

Active Learning Framework for Cost-Effective TCR-Epitope Binding Affinity Prediction

Pengfei Zhang
Fulton Schools of Engineering (SCAI)
Arizona State University
Tempe, AZ, United States
pzhang84@asu.edu

Seojin Bang
Google DeepMind
Mountain View, CA, United States
seojinb@google.com

Heewook Lee*
Fulton Schools of Engineering (SCAI)
Arizona State University
Tempe, AZ, United States
heewook.lee@asu.edu

Abstract—T cell receptors (TCRs) are critical components of adaptive immune systems, responsible for responding to threats by recognizing epitope sequences presented on host cell surface. Computational prediction of binding affinity between TCRs and epitope sequences using machine/deep learning has attracted intense attention recently. However, its success is hindered by the lack of large collections of annotated TCR-epitope pairs. Annotating their binding affinity requires expensive and time-consuming wet-lab evaluation. To reduce annotation cost, we present *ActiveTCR*, a framework that incorporates active learning and TCR-epitope binding affinity prediction models. Starting with a small set of labeled training pairs, *ActiveTCR* iteratively searches for unlabeled TCR-epitope pairs that are “worthy” for annotation. It aims to maximize performance gains while minimizing the cost of annotation. We compared four query strategies with a random sampling baseline and demonstrated that *ActiveTCR* reduces annotation costs by approximately 40%. Furthermore, we showed that providing ground truth labels of TCR-epitope pairs to query strategies can help identify and reduce more than 40% redundancy among already annotated pairs without compromising model performance, enabling users to train equally powerful prediction models with less training data. Our work is the first systematic investigation of data optimization for TCR-epitope binding affinity prediction.

Index Terms—Active Learning, TCR-epitope Binding Affinity.

I. INTRODUCTION

T cell receptors (TCRs) play a pivotal role in adaptive immune systems by recognizing epitope — a part of antigen—presented on cell surface via major histocompatibility complex and initiating potential immune responses to safeguard the host [1]. Understanding the binding affinity between TCR and epitope sequences is fundamental for developing immunotherapy strategies, where T cells are engineered/created and subsequently assessed for their binding results to target epitopes in a wet-lab setting [2]. However, such assessments are slow and expensive. To overcome these challenges, computational methods for predicting TCR-epitope binding affinity have emerged to streamline the assessment process and minimize expenses.

Many machine learning and deep learning models have been developed to improve the prediction performance of TCR-epitope binding affinity [3]–[6]. These models typically take

two sequences, a TCR and an epitope, as input and predict the binding affinity between them. Despite emergence of these machine learning models, there has been limited attention given to optimizing the underlying data. Most prediction models have primarily focused on exploring the impact of various neural network architectures, neglecting the crucial role of data in achieving optimal performance.

As databases of annotated TCR-epitope pairs continue to grow [7]–[9], two important research questions arise. First, *how can we reduce the cost of annotating new TCR-epitope pairs in the future?* Reducing the cost of each individual wet-lab experiment is challenging due to the inherent expenses associated with these processes. However, more informed decisions can be made by selectively annotating the “most important” or “most useful” pairs for prediction models. By intelligently choosing pairs for annotation, we can optimize the allocation of resources, ultimately leading to more efficient use of annotation budgets for TCR-epitope pairs. Second, *how can we optimize the use of those already annotated pairs to train powerful prediction models with less training data without compromising model performance?* This question arises because of the presence of identical or similar TCR-epitope pairs in the data. Some pairs, although not identical, can be semantically similar in the latent feature space and thus contribute minimally to model training. Furthermore, identical pairs, i.e., TCR-epitope pairs with matching sequences, could cause model overfitting and slow down the training process. A noticeable overlap among these pairs has been observed in most of the currently available databases. For example, roughly 49.41% of TCR-epitope pairs with binding scores greater than zero in VDJdb [8] are identical to pairs in IEDB [9], leading to unnecessary consumption of computational resources. These similar pairs are considered less important or less beneficial to the model. Therefore, our objective is to identify and reduce the redundancy inherent among these already annotated TCR-epitope pairs. By doing so, users can train equally powerful prediction models with less training data.

In this study, we present *ActiveTCR*, a novel framework designed to address research questions of reducing annotation cost for future unlabeled pairs and reducing data redundancy among already labeled data in the context of TCR-epitope

Corresponding author: Heewook Lee (heewook.lee@asu.edu).

binding affinity prediction. *ActiveTCR* employs an active learning approach, where a prediction model is initially trained on a small subset of the training data then the model continuously selects the most informative samples from the unlabeled TCR-epitope pool to be annotated and retrain the prediction model until satisfactory performance is achieved or no more samples are available.

The crux of our framework lies in querying the “most important” pairs for the TCR-epitope binding affinity prediction model. We utilized entropy as a heuristic measure of “importance” of each pair concerning prediction models. To investigate the effectiveness of *ActiveTCR*, we explored five distinct query strategies, including a random sampling baseline, three variants of entropy sampling, and a misclassification sampling strategy for reducing redundancy. We evaluated *ActiveTCR* in two scenarios. Our experimental results demonstrated that *ActiveTCR* effectively halved the annotation cost for future unlabeled TCR-epitope pairs and reduced over 40% data redundancy among those already annotated pairs. *ActiveTCR* provides a promising solution to the challenges of reducing the annotation cost for future data and reducing redundancy among already annotated data for TCR-epitope binding affinity prediction models.¹

II. RELATED WORK

A. Computational Approaches for Binding Affinity Prediction

In order to predict binding affinity between TCR and epitope sequences, researchers have spent large efforts in employing machine learning and deep learning techniques. Two primary approaches explored to improve prediction performance are 1) designing the neural network architecture of prediction models and 2) developing amino acid embedding models for TCR and epitope sequences. Early models focused on neural network structures of prediction models while using a simple embedding matrix, BLOSUM62 [10], to map amino acid residues to continuous numeric vector representations. For example, *NetTCR* [4] used a series of convolutional layers to learn features of input TCR and epitope sequences whereas *ERGO2* [3] proposed an LSTM-based approach to learn sequential information of amino acid sequences. Later, *ATM-TCR* [5] utilized multi-head attention modules to learn contextualized features of amino acids. While these models achieved fair prediction performance on known epitopes (AUCs of 72.0–77.3%), they performed poorly in correctly identifying binding TCRs for novel (unseen) epitopes that were not observed during training (AUCs of 47.0–54.2%).

To develop models that can better generalize to unseen epitopes, researchers incorporated binding affinity prediction models with advanced amino acid embedding techniques. Several such models have been proposed to learn representations from a large corpus of TCR sequences, which can be used as feature extractors for the input sequences of TCR-epitope binding affinity prediction models. *TCRBert* [11]

learned amino acid embeddings using a masked language model, with its architecture inspired by the well-known language model BERT [12]. *catELMo* [13], a model inspired by *ELMo* [14], learned amino acid embeddings by predicting the next token based on its previous tokens processed by a stack of bi-directional LSTM layers. Embeddings from *catELMo* led to significant performance gains for unseen epitope sequences [6], [13].

B. Active Learning

Active learning is a machine learning algorithm that interactively requests annotations of new unlabeled training samples. Unlike passive learning, where a model is trained on a fixed set of labeled data, active learning incorporates human annotators (oracles) to iteratively label the most informative samples. Various query strategies have been developed to identify such informative samples for annotation and inclusion in the training set. Uncertainty sampling [15] selects the samples that a model is most uncertain about. The intuition is that the samples that confuse the model the most will benefit its learning the most. A common example for selecting high uncertainty samples is entropy-based strategy [16], [17]. It selects samples for which the predicted probability distribution over all labels is uniformly spread out. Diversity sampling measures the prediction diversity of unannotated samples, using this diversity as an indicator of informativeness [18]–[20].

By updating the training set, active learning can improve model performance while reducing annotation costs. This approach is particularly useful when labeled data is costly and scarce but unlabeled data is abundant. Active learning has proven effective in a variety of fields, including natural language processing [21], computer vision [22], [23], and medical image analysis [24]. It has also been applied to improve molecular-level optimizations such as prediction of protein-protein interaction [25] and target-drug interaction [26], [27]. To the best of our knowledge, no prior work has explored active learning approaches in the context of TCR-epitope binding affinity prediction.

III. DATA

This section describes how we prepared TCR-epitope pairs and divided them into training and testing sets. As the third complementary-determining region (CDR3) of the TCR β chain is the most critical component that interacts with epitope sequences [28], we made use of the CDR3 of TCR β chain and referred to it as TCR unless otherwise specified.

A. TCR-epitope Pairs

1) *Positive Pairs*: We sourced TCR-epitope pairs of human species from three publicly available databases: *VDJdb* [8], *McPAS* [7], and *IEDB* [9]. These pairs are clinically known to bind to each other and were used as our positive data. We followed the same pre-processing procedure as *ATM-TCR* [5]. The data was filtered to only include pairs of human MHC class I epitopes and TCR β sequences. We kept pairs

¹Code and models are publicly available at <https://github.com/Lee-CBG/ActiveTCR>.

with linear epitope sequences and discarded pairs containing wildcards such as * or X in sequences. For VDJdb [8], we only included pairs with a confidence binding score greater than zero. We also eliminated any duplicated TCR-epitope pairs, resulting in 150,008 unique TCR-epitope pairs, with 982 unique epitopes and 140,675 unique TCRs.

2) *Negative Pairs*: Given the scarcity of clinically confirmed negative TCR-epitope pairs [4], generating negative pairs is a standard practice in the field of TCR-epitope binding prediction [4]–[6], [13]. In order to synthesize negative pairs, we adopted the strategy of pairing existing epitopes with newly sampled TCRs from repertoires, an approach that has been substantiated in prior studies [4], [6], [13]. The rationale behind this approach was to simulate training pairs where TCRs that are prevalent in healthy individuals do not bind with disease epitopes, thereby providing a comprehensive contrast to the positive pairs in our machine learning model. We first randomly sampled TCRs from 20 million TCR sequences of healthy control repertoires of ImmunoSEQ [29]. We then replaced the original TCRs of the positive TCR-epitope pairs with TCRs of the healthy repertoires. This resulted in 150,008 unique negative TCR-epitope pairs.

B. Training and Testing Set Splits

To accurately assess the prediction performance on the novel (unseen) TCRs and epitopes, we used the two data partition approaches in ATM-TCR [5]. The epitope split shared no common epitopes between training and testing sets, allowing us to evaluate the model’s generalizability on novel epitope sequences. Similarly, the TCR split shared no common TCRs between training and testing sets, enabling us to evaluate the model’s generalizability on novel TCR sequences. For both splits, we used 80% of the entire data as training data and the remaining 20% as testing data.

IV. METHODS

A. ActiveTCR

ActiveTCR is an active learning framework designed for rapid and cost-effective development of TCR-epitope binding affinity prediction models. The crux of ActiveTCR lies in its ability to interactively query a large size of pool to iteratively update the training set D by adding the most informative TCR-epitope pairs, thereby learning improved prediction models M . ActiveTCR has two use cases in the context of TCR-epitope binding affinity prediction: 1) reducing annotation costs of unlabeled TCR-epitope pairs and 2) reducing redundancy among already labeled pairs. Detailed procedure is provided in (Supplementary Alg. 1).

The goal of the first use case is to maximize prediction performance gains while minimizing the associated annotation costs. As shown in Fig. 1a, ActiveTCR iteratively queries the most informative TCR-epitope pairs to be annotated by wet-lab efforts. The process begins by training a TCR-epitope binding affinity prediction model M , referred to as a learner, on an initial training set of labeled TCR-epitope pairs D . The learner M then interactively queries unlabeled TCR-epitope pairs

from the pool P_u . In order to do so, it predicts binding affinities of the unlabeled pairs and selects the most informative pairs by a query strategy Q . Wet-lab annotators then provide the ground truth binding labels for the queried TCR-epitope pairs L . The newly labeled pairs are added to the training set D , and the model M is retrained on the updated training set D . Through this feedback loop, ActiveTCR continually updates model M until satisfactory prediction performance is achieved.

The goal of the second use case is to develop a data efficient model that can achieve the same or better prediction performance with a less number of TCR-epitope training pairs. As shown in Fig. 1b, ActiveTCR iteratively queries more labeled pairs L and adds them to the training set D to improve the prediction model M . Redundant pairs may include identical data entries found in different databases or semantically similar entries that do not significantly contribute to the model performance. Since the ground truth label is accessible in the pool P_l , we can leverage it for querying informative pairs and discard redundant pairs. The wet-lab annotators are no longer required in this setting.

B. Query Strategies

a) *Global and local entropy-based sampling*.: Entropy-based sampling [30] queries TCR-epitope pairs that the model is least certain about their binding status. Such pairs are typically located near the classification boundary, making them more informative about shape of the boundary than pairs located farther from it. Entropy [31] is an intuitive metric for assessing the uncertainty of machine learning model predictions. A low entropy score of a sample indicates that the model is more confident in its prediction result, while a high entropy score suggests less confidence. The entropy score of a model M ’s prediction for a pair of TCR and epitope (t_i, e_i) is defined as follows:

$$Q_{EN}(t_i, e_i; M) = H(\hat{y}_i) \\ = -(\hat{y}_i \log_2 \hat{y}_i + (1 - \hat{y}_i) \log_2(1 - \hat{y}_i)) \quad (1)$$

where \hat{y}_i is model M ’s prediction output for the TCR-epitope pair (t_i, e_i) . At each iteration, prediction model M predicts binding affinity \hat{y} for pairs in P_u (or P_l) and selects those with the highest entropy scores to construct new training sets D . We used two entropy-based query strategies: *global entropy sampling* and *local entropy sampling*. Global entropy sampling queries pairs from the entire pooled pairs P_u (or P_l) at each iteration, while local entropy sampling queries pairs from a random subset of the pool, $P_{us} \in P_u$ (or $P_{ls} \in P_l$). In our experiment, we set the size of P_{us} (or P_{ls}) to be twice the size of the query pairs. For example, if 1,000 query pairs are selected at each iteration, the size of P_{us} (or P_{ls}) would be 2,000. Global entropy sampling is more reliable than local entropy sampling as it calculates entropy scores across all unlabeled pairs. Meanwhile, local entropy sampling is faster as it only computes entropy for a subset of pooled pairs.

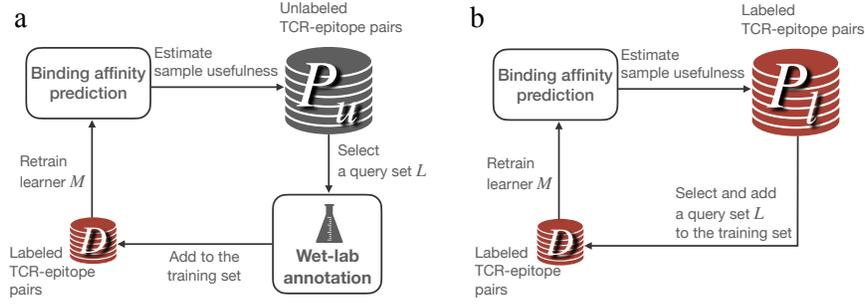


Fig. 1. ActiveTCR framework for TCR-epitope binding affinity prediction. It can be used **a)** to reduce annotation cost for unlabeled TCR-epitope pairs, and **b)** to reduce data redundancy among labeled TCR-epitope pairs. It interactively queries the most informative TCR-epitope pairs to iteratively construct a training set and to update a binding affinity prediction model.

b) Query-by-dropout-committee sampling: Query-by-committee [15] selects samples based on the level of disagreement among multiple prediction models. Each model, known as a committee member, learns its own decision boundary and predicts the binding affinity of a TCR-epitope pair. If the committee members largely disagree on their predictions, the pair is considered uncertain and added to training sets D . This strategy tends to be slower and more computationally demanding because it requires training multiple prediction models to serve as committee members. To address these limitations, researchers have used dropout during inference time to measure committee disagreement in active learning frameworks [19], [20]. This technique enables dropout layers during the prediction phase, introducing variation to the model’s prediction outputs by randomly dropping out neurons. As a result, the same model can produce different binding affinity scores for a pair of TCR and epitope without the need of training multiple committee models. This approach, referred to as dropout committee, is commonly used to measure the uncertainty of a neural network model with reduced computational burden [32].

We quantified the disagreement score for a TCR-epitope pair by a sum of Kullback-Leibler divergence [33] between each test time prediction and their average:

$$Q_C(t_i, e_i; M) = \sum_j D_{KL}(\hat{Y}_{i,j} \parallel \bar{Y}_i) \\ = \sum_j \sum_{k \in \{0,1\}} \hat{y}_{i,j,k} \log_2 \left(\frac{\hat{y}_{i,j,k}}{\bar{y}_{i,k}} \right) \quad (2)$$

where k is possible prediction outcomes (0 for non-binding and 1 for binding in our case), $\hat{y}_{i,j,k}$ is the binding affinity prediction score of (t_i, e_i) made by j -th committee member, and $\bar{y}_{i,k}$ is an average of all committee members’ prediction scores. In our experiment, we set the number of dropout committees as 10.

At each iteration, we applied *dropout committee sampling* to a randomly selected subset of pooled pairs P_{us} (or P_{ls}) for computational efficiency. To investigate whether this query strategy can enhance the performance of entropy-based

queries, we added it to the local entropy strategy, defining the measurement as follows.

$$Q_{EN,C}(t_i, e_i; M) = \omega Q_{EN}(t_i, e_i; M) + (1 - \omega) Q_C(t_i, e_i; M) \quad (3)$$

where ω is a weight between local entropy sampling and dropout committee sampling. For simplicity, we assumed the same importance for those two samplings and assigned equal weights to them in our experiments.

c) Misclassification sampling: We designed a simple yet effective query strategy called *misclassification sampling*. It selects samples from the labeled TCR-epitope pool P_l that are misclassified (both incorrectly labeled as positive and incorrectly labeled as negative) by the model M with a significant difference between the true label and the predicted label. The intuition is that samples that are misclassified by a large margin can potentially improve model performance by providing more information about the model’s weaknesses. The misclassification sampling score of a TCR-epitope pair (t_i, e_i) is defined as follows:

$$MC(t_i, e_i, y_i; M) = |y_i - \hat{y}_i| \text{ for } i \in P_l \quad (4)$$

where y_i is the ground truth binary label and $\hat{y}_i = M(t_i, e_i)$ is the binding affinity prediction score between 0 and 1. Note that this approach is only applicable to the use case of reducing redundancy among annotated TCR-epitope pairs as it requires ground truth labels for them in P_l to determine misclassification.

C. Binding Affinity Prediction Model

Our learner M is a TCR-epitope binding affinity prediction model that predicts the likelihood of binding between a TCR and an epitope. It takes two sequence inputs (a TCR and an epitope) and outputs a probability that the TCR will bind to the epitope. In our active learning framework, users have the flexibilities to select the binding affinity prediction model that best suits their needs, for instance, ATM-TCR [5], ERGO2 [3], NetTCR [4], and PiTE [6].

We used the 3-linear-layered model from catELMo [13] as our prediction model M because it performed similarly to the state-of-the-art performance of PiTE while being faster and more lightweight [6]. It is composed of two main steps:

amino acid embedding and binding affinity prediction. In the first step, the TCR and epitope sequences were represented as numeric vectors of size 1,024 using catELMo, a state-of-the-art amino acid embedding model. In the second step, a simple prediction model with three linearly connected layers was trained to predict the binding affinity between the embedded TCR and epitope sequences. The embedded vectors were used as input for the model. Each was first processed through a linear layer with 2,048 neurons followed by a Sigmoid Linear Units (SiLU) activation function [34], batch normalization [35], and 0.3 rate dropout [36]. The two processed sequences were then concatenated and fed to a linear layer with 1,024 neurons, followed by a SiLU activation function, batch normalization, and 0.3 rate dropout. Finally, the last linear layer with a neuron followed by a sigmoid activation function produced a binding affinity score between 0 and 1. Binary cross-entropy loss and Adam optimizer [37] were used to train the model. The batch size was 32 and the learning rate was 0.001. The training continued for 200 epochs or was early terminated if the validation loss did not decrease for 30 consecutive epochs.

V. RESULTS ON REAL DATA

A. Study Design

In this section, we experimentally assess how ActiveTCR contributes to the annotation cost reduction for unlabeled TCR-epitope pairs and to the data redundancy reduction among labeled pairs.

We randomly selected 10% (about 24,006 pairs) of the training TCR-epitope pairs as the initial training data D . The remaining 90% (about 216,054 pairs) of the training set was treated as pool data (P_u or P_l). Our learner model M predicted the binding affinities \hat{y} and identified the most informative pairs based on the different query strategies. At each iteration, 24,006 additional pairs L were queried from pool data. For the use case of reducing annotation costs of unlabeled data, we removed the binding labels of TCR-epitope pairs in P_u to simulate the presence of future unlabeled pairs. The ground truth labels were provided to the model only after it made queries L , and served as the wet-lab annotation. For the use case of reducing data redundancy, we allowed the query strategies to leverage ground truth label information to select pairs from P_l . ActiveTCR continuously expanded the size of dataset D by adding the queried samples L . The prediction model M was interactively retrained until all pairs in the pool were exhausted or satisfactory results were achieved. The testing set was held out to assess the performance of the model M at each iteration. We reported AUCs of prediction models using different query strategies on the testing set. Each query strategy had 10 independent runs with a seed value of 42 unless otherwise specified.

B. Reduce Annotation Cost by Interactively Annotating Unlabeled Samples

ActiveTCR significantly reduced annotation costs for future unlabeled TCR-epitope pairs compared to the random

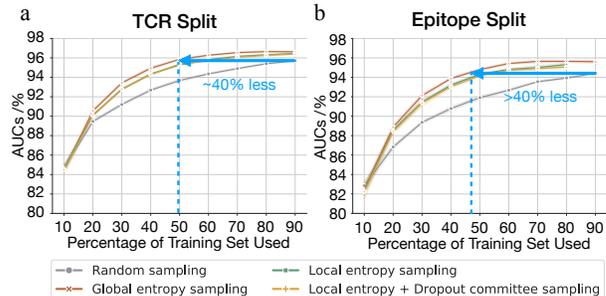


Fig. 2. Performance of ActiveTCR in reducing annotation costs, measured on a) TCR split and b) epitope split. Average (solid line) and standard error (band) AUC of 10 independent runs for each query strategy are reported. ActiveTCR using global entropy sampling reduced approximately 40% annotation costs compared to the random sampling baseline. Two of the query strategies required sub-sampling the unlabeled pool, which was earlier stopped at iteration 8 (80% training set) in epitope split due to insufficient TCR-epitope pairs for sub-sampling.

sampling baseline. We simulated a realistic scenario in which not all labels are available at the beginning. We assumed that only 10% samples are labeled, which is our initial set. We then iteratively added back the “labels” of each additional 10% training samples (to mimic the wet lab procedure of annotating additional TCR-epitope pairs) by ActiveTCR. The key question we addressed in this experiment is whether ActiveTCR, as an active learning framework, can select the most informative TCR-epitope pairs to be annotated, such that it achieves a higher performance gain compared to a random sampling strategy.

Among the strategies, global entropy sampling reduced at least 40% annotation cost in both TCR and epitope splits (Fig. 2). Global entropy sampling, while being more computationally intensive than local entropy sampling, was found to be more effective in reducing annotation cost and improving the model performance. This suggested that samples queried from global entropy were more informative and therefore contributed more to the model’s learning than those queried using local entropy sampling. Local entropy with dropout committee sampling did not appear to offer additional benefits over local entropy sampling in both TCR and epitope splits, indicating that using a dropout layer to introduce uncertainty does not necessarily improve model learning. A possible explanation is that the weight (w in Equation 3) we assigned between them is suboptimal. Misclassification sampling was excluded from the comparison as it required ground truth labels for pairs.

As the model queried new samples with their binding affinity ground truth hidden at each iteration, we did not have control over the ratio of queried positive and negative pairs at each iteration. To understand the preference of each query strategy over positive and negative pairs, we visualized the amount of queried positive and negative pairs in L for each iteration in both TCR and epitope splits (Fig. S1). We found that global entropy sampling, the best-performing strategy, had a relatively uneven distribution of positive and negative pairs at each iteration but maintained a relatively balanced positive-

negative ratio for cumulative pairs D (Fig. S2).

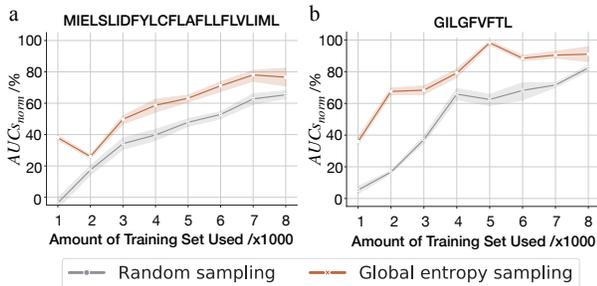


Fig. 3. Epitope-specific performance of ActiveTCR in reducing annotation costs. The performance is measured on epitope split for two individual epitopes **a)** MIELSLIDFYLCFLAFLFLVLIIML and **b)** GILGFVFTL. The average (solid line) and standard error (band) of AUC_{norm} from 5 independent runs for each query strategy are reported. AUC_{none} for MIELSLIDFYLCFLAFLFLVLIIML and GILGFVFTL are 81.80% and 92.98% when fine-tuning M_0 on zero pairs, respectively. AUC_{all} for MIELSLIDFYLCFLAFLFLVLIIML and GILGFVFTL are 89.04% and 95.47% when fine-tuning M_0 on the entire new unlabeled pool, respectively. ActiveTCR consistently outperforms the random sampling baseline with however many TCR-epitope pairs queried and annotated.

We also showed that ActiveTCR consistently improved prediction performance for individual novel epitopes by fine-tuning the initial prediction model M_0 on queried pairs. Such investigation is particularly beneficial when the goal is to optimize the prediction performance of novel or rarely observed target epitopes associated with specific diseases. We emulated a situation where a single laboratory, equipped with an initial prediction model trained on existing TCR-epitope pairs, aims to explore a novel target epitope. This setup aligns with the laboratory’s research interests as it represents a common real-world scenario where labs are often interested in studying novel or disease-specific epitopes. We assumed that the initial prediction model was trained on an existing database of multiple epitopes and TCRs, and that the target epitope was novel to the model. To ensure the target epitope has never been seen by the model, we selected the epitope from the testing set of epitope split. The initial prediction model M_0 was trained on 10% of randomly selected pairs from the epitope split training set. We did not utilize the remaining pairs from the training set as we assumed only a limited number of training pairs are available. We then used 90% of the randomly selected TCR and the target epitope pairs from the epitope split testing set as the unlabeled pool, and the remaining 10% as testing pairs.

We fine-tuned the initial model M_0 on different sizes of query sets $(1, \dots, 8K)$. A 10x smaller learning rate of 0.0001 was used to prevent model overfitting. We compared the best-performing global entropy sampling query with a random sampling baseline and quantified the performance improvements. Fig. 3 shows results for the two most abundant epitopes in the testing set, MIELSLIDFYLCFLAFLFLVLIIML and GILGFVFTL. We reported the normalized AUC (AUC_{norm}) and defined it as follows:

$$AUC_{norm} = (AUC - AUC_{none}) / (AUC_{all} - AUC_{none}) \quad (5)$$

where AUC_{none} is the AUC of the initial model that fine-tuned on zero queried pairs, and AUC_{all} is the AUC of the fine-tuned model on the entire unlabeled pool pairs of an epitope. Our results indicated that ActiveTCR with global entropy sampling consistently outperforms the random sampling baseline in terms of prediction performance for individual epitopes, regardless of the number of TCR-epitope pairs queried and annotated. This experimental setup illustrated how ActiveTCR’s query aligns with a laboratory’s interests and capabilities, assisting in the selection of the most “promising” TCR-epitope pairs for testing, thereby reducing the experimental workload while maximizing the potential discovery of useful knowledge.

C. Identify and Reduce Redundancy among Labeled Samples

ActiveTCR identified and removed at least 40% of training data as redundancy (Fig. 4) while matching the performance of passive learners with random sampling. This significant reduction directly translates into computational savings for future model training tasks. In this experiment, we queried positive and negative pairs separately. This gave us control over the distribution of queried positive and negative pairs. We queried 5% (12,003 pairs) from each at each iteration. ActiveTCR with global entropy sampling with only 50% of the training data achieved similar performance to random sampling with 90% of training data, indicating that more than 40% of the training data were not necessarily contributing to the model.

We noticed that, with misclassification sampling, ActiveTCR struggled to achieve high AUC prediction scores with a larger number of queried training pairs at first. This may be because the model was being fed with “difficult” samples, making it challenging to generalize well on the testing set. However, by the third iteration, the performance score of the method started to improve and even surpassed that of global entropy sampling in the final iterations.

We also observed that global entropy consistently outperformed local entropy in reducing data redundancy, suggesting that samples queried from global entropy are less redundant and could improve the model’s performance more. Despite its relatively lower performance compared to other query strategies, local entropy with the dropout committee sampling was still able to reduce approximately 30% of redundant data in the epitope split and 20% in the TCR split. We speculated that the reason for its suboptimal performance is the challenge of determining the weights of local entropy and the dropout committee sampling method. Nonetheless, this approach can still provide a significant reduction in annotation costs. Overall, we demonstrated that ActiveTCR is effective for reducing redundancy among those TCR-epitope pairs with annotated binding results. It can significantly reduce the amount of training data required to match comparable performance to passive learning.

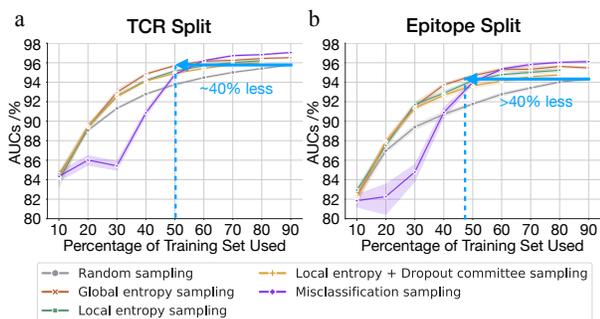


Fig. 4. ActiveTCR performance comparison in reducing computational cost by reducing redundancy for TCR-epitope pairs using five different query strategies for **a)** TCR split and **b)** epitope split. Average (solid line) and standard error (band) AUC of 10 independent runs for each query strategy are reported. ActiveTCR using global entropy sampling reduced approximately 40% redundancy among labeled data compared to the random sampling baseline. Two of the query strategies required sub-sampling the unlabeled pool, which was earlier stopped at iteration 8 (80% training set) in both TCR and epitope split due to insufficient TCR-epitope pairs for sub-sampling.

It should be noted that the goal of this section is to 1) demonstrate that there are many redundant TCR-epitope pairs that do not necessarily add to performance gains of the prediction model and to 2) identify a compact, highly informative subset of the data (which we refer to as the “primal dataset”). When new TCR-epitope pairs are added to a dataset, the “primal data” computed prior to the addition of new pairs can be used as a starting point and querying can be done only from the new pairs. This avoids the need to retrain prediction models from scratch using the entire data including the newly added pairs. We conducted an additional experiment and observed no significant performance differences between utilizing an identified primal dataset and running from scratch. In this experiment, we used 60% of the training data (refer to as A) as initially available TCR-epitope pairs with labels. Running ActiveTCR on this dataset (A), a subset was identified as primal data (A_{Primal}) after removing $A_{Redundant}$, where $A = A_{Primal} \cup A_{Redundant}$. Then, we randomly selected a set of additional TCR-epitope pairs with labels (referred to as B). The total number of pairs in B is one third of the number of pairs in A . Note that A and B are disjoint. With a larger number of labeled pairs, in order to newly define the primal data ($(A \cup B)_{Primal}$), one can either rerun ActiveTCR on the entirety of $A \cup B$, or utilize the previously defined primal data (A_{Primal}) as the starting point and query only from the newly added dataset (B). Utilizing the previously defined primal data only requires iterative model training sweeping through the additional data B . However, training from scratch requires more iterations of model training. This reduction in training iterations grows with each occasion of obtaining new labeled data. In the TCR split, the former approach yielded an average AUC score of 96.09%, while the latter resulted in a score of 95.35%. Similarly, in the epitope split, the former method achieved an AUC of 95.42%, and the latter, 94.41%. These similar performances indicate ActiveTCR allows users to expand upon pre-existing datasets without the need to rerun

ActiveTCR from the beginning.

VI. DISCUSSION

We proposed ActiveTCR, an active learning framework for cost-effective TCR-epitope binding affinity prediction. We investigated five query strategies and demonstrated the advantages of ActiveTCR in two practical scenarios, reducing annotation costs of unlabeled pairs and decreasing computational cost by reducing redundancy among annotated TCR-epitope pairs.

Despite the significant reduction in annotation cost and data redundancy achieved by ActiveTCR, retraining the prediction model M at each iteration is relatively slow. To address this, we investigated an alternative training strategy: fine-tuning of the previous prediction model on newly queried samples L , instead of training a new model on updated training sets D . We found that the fine-tuning method converges much faster but severe overfitting problems were also observed for all query strategies.

As ActiveTCR randomly selected the initial training set, it may cause the “cold start” problem in active learning. This means that different initial training data may lead to different initial TCR-epitope binding affinity prediction models M , and in turn, may affect subsequent models and queried pairs in following iterations. To demonstrate the robustness of our method under different starting points, we conducted additional experiments based on three different random seeds and consistently observed significant amounts of annotation cost reductions (Fig. S3).

One limitation of our study is that our experimental setting may not be a perfect reflection of the real-world distribution of TCR-epitope pairs. The dataset we prepared for this study has a 1:1 ratio of positive and negative pairs, while datasets obtained in real-world scenarios may differ. In a real-world scenario, researchers may focus on improving model performance for epitopes rising from a set of target diseases. Annotating a collection of TCR-epitope pairs based on a set of target epitopes may require analysis of TCR repertoire across disease and control subjects to extract TCRs that are likely to bind to the target epitopes. Consequently, there are likely to be more positive pairs and fewer negative pairs. To demonstrate how the original dataset distribution affects the query sample distributions, we ran ActiveTCR with global entropy sampling on a TCR-epitope dataset with different positive and negative ratios. We prepared the dataset with the ratio of positive and negative pairs as 2:1, 4:1, and 9:1, respectively. We kept the number of positive TCR-epitope pairs constant and randomly selected half, one-fourth, and one-eleventh of the negative pairs, then concatenated them together. Fig. S4 shows the queried sample distributions for different positive-negative ratios. While it is difficult to pinpoint the reasons, we found that the global entropy sampling query strategy attempted to query a more balanced number of positive and negative pairs than random sampling from iterations 1 to 4. This finding highlighted the potential advantage of this query strategy and we believe it deserves further investigation in

future studies. We also observed that the model preferred to query positive pairs at later iterations, possibly because more positive pairs were available and negative pairs had been used up as it iterates. Additionally, the annotated TCR-epitope pairs were not specifically designed for machine learning purposes but were collected from various biological or clinical studies and may contain biases towards certain diseases.

REFERENCES

- [1] M. Attaf, M. Legut, D. K. Cole, and A. K. Sewell, "The T cell antigen receptor: the Swiss army knife of the immune system," *Clinical & Experimental Immunology*, vol. 181, no. 1, pp. 1–18, 2015.
- [2] T. N. Schumacher, "T-cell-receptor gene therapy," *Nature Reviews Immunology*, vol. 2, no. 7, pp. 512–519, 2002.
- [3] I. Springer, N. Tickotsky, and Y. Louzoun, "Contribution of T cell receptor alpha and beta CDR3, MHC typing, V and J genes to peptide binding prediction," *Frontiers in Immunology*, vol. 12, p. 664514, 2021.
- [4] A. Montemurro, V. Schuster, H. R. Povlsen, A. K. Bentzen, V. Jurtz, W. D. Chronister, A. Crinklaw, S. R. Hadrup, O. Winther, B. Peters *et al.*, "NetTCR-2.0 enables accurate prediction of TCR-peptide binding by using paired TCR α and β sequence data," *Communications Biology*, vol. 4, no. 1, p. 1060, 2021.
- [5] M. Cai, S. Bang, P. Zhang, and H. Lee, "ATM-TCR: TCR-epitope binding affinity prediction using a multi-head self-attention model," *Frontiers in Immunology*, vol. 13, 2022.
- [6] P. Zhang, S. Bang, and H. Lee, "PITE: TCR-epitope binding affinity prediction pipeline using Transformer-based sequence encoder," in *PACIFIC SYMPOSIUM ON BIOCOMPUTING 2023: Kohala Coast, Hawaii, USA, 3–7 January 2023*. World Scientific, 2022, pp. 347–358.
- [7] N. Tickotsky, T. Sagiv, J. Prilusky, E. Shifrut, and N. Friedman, "McPAS-TCR: a manually curated catalogue of pathology-associated T cell receptor sequences," *Bioinformatics*, vol. 33, no. 18, pp. 2924–2929, 2017.
- [8] M. Shugay, D. V. Bagaev, I. V. Zvyagin, R. M. Vroomans, J. C. Crawford, G. Dolton, E. A. Komech, A. L. Sycheva, A. E. Koneva, E. S. Egorov *et al.*, "VDJdb: a curated database of T-cell receptor sequences with known antigen specificity," *Nucleic Acids Research*, vol. 46, no. D1, pp. D419–D427, 2018.
- [9] R. Vita, S. Mahajan, J. A. Overton, S. K. Dhanda, S. Martini, J. R. Cantrell, D. K. Wheeler, A. Sette, and B. Peters, "The immune epitope database (IEDB): 2018 update," *Nucleic Acids Research*, vol. 47, no. D1, pp. D339–D343, 2019.
- [10] S. Henikoff and J. G. Henikoff, "Amino acid substitution matrices from protein blocks," *Proceedings of the National Academy of Sciences*, vol. 89, no. 22, pp. 10915–10919, 1992.
- [11] K. Wu, K. E. Yost, B. Daniel, J. A. Belk, Y. Xia, T. Egawa, A. Satpathy, H. Y. Chang, and J. Zou, "TCR-BERT: learning the grammar of T-cell receptors for flexible antigen-xbining analyses," *bioRxiv*, 2021.
- [12] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," Association for Computational Linguistics, Jun. 2019, pp. 4171–4186.
- [13] P. Zhang, S. Bang, M. Cai, and H. Lee, "Context-aware amino acid embedding advances analysis of TCR-epitope interactions," jul 2023. [Online]. Available: <https://doi.org/10.7554/elife.88837.1>
- [14] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, "Deep contextualized word representations." Association for Computational Linguistics, Jun. 2018, pp. 2227–2237.
- [15] H. S. Seung, M. Opper, and H. Sompolinsky, "Query by committee," in *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, 1992, pp. 287–294.
- [16] A. Holub, P. Perona, and M. C. Burl, "Entropy-based active learning for object recognition," in *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. IEEE, 2008, pp. 1–8.
- [17] F. Ozdemir, Z. Peng, C. Tanner, P. Fuernstahl, and O. Goksel, "Active learning for segmentation by optimizing content information for maximal entropy," in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4*. Springer, 2018, pp. 183–191.
- [18] B. Settles, M. Craven, and S. Ray, "Multiple-instance active learning," *Advances in Neural Information Processing Systems*, vol. 20, 2007.
- [19] M. Ducoffe and F. Precioso, "Active learning strategy for CNN combining batchwise dropout and query-by-committee," in *The European Symposium on Artificial Neural Networks*, 2017.
- [20] E. Tsymbalov, M. Panov, and A. Shapeev, "Dropout-based active learning for regression," in *Analysis of Images, Social Networks and Texts: 7th International Conference, AIST 2018, Moscow, Russia, July 5–7, 2018, Revised Selected Papers 7*. Springer, 2018, pp. 247–258.
- [21] A. Siddhant and Z. C. Lipton, "Deep Bayesian active learning for natural language processing: Results of a large-scale empirical study," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Brussels, Belgium: Association for Computational Linguistics, Oct.-Nov. 2018, pp. 2904–2909. [Online]. Available: <https://aclanthology.org/D18-1318>
- [22] E. Haussmann, M. Fenzi, K. Chitta, J. Ivanecky, H. Xu, D. Roy, A. Mittel, N. Koumchatzky, C. Farabet, and J. M. Alvarez, "Scalable active learning for object detection," in *2020 IEEE Intelligent Vehicles Symposium (iv)*. IEEE, 2020, pp. 1430–1435.
- [23] T. Yuan, F. Wan, M. Fu, J. Liu, S. Xu, X. Ji, and Q. Ye, "Multiple instance active learning for object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 5330–5339.
- [24] S. Budd, E. C. Robinson, and B. Kainz, "A survey on active learning and human-in-the-loop deep learning for medical image analysis," *Medical Image Analysis*, vol. 71, p. 102062, 2021.
- [25] T. P. Mohamed, J. G. Carbonell, and M. K. Ganapathiraju, "Active learning for human protein-protein interaction prediction," *BMC Bioinformatics*, vol. 11, no. 1, pp. 1–9, 2010.
- [26] R. F. Murphy, "An active role for machine learning in drug development," *Nature Chemical Biology*, vol. 7, no. 6, pp. 327–330, 2011.
- [27] M. Eisenstein, "Active machine learning helps drug hunters tackle biology," *Nature Biotechnology*, vol. 38, no. 5, pp. 512–515, 2020.
- [28] N. L. La Gruta, S. Gras, S. R. Daley, P. G. Thomas, and J. Rossjohn, "Understanding the drivers of MHC restriction of T cell receptors," *Nature Reviews Immunology*, vol. 18, no. 7, pp. 467–478, 2018.
- [29] S. Nolan, M. Vignali, M. Klinger, J. N. Dines, I. M. Kaplan, E. Svejnova, T. Craft, K. Boland, M. Pesesky, R. M. Gittelman *et al.*, "A large-scale database of T-cell receptor beta (TCR β) sequences and binding associations from natural and synthetic exposure to SARS-CoV-2," *Research Square*, 2020.
- [30] D. D. Lewis, "A sequential algorithm for training text classifiers: Corrigendum and additional data," in *ACM SIGIR Forum*, vol. 29, no. 2. ACM New York, NY, USA, 1995, pp. 13–19.
- [31] C. E. Shannon, "A mathematical theory of communication," *The Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, 1948.
- [32] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *International Conference on Machine Learning*. PMLR, 2016, pp. 1050–1059.
- [33] S. Kullback and R. A. Leibler, "On information and sufficiency," *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, 1951.
- [34] S. Elfving, E. Uchibe, and K. Doya, "Sigmoid-weighted linear units for neural network function approximation in reinforcement learning," *Neural Networks*, vol. 107, pp. 3–11, 2018.
- [35] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International Conference on Machine Learning*. PMLR, 2015, pp. 448–456.
- [36] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [37] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014.