removed. Because the proximity decoder directly receives the target-bearing projected edge representation, this table supports the joint training recipe as an empirical package, and the edge-proximity term is interpreted as representation regularization.