frontiers | Frontiers in Immunology

TYPE Original Research
PUBLISHED 20 January 2023
DOI 10.3389/fimmu.2022.1031011

Check for updates

OPEN ACCESS

EDITED BY
Roi Gazit,
Ben Gurion University of the Negev, Israel

REVIEWED BY
Benyamin Rosental,
Ben-Gurion University of the Negev, Israel
Uri Hershberg,
University of Haifa, Israel

*CORRESPONDENCE
Yoram Louzoun
louzouy@math.biu.ac.il

SPECIALTY SECTION
This article was submitted to Cancer Immunity and Immunotherapy, a section of the journal Frontiers in Immunology

RECEIVED 29 August 2022
ACCEPTED 27 December 2022
PUBLISHED 20 January 2023

CITATION
Akerman O, Isakov H, Levi R, Psevkin V and Louzoun Y (2023) Counting is almost all you need.
Front. Immunol. 13:1031011.
doi: 10.3389/fimmu.2022.1031011

COPYRIGHT
© 2023 Akerman, Isakov, Levi, Psevkin and Louzoun. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Counting is almost all you need

Ofek Akerman[1,2], Haim Isakov[1], Reut Levi[1], Vladimir Psevkin[1] and Yoram Louzoun[1]*

[1]Department of Mathematics, Bar-Ilan University, Ramat Gan, Israel, [2]Department of Computer Science, Bar-Ilan University, Ramat Gan, Israel

The immune memory repertoire encodes the history of present and past infections and immunological attributes of the individual. As such, multiple methods were proposed to use T-cell receptor (TCR) repertoires to detect disease history. We here show that the counting method outperforms two leading algorithms. We then show that the counting can be further improved using a novel attention model to weigh the different TCRs. The attention model is based on the projection of TCRs using a Variational AutoEncoder (VAE). Both counting and attention algorithms predict better than current leading algorithms whether the host had CMV and its HLA alleles. As an intermediate solution between the complex attention model and the very simple counting model, we propose a new Graph Convolutional Network approach that obtains the accuracy of the attention model and the simplicity of the counting model. The code for the models used in the paper is provided at: https://github.com/louzounlab/CountingIsAlmostAllYouNeed.

KEYWORDS

repertoire classification, immune repertoire, machine learning, attention, graphs, T cells, immunology

1 Introduction

Following recent developments in immune sequencing technology (1–3), large T-Cell Receptor (TCR) repertoires can be sampled. Given the association of diseases and TCRs, such repertoires could in theory be used for systemic detection of disease history. However, methods to decipher the disease history from these repertoires (currently denoted as "reading the repertoire") are still limited. Recently, Bayesian approaches and machine learning methods to read repertoires (4–9) were proposed in this field, with good accuracy. However, even those do not reach the accuracy required for clinical usage.

From a computational point of view, the repertoire classification problem is a Multiple Instance Learning (MIL) task. MIL problems arise when the training examples are of varying sizes. In MIL problems, a set or bag is labeled instead of a single object. In the standard definition, a bag $X=\{x_i\}$ receives a label $Y_X=max\{y_i\}$, where y_i is the label of x_i. Here, $y_i \in \{0,1\}$. However, this can be extended to any label. During training, we are unaware of y_i. Only Y_X, the class of each bag in the training set, is known. Examples of MIL problems are video classification, where each frame is an instance, text classification, where each word is an instance, 3D object classification, where each point is an instance, and more (10, 11).

The standard MIL assumption can be expanded to address tasks where positive bags cannot be identified by a single instance. However, the bag can still be classified by the distribution, interaction, or accumulation of the instances in the bag (10).

To formulate the TCR repertoire classification task as an MIL task, a repertoire can be viewed as a bag of TCR sequences, of which a very small fraction is associated with the class of interest. We use the following notations in the current analysis: $T=\{t_1,t_2,t_3,\ldots,t_R\}$ is the group of all TCRs in all samples (training or test) that may be very large. $X_j=\{t_{j_1},t_{j_2},t_{j_3},\ldots,t_{j_N}\}$ is a specific repertoire and $Y(X_j)\in\{0,1\}$ is the binary label of the repertoire X_j. We further assume for the sake of notation simplicity that a TCR t can either bind or not bind any peptide p, with some arbitrary binding cutoff. We denote the set of TCR that binds the peptide p by $T(p)$.

The TCR repertoire classification problem includes unique difficulties compared with classical MIL problems:

- **Low overlap** - The immune repertoire overlap of different individuals is low [Greiff et al. (12)Elhanati et al. (, 13)]. Given two repertoires X_j,X_j, $|X_j\cap X_k|$ is very small.
- **Non-injectivity of TCR-peptide binding** - Multiple sequences can bind to the same pathogen (14). $|T(p)|>1$ for most target peptides.
- **Large TCR diversity** - Recent studies suggest that the human body can have $>10^{14}$ unique TCR sequences (15). $|T|\geq 10^{14}$.
- **An extremely low Witness Rate (WR)** - In MIL problems, the WR is defined by the percentage of discriminating instances within a bag. A WR of 1-5% is considered low in MIL tasks (8). We analyze here a large CMV binding dataset, used by multiple groups (8, 13, 16, 17). Each immune repertoire in the dataset has an average of 192,515 (± 80,630 s.d) unique TCR sequences (4), of which we further estimate only an order of 100 are associated with CMV (4, 18), i.e., the WR can be lower than 0.0001%. Formally, for each repertoire X_j and target peptide p $\frac{|T(p)|}{|X_j|}$ is very small.

We here show that counting arguments actually produce better results than the current SOTA ML or Bayesian methods. We then further improve on that by including the similarity between different TCRs using the combination of a Variational AutoEncoder (VAE) (19), and a novel attention model to include not only the relative importance of positive samples but also their quantity, named attTCR (attention TCR). Finally, we propose an intermediate solution between the counting and attTCR - gTCR that uses a graph of the TCR repertoire co-occurrences to predict the class of a sample.

2 Related work

In recent years, ML and statistical data analysis tools have been proposed to solve the repertoire classification problem. Emerson et al. (4) released a dataset composed of 786 immune repertoires, most of them with a CMV negative/positive classification as well as low-resolution class-I HLA typing (for a detailed data description see section 5.9). They use a Fisher exact test to score TCRs based on their association with positive and negative repertoires and classify TCR repertoires as either positive or negative to CMV or for a given HLA allele.

Their work has been enlarged by TCR-L (5) who evaluate the association between the TCR repertoire and clinical phenotypes. TCR-L expands on Emerson and also uses information about the structure of the TCR sequences and other information about the patient.

Machine learning models, and specifically, attention-based machine learning models, were also proposed as immune repertoire classifiers. deepTCR (6) implements multiple deep learning methods, and a basic form of attention-based averaging. deepTCR encodes each TCRβ chain with a combination of its Vβ, Dβ and Jβ genes using a Convolutional Neural Network (CNN) that extracts sequence motifs. This information is further encoded using a VAE. Then, an attention score is given to each TCR using a custom attention function they designed called AISRU. Finally, a fully connected network (FCN) classifier determines the immune repertoire's status.

Another recently developed model (7) uses 4-mers, sub-sequences of the TCRs CDR3. A logistic regression model is trained on the 4-mers as inputs. Similarly, MotifBoost (20) uses 3-mers to classify the repertoire, using GBDT (gradient-boosted decision trees).

Finally, Deep-RC (8) implements an attention model and uses 1D CNNs in order to embed every TCR to a fixed dimension. Those embeddings are forwarded to more FCN layers and awarded attention scores using a Transformer-like (21) attention equation.

3 Novelty

The algorithms presented here present multiple novel aspects to improve the accuracy of repertoire association studies.

First, we show that a simple counting argument obtains a higher accuracy than all previous methods.

We then propose a novel attention method that on the one hand gives a different importance to different components, but on the other hand, counts them. This is obtained through the sum over the attention of each TCR, with no softmax, but with sigmoid. We show that in contrast with classical attention models, the attention scoring with a non-constant sum improves performance over the simple counting algorithm. The only normalization performed is on the sum of the attention scores, to put it in the active range for the loss function.

Finally, we combine the counting and attention in the Graph Neural Network (GNN) based gTCR model. We use a GNN to classify the repertoire. To the best of our knowledge, this is the first usage of GNN in TCR repertoire classification. The proposed GNN has two novel methodological aspects. First, the contribution of self-edges in the modified adjacency matrix is learned with the weights. Second, we use vertex identity-aware graph classification. The combination of these two methods obtains the accuracy of the attention model with the simplicity of the counting one.

At the technical level, attTCR offers several improvements over Deep-RC (8) and deepTCR (6). The embedding method of each TCR using a cyclic variational autoencoder has never been used on TCRs.

The combination of these methods produces three levels of complexity for the model, where even the simplest model is more accurate than current state of the art (SOTA) models.

10.3389/fimmu.2022.1031011

4 Results

4.1 Positive selection and detection of TCRs associated with a condition

Although the TCR repertoire is very diverse, with most positions along the CDR3 highly variable (15, 22), still a large number of TCRs are shared among multiple patients.

We computed sharing of TCRs between samples in the Emerson dataset (4) (further denoted ECD), where a TCR is defined as the combination of *V*β, and *J*β genes and a CDR3 amino acid sequence (even with different nucleotide sequence). While most of the TCR sequences appear in a single repertoire, there are ~10^5 unique TCRs that appear in more than 10 different repertoires, and hundreds of TCRs that appear in more than a 100 repertoires (Figure 1A). As such, there is enough intersection between different TCRs to perform classification algorithms.

One can assume that following T cell clonal expansion, TCRs that bind to specific diseases are more frequent, and as such are likely to appear in repertoires of people who are or were infected by the disease. However, while we expect some TCRs to be positively associated with a disease or a condition, there is no a-priori reason for any TCR to be negatively associated with a condition (i.e., that its absence is evidence for a condition). To test the absence of negative selection by pathogen, we split the data into a training and a test set (see 'Experimental setup'), and calculated the χ^2 score between the expected and observed number of CMV-positive patients that carry a TCR for both the train and test sets (see section 5.3). We then multiplied the score by the sign of the difference between the expected and observed number of CMV+ patients carrying the TCR (i.e., TCRs less present in positive samples than expected have a negative sign - Figure 1B).

For the vast majority of the TCRs, the χ^2 score is distributed around 0. However, there are some outliers with high χ^2 scores in the training set. Many of those also have a high χ^2 score in the test set (red points). More interestingly, the deviation is only on the positive side. In other words, some TCRs are strongly positively associated with the CMV+ patient class. However, as expected, there are no TCRs associated with the CMV- patient class. We propose to use (only) the TCRs positively associated with the condition (CMV in this case) in the training set to classify patients.

4.2 No systemic difference between CMV+ and CMV- samples

High χ^2 score reactive TCRs are obviously more likely to be shared between more repertoires than the other TCRs (Figure 1C), since a non-shared receptor per definition has a low χ^2 score. Although reactive TCRs go through clonal expansion, checking which TCRs have a large frequency within the repertoire of each donor is not a sufficient method to find such reactive TCRs. Figure 1D demonstrates the lack of correlation between the χ^2 score of each TCR and its average frequency in the samples where it is present.

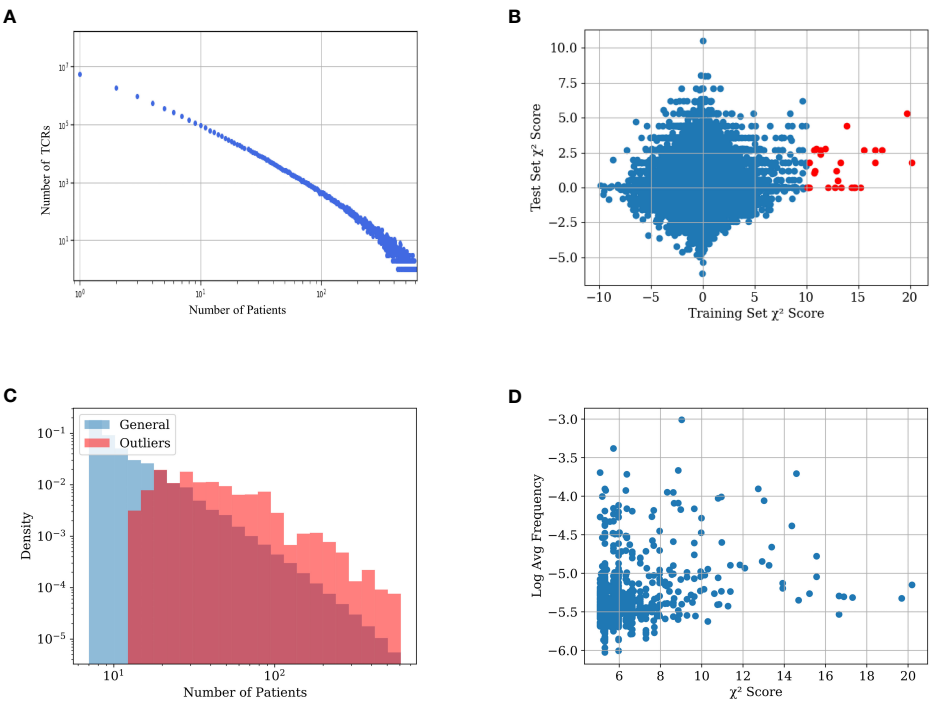

FIGURE 1
(A) TCR number as a function of the number of the patient repertoires that have them in the training set. **(B)** Distribution of the TCRs' χ^2 scores in the training and test sets. The x-axis value is the χ^2 score of the TCR on the training set, the y-axis value is the χ^2 score of the same TCRs on the test set. TCRs with an absolute χ^2 score of over 10 in the training set are colored red. Notice that there are only such points on the positive side of the axis. **(C)** Distribution of average frequency per sample reactive and general TCRs in the dataset. General TCRs refers to all the TCRs in the dataset included in at least 7 repertoires, and reactive TCRs refer to the 200 TCRs with the highest χ^2 score. The distribution of reactive receptors is clearly shifted to the right. **(D)** Scatter plot of different TCRs in the dataset. The x-axis represents the χ^2 score of each TCR, and the y-axis represents its log average frequency in the repertoires it appears in. No correlation is observed between the two.

Instead of focusing on a specific TCR, one could propose to use more generic features of the repertoire to distinguish between CMV+ and CMV- patients [Tickotsky-Moskovitz et al. (23) Gordin et al. (24) Snir and Efroni (25)]. This may be true for lytic conditions, but not for latent or historical conditions. We expect no difference in the general properties of the peripheral repertoire. For events in the distant past, most of the TCRs that were active during the immune response are no longer in the blood in high quantities, and when looking at the general data distribution in the repertoire, there is no difference between positive and negative repertoires (see the Appendix for comparison between V, and J gene distributions and the CDR3 compositions of CMV+ and CMV- patients).

4.3 Counting is all you need

Given the association of specific TCRs with a condition, one could propose different methods to combine reactive TCRs into a classifier for disease history. We here argue that counting the number of such TCRs in a repertoire is a better classifier than existing complex ML classifiers.

To clarify that, we propose a simplistic model that captures the essence of the problem. Assume a general, very large set of TCRs, where each patient has a random subset of these TCRs. Within the large set of TCRs, there is a small subset associated with the disease, and patients that had the disease have a higher than random chance of having these TCRs (see Figure 2 for a description of the model). The data generation process uses 3 probabilities: p_0 - the probability that a TCR would be selected in any patient, p_1,p_2 - the probability that a selected TCR associated with CMV is added to a repertoire in CMV positive and negative samples (Figure 2).

In this model, all TCRs are independent (the presence or absence of different TCRs are not correlated). In such a model,

$$log(P(CMV+|X_j)) = log(P(CMV+)) + log(P(X_j|CMV+)) - log(P(X_j)) = log(P(X_j|CMV+)) + C. \quad (1)$$

Since the TCRs are independent,

$$log(P(X_j|CMV+)) = \sum_i log(p(t_{j_i}|CMV+)) \quad (2)$$

$p(t_{j_i}|CMV+)$ are sampled from a binomial distribution. For reactive TCRs $E(p(x_{j_i}|CMV+))=p_0p_1$, whereas $E(p(x_{j_i}|CMV-))=p_0p_2$. Since $p_2 \ll p_1$, the negative component can be ignored. Since the χ^2 index awards a high score to TCRs that appear in more positive repertoires than negative TCRs, we can expect that by picking a conservative threshold, most of the TCRs that have a high enough χ^2 are truly reactive (as can be observed from the absence of TCRs with parallel negative scores). However, since general non-reactive TCRs appear in large amounts in both positive and negative repertoires, some might still pass the threshold and be falsely classified as reactive TCRs. When the value of $p_0{*}p_1$ is large enough so that there are many more true reactive TCRs found than false reactive TCRs, we expect that classification to be correct.

We calculated the number of false and true reactive TCRs that are extracted by the χ^2 scoring for different $p_0{*}p_1$ values, using the binomial distribution above (Figure 3A). In the specific sample sizes (see Methods for details of simulations) used here, one can clearly see that by a value of $p_0{*}p_1>0.06$ there are considerably more true reactive than false reactive TCRs detected. Below this value, classification would be impossible, while above this value, it should be straightforward. To test that, we applied a straightforward algorithm, where we counted the number of significant TCRs as defined by the

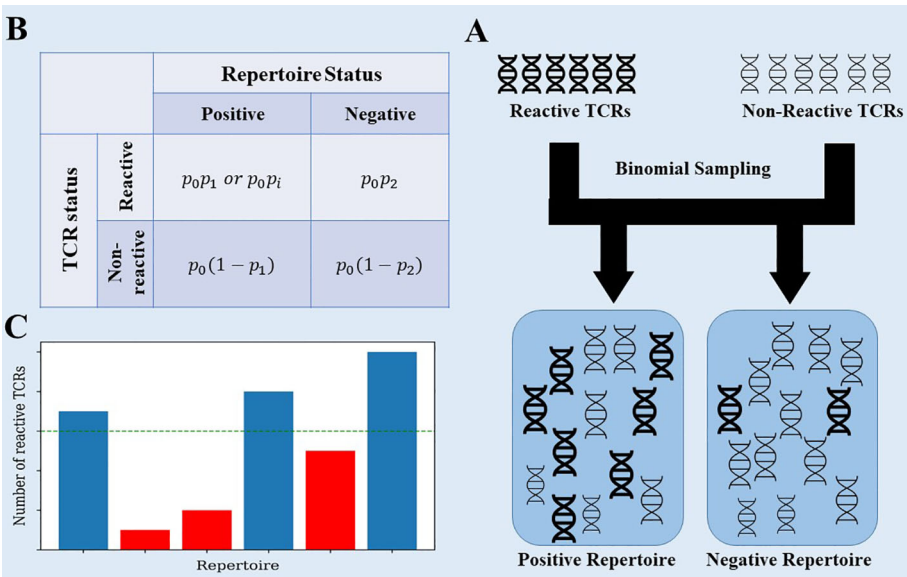

FIGURE 2
(A) The data generation process of the toy model. Each generated repertoire is created using binomial sampling from a collection of positive and negative TCRs. **(B)** The data generation process uses 3 probabilities: p_0 - the probability that a TCR would be selected in any patient, p_1,p_2 - the same for TCRs associated with CMV in CMV+ and CMV- samples. We also tested a model where we replaced p_1 with p_i~N(p_1,σ^2) for each positive TCR t_i.
(C) When classifying the generated repertoires, the reactive TCRs are extracted from each repertoire using the χ^2 score on the training set, and then counted in the test set. Repertoires with a large enough number of reactive TCRs are classified as positive.

10.3389/fimmu.2022.1031011

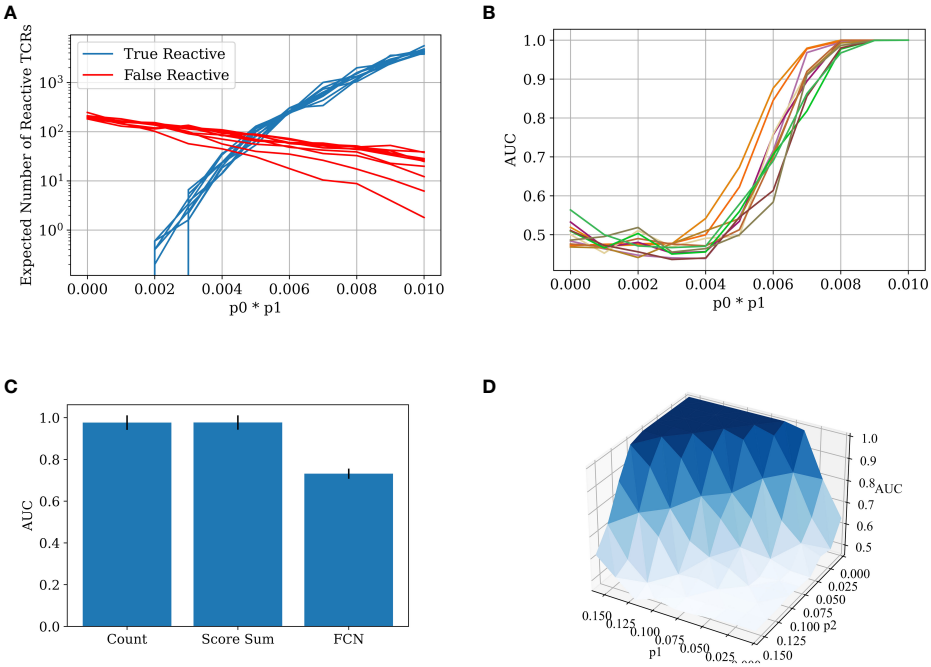

FIGURE 3
(A) The number of true reactive and false reactive TCRs extracted by the χ^2 scoring. The number is the average of 5 calculations on the training set over a 5 CV split. Each line represents a constant $p_0 \sim U[0.01,0.1]$ value with different p_1 values. The x-axis is the product of p_0 and p_1. The other parameters are constant: N=100,000, p_2=0.002. **(B)** The AUC score for data generated with different p_0, p_1 probabilities (5-CV fold). The classification was obtained using the counting method. The colors represents different$p_0 \sim U[0.01,0.1]$ values with different p_1 values. The x-axis is the product of p_0 and p_1. The other generation parameters are as above. **(C)** Bar plot of the AUC results for different models on the 5 CV above. In all the models, meaningful TCRs are extracted by calculating the χ^2 score for each TCR in the test set, and then taking only TCRs above a certain threshold (in this case, 3.84). The counting model counts the relevant TCRs in each test set sample and classifies it by the number of relevant TCRs in each repertoire. The score sum model sums the χ^2 score for the relevant TCRs in the test repertoires and classifies them according to the sum. The FCN model trains a 2-layer FCN over the training repertoires and then makes a prediction on the test repertoires using the TCR one-hot vectors as input. The parameters used in the generation of the repertoires are N=100,000, p_0=0.1, p_1=0.08, and p_2=0.002. **(D)** A surface plot that presents the AUC of the counting model for different p_1 and p_2 combinations. Here, p_1 is not constant for each TCR. Instead, p_i is sampled for each TCR t_i (see Figure 2) from a normal distribution. The other generation parameters are constant: p_0=0.01,σ=0.03,N=100,000.

training set in each test sample and used the count as a classification score. One can see that the transition between the points that there are more false reactive TCRs than true reactive TCRs to there being orders of magnitude more true reactive TCRs than false reactive TCRs is sharp, and the AUC transition is expected to be similar. As such, either classification is trivial and then counting is enough, or it is impossible and then all other algorithms will also fail. The same holds for all parameter regimes of p_2 and p_1.

The test generated data (500 positive repertoires, 500 negative repertoires) and was split into a test and a training set. Reactive TCRs were extracted from the training set and counted in each sample in the test set. Then, an AUC score was calculated using the number of positive clones present in each repertoire in the test set. We ran the counting model on the generated data with different parameters. As expected from the argument above, when trying to classify the generated data with a low value ofp_0*p_1, the classification is impossible. With a high enough value of p_0*p_1, the classification is almost trivial, and a simple counting model can achieve a perfect AUC (Figure 3B). More importantly, the range between the two extremities is very narrow, either you can or cannot classify the repertoires using counting. Since there is no *a priori* reason to assume for any disease and sampling level in any given experiment that they are exactly in this narrow range, one can argue that in general for any disease, either classification is impossible, or a simple counting argument can obtain a high accuracy.

Given this simple argument, one would expect other methods to simply overfit in the simulation above. To test for that, we compared the counting with more complex methods (see Methods). Indeed, counting the relevant TCRs is the best repertoire classification method. The introduction of machine learning methods often only reduces the classification accuracy, following over-fitting on the training set (Figure 3C).

To ensure that the results are not an artifact of the highly simplified model, where all the positive TCRs have the same probability, we further enlarged the model to contain a different *a priori* probability for each positive TCR to appear (see Methods). Figure 3D shows that the conclusion of the sharp transition is true even with looser conditions. Even when p_2 is changing, and when the reactive TCRs are sampled in a non-constant distribution, there is still a clear and sharp "tipping point" between impossible and easy classification, suggesting that this argument may apply to real sampled data.

4.4 Application to real data

To show that the counting argument works in general even when not all TCRs are independent, we analyzed the immune repertoire ECD (4).

To test for the CMV classification, we split the data into a training: validation:test split ratio of 8:1:1, and used 9 cross-validations on the

10.3389/fimmu.2022.1031011

training and validation (the test set was either not changed or ever used in the training). We then applied the counting method:

1. Calculate the χ^2 score for each TCR in the training set.
2. Extract the top-k TCRs with the highest χ^2 score. In this case k=100. One could alternatively use a p value cutoff with similar values, but we have here tried to minimize the hyperparameter optimization to show how generic the counting algorithm is.
3. Count the number of reactive TCRs in each test sample.
4. Calculate AUC on the test set using the counts above.

Again, the counting model outperformed all published models, including the Emerson et al. (4) model on the same test set for different training set sizes (Figure 4). The advantage of the counting algorithm is further obvious in small training sample sizes. In contrast with Emerson et al. (4) and deepRC (8), the counting method can obtain a signal even for 100 training samples.

The current approach has multiple limitations. First, it is based on the assumption of positive selection. This is probably true in all infective diseases, but may not be true in other conditions, such as autoimmunity. Another more significant limitation is the limitation within batch predictions. The observed repertoire may be affected by the details of the experiment protocol, which in turn may affect the frequency of different TCRs. We now plan to test this issue.

4.5 TCRs correlations

In contrast with the simplistic model, TCR usage in real samples can be correlated. The counting method, as adequate as it is, neglects the information that may be available in this correlation. As such it does not reach a perfect AUC in the ECD. To check the co-expression of reactive TCRs, we computed the Spearman correlation between the appearance vector of each TCR in each sample (1 if the TCR is in the sample and 0 otherwise (Figure 5), and clustered the samples based on their correlation). The clusters of related TCRs are very clear. To test the significance, we used a Mann-Whitney U-test between the correlation matrix and a correlation matrix of random shuffled vectors (p<1.e−100).

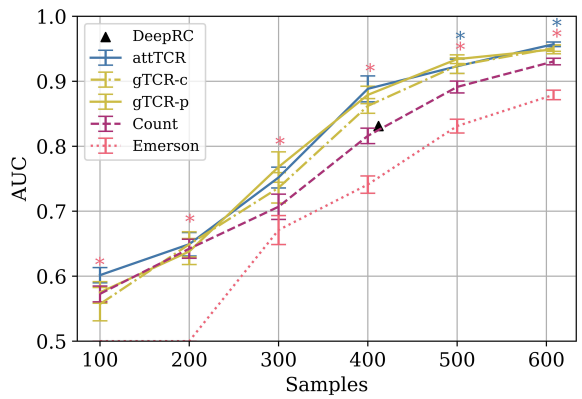

FIGURE 4
The AUC results of different models on different train sample sizes on the ECD (4). The results are over a 9 CV split of the training and the dev sets. The test set is the same for every model. Stars are used to mark statistically significant results using a t-test (p<0.05). Pink stars represent the t-test between the Emerson model and the counting model, and blue stars represent the t-test between attTCR and the counting model. The results were also compared to the results reported for deepRC (8) (with a different experimental setup). For further result comparison to DeepRC and MotifBoost on the ECD, which have even lower AUC, see (20).

4.6 Autoencoder projections

To address the similarity between TCRs, one can use either a sequence similarity (how similar are the TCR CDR3 and V sequences), or a functional similarity (how often they co-appear in the same sample). For the sequence similarity, we projected each sequence using an improvement of the ELATE (Encoder-based LocAl Tcr dEnsity) TCR autoencoder (19). ELATE was enlarged to become a cyclic variational autoencoder, and the TCR representation method was improved (see Methods).

To confirm that the autoencoder projection is associated with the class of the TCRs, we sampled 100 TCRs out of the 200 TCRs with the highest χ^2 score, and 100 random TCRs, and computed the average nearest neighbor euclidean distance between the projections within each group (with 30 cross-validations). The distance between reactive TCRs is significantly lower than random TCRs (12.95 vs 13.886, T-test p<1.e−10), suggesting that reactive TCRs are evenly distributed among all TCRs.

4.7 attTCR

In order to combine the projections into a classifier, we propose an attention model. However, classical attention models sum the positive attention scores to 1. As such, these models would fail to count the number of reactive TCRs in a sample. Instead, they would focus on the relative importance of reactive TCRs. We thus propose a novel attention model that does not apply a softmax to the score assigned to each reactive TCR (see Methods for details), but sigmoid. As such it allows us to estimate the relative importance of reactive TCRs and on the other hand to count them. The sum is then normalized to be in the active range for the loss function. We then tested the combination of the projection and the attention on the ECD, and the results are significantly better than the counting algorithm (Figure 4, for every training set size, and p values of differences therein), and much better than the existing leading models compared in this paper. Many of the parameters here were not optimized to avoid any overfitting. However, a higher performance could be obtained through hyperparameter optimization.

4.8 gTCR

attTCR has an impressive precision. However, it is complex and its training is costly (in GPU time). An alternative method to incorporate the relation between TCR would be purely based on their co-occurrence in samples. To address that, we propose a novel GNN formalism that we denote Graph TCR (gTCR). We define a graph connecting TCRs based on the correlation between their co-occurrence patterns (two TCRs are connected if the Spearman correlation coefficient between their co-occurrence vector is above 0.2). Then the occurrence vector of each TCR in a given sample is the input of this GNN In parallel, the log frequency of each TCR is included as the input to an FCN and the last layers of the two are the input of a final FCN layer that combines the

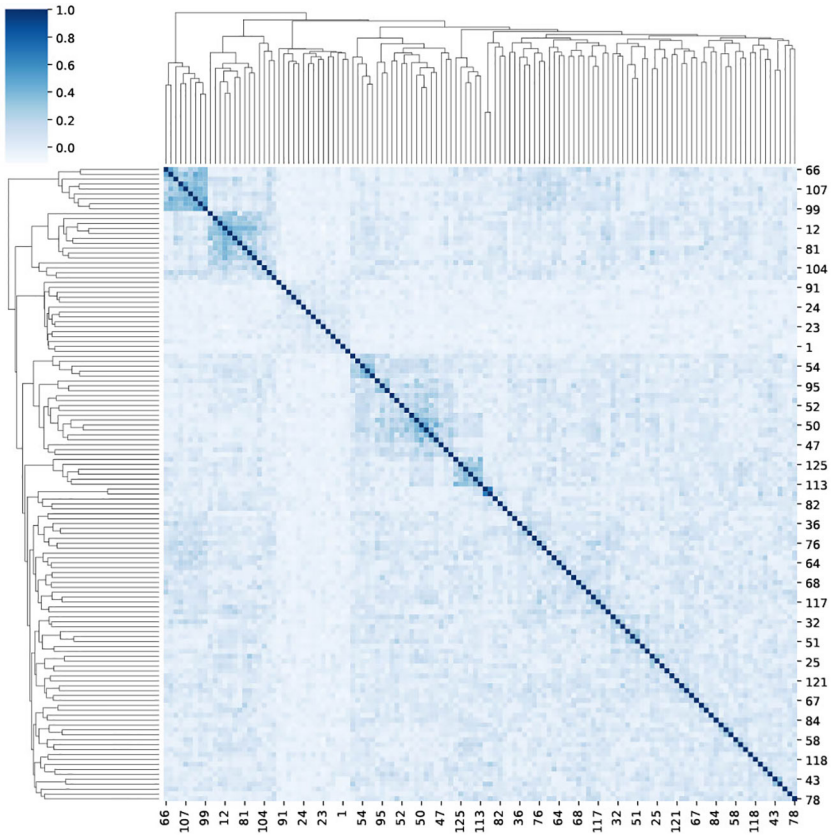

FIGURE 5
A clustermap of the Spearman correlation between 125 reactive TCRs. For each reactive TCR, we extracted from the ECD (4), and we assigned a one-hot vector that represents the appearance of the TCR in different repertoires in the data. Then, for each TCR pairing, we calculated the Spearman correlation between their one-hot vectors.

interaction map with the co-occurrence and the log frequency. The results of gTCR are close to the results of attTCR, with no significant difference (Figure 4). Note that similar results can be obtained by producing a graph using the similarity of the TCRs projection (denoted in the figure gTCR-p in contrast with gTCR-c).

The difference between the two gTCR models is simply the interaction matrix, which can be either based on the sequence or the appearance similarity. The resulting interaction matrices are very different (Jaccard index =0.002 ± 0.003 in 10 training/test divisions). Thus, information seems to be available through both distance definitions.

4.9 HLA allele repertoire classification

The ECD (4) provides the low-resolution A and B HLA-alleles of most samples. We further tested the algorithms above HLA prediction accuracy. From a MIL point of view, this is equivalent to CMV classification. Indeed, the counting model handles this classification task very well, especially with very frequent HLA alleles (Figure 6A). The difference between the counting model and the Emerson model is statistically significant (p=0.017 with Mann-Whitney U-test). We use the counting model with k=100. We use a number cutoff instead of a threshold cutoff to ensure that we find reactive TCRs for rare HLA alleles. Those TCRs receive a relatively low χ^2 score to the reactive TCRs since very few samples have them.

The counting model has a higher accuracy than the Emerson model on most HLA alleles (Figure 6B). Machine learning models, specifically attTCR and gTCR-c, have similar results to the counting model for common HLA alleles (Figure 6C), but over-fit for rare HLA alleles (Figure 6D). For full results over all the HLA alleles, see the Supplementary material.

Multiple other comparisons were proposed, such as taking the MIRA (27) Covid-19 samples as positive repertoires and the ECD as negative repertoires (since there was no COVID-19 at the sampling time), and the counting method obtains an AUC of 1 on this comparison. However, this may be a batch effect, since the two samples may have differences in the sampling and analysis protocol.

5 Methods

5.1 Machine learning definitions

To clarify the machine learning terminology, we include Table 1.

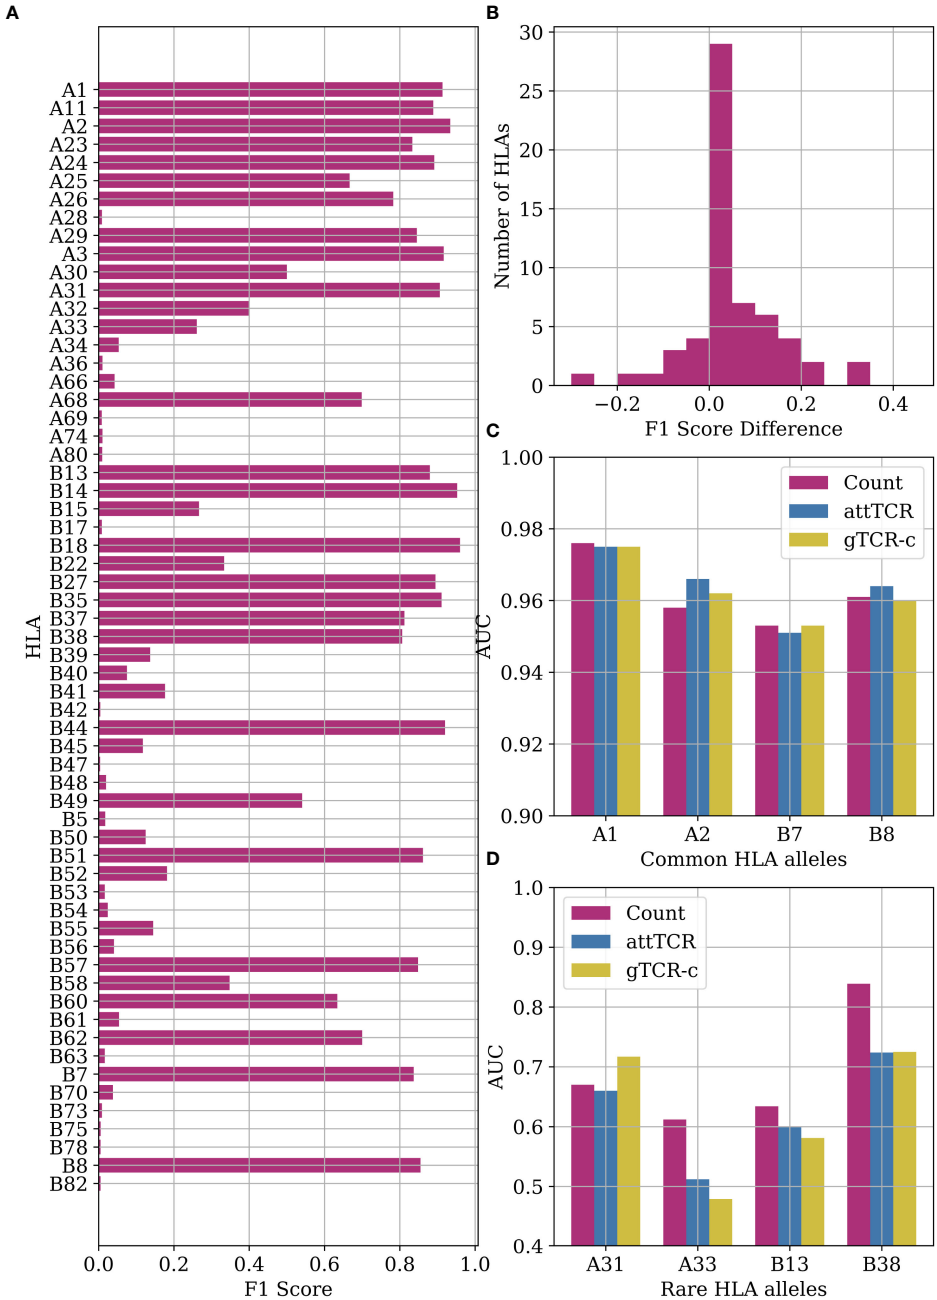

FIGURE 6
(A) F1 score results for the counting model on HLA classification. We performed a leave-one-out split over the entire dataset. **(B)** A histogram of the F1 score differences between the classification results of the counting model and the Emerson model. The difference between the counting model and the Emerson model is statistically significant (p=0.017 with Mann-Whitney U-test) (4) on the same HLA alleles. **(C)** Comparison of AUC results of the counting model, attTCR, and gTCR-c on the repertoire HLA classification task. A 5-fold CV was used, and the AUC was calculated using prediction pooling instead of averaging (26). The HLA alleles presented are the most frequent HLA alleles in the dataset. **(D)** Comparison of AUC results of the counting model, attTCR, and gTCR-c on the repertoire HLA classification task. The HLA alleles presented are the least frequent in the dataset.

5.2 Simulated samples

In order to analyze the performance of the classification methods, we propose a simple simulation that captures the essence of the classification problem. Assume a general, very large set of TCRs, where each patient has a random subset of these TCRs. Within the large set of TCRs, there is a small subset associated with the disease, and patients that had the disease have a higher than random chance of having these TCRs (see Figure 2 for a description of the model). The data generation process uses 3 probabilities: p_0 - the probability that a TCR would be selected in any patient, p_1,p_2 - the probability that a chosen TCR is associated with positive and negative samples. We also tested a model where we replaced p_1 with $p_i \sim N(p_1,\sigma^2)$ for each reactive TCR t_i. Note that p_1 may not follow a truncated normal distribution. It could for example follow a log-normal or scale-free distribution. However, this is a toy model, and we wanted to propose the simplest model.

TABLE 1 Basic machine learning terms used in this paper and their definitions.

Term	Definitions
Training set	The set of data points used to optimize our model's weights.
Validation set	The set of data points used to internally evaluate the model's performance, and optimize the hyper-parameters of the model.
Test set	The set of data points used to externally evaluate the model. The model was never exposed to the test set samples until the final test.
AUC (Area Under Curve)	The ROC curve is created by plotting the true positive rate ($TPR = \frac{TP}{P}$) against the false positive rate ($FPR = \frac{FP}{N}$) for different threshold values. *TP*, and *FP* are the fraction of samples classed aspositive, and are truly positive or negative, respectively; and *P* and N are the total number of real positives and negatives, respectively. A positive/negative classification is an above/below threshold score. The AUC is the area under the curve. It is 1 for a perfect classifier and 0.5 for a random classifier [28]
FCN (Fully Connected Network)	A type of machine learning model that is composed of layers, and weights (matrices) connecting these layers. An FCN is usually composed of an input layer, hidden layers, and an output layer.
Hidden layer	An internal layer of an FCN
Loss Function	A differentiable function that scores the predicted output of a model in relation to the real label. The function is used in the model parameter training.
MSE (Mean Squared Error) Loss	Average squared difference between two vectors.
Activation Function	A non-linear function used between layers of a network.
ReLU (Rectified Linear Unit)	An activation function. $f(x)$=max(0,x).
Sigmoid	An activation function that scores the input between 0 and 1. $f(x) = \frac{1}{1+e^{-x}}$. Frequently used as an output function of a model.
Softmax	A function that receives a vector and normalizes its values to sum to one.
Dropout	A regularization tool used on the models' layers. Each layer zeros some of its values randomly according to the dropout probability
Graph Convolutional Network	A Neural network structure where the input of each node in each layer is a combination of its neighbors in a graph.

In the different trials performed in the current analysis, we generated 1,000 different repertoires (500 positives, 500 negatives) using differing generation probabilities (p_0, p_1, p_2). All the experiments were performed using a 4:1 training:test split, using 5 different data splits.

5.3 χ2 Score

To extract reactive TCRs from a repertoire, we use a simple scoring method. For each TCR t_i, the χ^2 formula uses the following values:

- $Npos_i$ - The number of positive repertoires that contain t_i.
- N_i - The total number or repertoires that contain t_i.
- $Npos$ - The total number of positive repertoires in the data.
- N - The total number of repertoires in the data.

The χ^2 score for TCR t_i is calculated using Equation 3.

$$\chi^2 i = \frac{sign(Npos_i - \frac{NposN_i}{N})(Npos_i - \frac{NposN_i}{N})^2}{\frac{NposN_i}{N}} \quad (3)$$

The difference with the regular χ^2 is simply the sign of the deviation. Since we do not only want to detect deviations from the null model, but also whether the number of positive TCRs is higher or lower than expected since only positive selection is of interest. Also, note that we assume $N_pos * N_i/N$ is larger than 1 so that there are at least a few expected positive TCRs.

5.4 Counting model

The counting model is a simple model that effectively manages to distinguish between positive and negative repertoires on the test set. The counting model has the following steps:

1. Calculate the χ^2 score for each TCR in the training set.
2. Extract all the TCRs with a χ^2 score over a certain threshold. The threshold can be either a *p*-value (i.e. define a threshold based on the translation of a p-value to a cutoff), or a fixed number of TCRs *k*.
3. Count the number of significant reactive TCRs in each file of the test set.
4. Calculate AUC on the test set using the counts.

10.3389/fimmu.2022.1031011

5.5 TCR autoencoder

A TCR autoencoder is a model that preserves the information about input t_i V gene and CDR3 sequence, while reducing the dimension to a low dimension representation z_i. The training of the TCR autoencoder includes several steps of data processing (19). The first step is representing each of the amino acid per position as well as the V genes by an embedding vector. There are twenty possible amino acids and an additional end signal is required. Each instance is then processed by an autoencoder network and encoded to size R^{30} (we have previously checked that adding dimensions beyond 30 had a very limited contribution to the accuracy).

The autoencoder network contains three layers of 800, 1100, and 30 neurons as the encoder and a mirrored network as the decoder. The network is trained with a dropout of 0.2 and a ReLU. An MSE loss function is implemented to compare each input sequence to the resulting decoded sequence (19). The current version differs from the ELATE encoder (19), since it includes a variational term. Instead of encoding an input as a single point, we encode it as a distribution over the latent space. The model is then trained as follows: First, the input is encoded as a distribution over the latent space; second, a point from the latent space is sampled from that distribution, Then the sampled point is decoded and the reconstruction error can be computed; finally, the reconstruction error is back-propagated through the network. The VAE loss function is the same as ELATE with a Kulback-Leibler divergence between the returned distribution and a standard Gaussian.

The problem with the standard VAE is that the KL term tends to vanish. A recent work (29) studied scheduling schemes for β, and showed that KL vanishing is caused by the lack of good latent codes in training the decoder at the beginning of optimization. To remedy this, we used a cyclical annealing schedule, which repeats the process of increasing β multiple times. This new procedure allows the progressive learning of more meaningful latent codes, by leveraging the informative representations of previous cycles as a warm restart.

5.6 attTCR

attTCR receives as an input the reactive TCRs of each repertoire, and as an output a score between 0 and 1 that predicts whether the repertoire is positive or negative. The model is composed of an encoder network, an attention scorer, and a normalization layer. For detailed model, architecture see Figure 7. The encoder was explained above.

5.6.1 Attention network

Each TCR t_i is assigned an attention score a_i, such that $a_i \in [0,1]$. TCRs that are more important to the classification should receive higher attention scores. The attention network takes as an input the embedding of each TCR by the encoder network and is composed of 2 hidden layers of size q. The output of the attention network for each TCR sequence is a single attention score. Therefore, for the entire repertoire, the network outputs a vector v of dimension N (the number of reactive TCRs). A sigmoid function is used to produce an attention score between 0 and 1 for each reactive TCR in the repertoire. We have here used the traditional Transformer (Vaswani et al. (21)) notation. We use the following matrices and vectors to describe the attention process:

- $Q \in R^{N \times q}$ - The queries matrix. In our model, the matrix is created after the 2 hidden layers of the attention network.
- $\xi \in R^{q \times 1}$ - The keys vector. The weights of the output layer of the attention network.

The attention score calculation applied by equation 4, where σ is the sigmoid function:

$$\sigma\left(\frac{\xi^T Q^T}{\sqrt{q}}\right) \tag{4}$$

Note that unlike traditional attention models, we do not use the softmax function on the resulting attention vector, nor do we multiply each attention by the TCR representation. We are not interested in performing a weighted average. Instead, we want to score each TCR and still keep the information about the number of reactive TCRs in the repertoire, i.e., N. Thus the score of an entire repertoire is simply the sum of the attention values for all the reactive TCRs in this sample.

5.6.2 Normalization layer

The input of the normalization layer in a vector $v \in R^N$ with the scores of each TCR in the repertoire. The Normalization layer's goal is

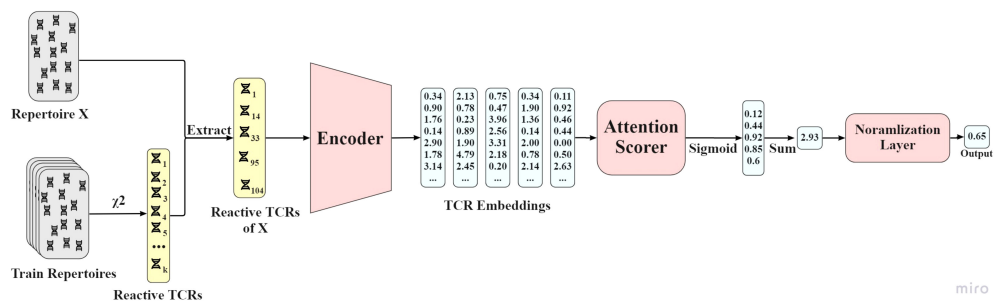

FIGURE 7
AttTCR's architecture. First, the reactive TCRs are sampled from all the train repertoires using the χ^2 method. Then, for each repertoire X, the reactive TCRs contained in X are extracted. Each reactive TCR is projected by the encoder. The projections are then scored by the attention scorer. The scores are summed and normalized. The output of the model is a number between 0 and 1 that indicates the confidence of the model on whether the repertoire is positive.

10.3389/fimmu.2022.1031011

to convert the sum over the vector v to a number between 0 and 1, so we can train the model using BCE loss. Just putting Σv into a sigmoid function is not going to work, since the sum of N scores $a_i \in [0,1]$ is very likely to be too large for the sigmoid function. As a result, all the repertoires would output a number very close to 1, which might hurt the training process. Therefore, we use 2 learned parameters: γ_1, γ_2, to normalize the sum before the sigmoid function. In conclusion, the normalization layer performs Equation 5.

$$\sigma(\gamma_1 \sum v + \gamma_2) \tag{5}$$

5.7 gTCR

5.7.1 Graphs

A GCN is a classic machine learning model used in supervised machine learning with relational data represented as a graph. Kipf and Welling (30) proposed a classical "Graph Convolution Network" (GCN) design, in which layer-wise propagation rules are used to characterize features around each node. For the graph classification task, a GCN model is applied to the graph and vertices' features. Then, the aggregated information was gathered from the vertices' features as well as graph structure is fed into the classification machine.

5.7.2 TCR similarity graph

We define here two ways of modeling the TCR graph. Both ways consist of two stages, a definition of the similarity matrix between the reactive TCRs followed by a zeroing stage where rows and columns from the similarity matrix are filled with zero value if the reactive TCR is absent from the sample's repertoire.

One way of modeling such a similarity matrix between reactive TCRs is obtained using the Spearman correlation matrix between the training samples' presence vectors. These sample's presence vectors contain 0 or 1 according to the presence of each reactive TCR in the sample's repertoire. Another way of modeling a similarity matrix between reactive TCRs is obtained using the inverse of the euclidean distance between the projection of the reactive TCRs obtained from the autoencoder.

5.7.3 gTCR

The gTCR (graph TCR) model combines the information from the frequencies vector as well as the graph represented by the normalized adjacency matrix as can be seen in Equation 6. An embedding vector of the log frequencies is obtained from a 2-layer FCN, each followed by a tanh activation function and dropout layer (Equation 8). In parallel, one layer of a GCN model is applied (Equation 9) to the reactive TCR presence vector. Then the outputs of the two networks are concatenated and serve as the input of a 2-layer FCN to predict a binary condition (Equation 10).

$$\tilde{A} = D^{-\frac{1}{2}} A D^{-\frac{1}{2}} \tag{6}$$

$$\text{D is diagonal matrix such that } D_{ii} = \sum_j A_{ij} \tag{7}$$

$$ev \text{ is the frequencies vectors embedding}: ev = FCN(f) \tag{8}$$

$$eg \text{ is the graph embedding}: eg = (\tilde{A} + \alpha \cdot I) \cdot sign(f) \cdot W \tag{9}$$

$$Concatenate(ev, eg) \Rightarrow FCN \tag{10}$$

α is a learned scalar regulating the importance given to the vertex's feature compared to its neighbors' features. α is initialized with 1 plus Gaussian $N(0,0.1)$.

5.8 Comparison to other methods

The counting method was compared to 2 other classification methods:

- **Score Sum** - This method is almost entirely similar to the counting model. The only difference is that instead of classifying the repertoires based on the number of reactive TCRs found in the repertoires, we classify them by the sum of the χ^2 scores of the reactive TCRs in each repertoire.
- **FCN** - After extracting the reactive TCRs from the data, each repertoire is embedded to a vector of the dimension of the number of reactive TCRs. Each dimension in the vector represents a different reactive TCR, and its value is set to 1 if the repertoire contains the TCR and 0 otherwise. Then, a 2-layer FCN is fitted on the vector training set, and tested on the test set.

5.9 Data

The Emerson dataset contains 786 immune repertoires (4). Each repertoire contains between 4,371 to 973,081 (avg. 299,319) TCR sequences with a length of 1 to 27 (avg. 14.5) amino acids. The V and J genes and the frequency are saved for each TCR. 340 repertoires are labeled CMV+, 421 are labeled CMV-, and 25 are of unknown status. We only use the repertoire with a known CMV status, 761 repertoires in total. In addition 626 of the repertoires have HLA allele information available.

5.10 Preprocessing

The Emerson dataset (4) is composed of 786 repertoires in total. However, since the task at hand is a supervised classification task, the 25 repertoires without a CMV classification are not beneficial to the learning process, so they are removed from the dataset. Then, all the TCRs that have missing CDR3 amino acid information are discarded. In the following step, the TCRs are filtered based on prevalence in different repertoires. Only TCR sequences that appear in 7 different repertoires or more remain in the repertoires after the filtration.

5.11 Experimental setup

When predicting CMV status of the repertoires, the models are tested with a test size that contains 10% (77 samples) of the data. For all the models tested, the test set is the same. attTCR is trained using a 9-fold CV between the training set and the validation set, while gTCR

is trained over 20 different splits. In the counting model, the validation set is not used. All the models are evaluated using an AUC score on the test set (28). The significance of the results is tested against a p-value threshold of 0.05.

The HLA allele repertoire classification in Figure 6A was evaluated using an F1 score. In Figures 6C, D, the measure was changed to AUC over a 5-fold CV with a train:validation:test split of 3:1:1. The AUC was calculated using the pooling method, i.e., calculated once over all the predictions (26).

6 Discussion

We have here proposed three novel methods with different levels of complexity, and shown that even the simplest of these models outperform the current State of The Art (SOTA) for repertoire classification. The simplest model is simply counting reactive TCRs, followed by a novel attention model that combines classical attention models with counting, and finally a combination of graph-based machine learning with MIL. All the models presented in the paper rely on the assumption that TCRs are only positively selected, and that there are no TCRs negatively associated with a condition. Note that Emerson et al. (4) used a Fisher exact test to score TCRs based on their association with positive and negative repertoires. It also classifies each significant TCR as either a positive or a negative selected TCR. However, the assumption that there are any negatively selected TCRs does not make much immunological sense. TCR expansion occurs when a certain TCR binds to an antigen-peptide. There is no equivalent process for TCRs that do not bind to antigen-peptides. Thus, in theory, the abundance of a TCR in a repertoire can only indicate that the TCR was positively selected.

Some TCRs are highly abundant in different individuals (31), and have initial production probability (3, 13). Therefore, positively selected TCRs exist in various frequencies in positive immune repertoires, some especially common ones might appear randomly in negative instances as well. Thus, the relative abundance of a TCR in many repertoires does not automatically make it more indicative than a TCR that appears in a few repertoires. Once a TCR is proven to be positively selected, its frequency does not matter much when it comes to repertoire classification. Hence, the counting model is a good way to classify the repertoires given the reactive TCRs.

We have shown in the data that TCRs are indeed only positively selected, and that it improves on existing models in both theory and real data. There are distinctions to be made between the real repertoire data, and the generated one. The most obvious is that real TCR presence in a repertoire does not follow a binomial distribution. Real TCRs have a scale-free distribution. Some TCRs are public TCRs and are very common (31), and others are very rare. The pool size of positive and negative TCRs to draw from is also vastly different in size. Statistically, there are many more possible negative TCRs than TCRs that bind to an epitope-peptide of a specific disease. In addition, TCRs are sampled in varying sizes, whereas the repertoires in the generative model are all around p_0N. Despite these differences, we believe that the conclusions of the toy model are still true on real repertoire data. However, these differences do not affect the validity of counting and its extension.

The current approach is purely based on the observed TCR presence and absence and on their sequence. It completely ignores the antigen or MHC properties. However, multiple algorithms were proposed for both TCR-peptide (6, 13, 32–44) and TCR-MHC binding (45–56). While the accuracy of such algorithms keeps improving, it is still too, it may be too early to use such algorithms for repertoire classification.

The ML models presented in the paper, especially attTCR, can also be used in a large variety of problems. attTCR presents a new approach to attention scoring, that can be used in every MIL task that involves counting. Further research has to be done on the quality of the proposed ML models on other non-related tasks. However, we propose that these three levels of modeling - counting, counting attentions models and GNNs on selected shared samples may be a general approach to all MIL problems.

Data availability statement

Publicly available datasets were analyzed in this study. This data can be found here: https://clients.adaptivebiotech.com/pub/emerson-2017-natgen.

Author contributions

OA did most of the analysis and wrote a part of the paper. RL developed the VAE. YL supervised the work and wrote a part of the paper. HI developed the Graph-based classification. VP did a part of the statistical analysis. All authors contributed to the article and approved the submitted version.

Funding

This work was funded by ISF grant 20/870 and by the BIU DSI Grant.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be foundonline at: https://www.frontiersin.org/articles/10.3389/fimmu.2022.1031011/full#supplementary-material

SUPPLEMENTARY IMAGE 1

(A) A swarm plot of the different repertoires in the data. Each dot represents a repertoire. The y-axis represents the average count of a TCR in a repertoire, where the count of a TCR is defined as the number of clones the TCR has in the repertoire. It is clear that there is not a big difference in the count distribution between positive and negative repertoires. **(B)** A swarm plot of the different repertoires in the data. Each dot represents a repertoire. The y-axis represents the average frequency of a TCR in a repertoire. It is clear that there is not a big difference in the frequency distribution between positive and negative repertoires.

SUPPLEMENTARY IMAGE 2

(A) A histogram of the different Vb-genes in the data. Each column represents the average frequency of a Vb-gene in positive and negative repertoires. It is clear that the v-gene distribution between negative and positive repertoires is very similar. **(B)** A histogram of the different Jb-genes in the data. Each column represents the average frequency of a Jb-gene in positive and negative repertoires. It is clear that the J-gene distribution between negative and positive repertoires is very similar.

References

1. Georgiou G, Ippolito GC, Beausang J, Busse CE, Wardemann H, Quake SR. The promise and challenge of high-throughput sequencing of the antibody repertoire. *Nat Biotechnol* (2014) 32:158–68. doi: 10.1038/nbt.2782

2. Brown AJ, Snapkov I, Akbar R, Pavlović M, Miho E, Sandve GK, et al. Augmenting adaptive immunity: progress and challenges in the quantitative engineering and analysis of adaptive immune receptor repertoires. *Mol Syst Design Eng* (2019) 4:701–36. doi: 10.1039/C9ME00071B

3. Benichou J, Ben-Hamo R, Louzoun Y, Efroni S. Rep-seq: uncovering the immunological repertoire through next-generation sequencing. *Immunology* (2012) 135:183–91. doi: 10.1111/j.1365-2567.2011.03527.x

4. Emerson RO, DeWitt WS, Vignali M, Gravley J, Hu JK, Osborne EJ, et al. Immunosequencing identifies signatures of cytomegalovirus exposure history and HLA-mediated effects on the T cell repertoire. *Nat Genet* (2017) 49:659–65. doi: 10.1038/ng.3822

5. Liu M, Goo J, Liu Y, Sun W, Wu MC, Hsu L, et al. TCR-l: an analysis tool for evaluating the association between the T-cell receptor repertoire and clinical phenotypes. *BMC Bioinf* (2022) 23:1–16. doi: 10.1186/s12859-022-04690-2

6. Sidhom J-W, Larman HB, Pardoll DM, Baras AS. DeepTCR is a deep learning framework for revealing sequence concepts within T-cell repertoires. *Nat Commun* (2021) 12:1–12. doi: 10.1038/s41467-021-21879-w

7. Ostmeyer J, Christley S, Toby IT, Cowell LG. Biophysicochemical motifs in T-cell receptor sequences distinguish repertoires from tumor-infiltrating lymphocyte and adjacent healthy tissue. *Cancer Res* (2019) 79:1671–80. doi: 10.1158/0008-5472.CAN-18-2292

8. Widrich M, Schäfl B, Pavlović M, Sandve GK, Hochreiter S, Greiff V, et al. DeepRC: immune repertoire classification with attention-based deep massive multiple instance learning. *bioRxiv* (2020). doi: 10.1101/2020.04.12.038158

9. Zhang H, Zhan X, Li B. Giana allows computationally-efficient tcr clustering and multi-disease repertoire classification by isometric transformation. *Nat Commun* (2021) 12:1–11. doi: 10.1038/s41467-021-25693-2

10. Carbonneau M-A, Cheplygina V, Granger E, Gagnon G. Multiple instance learning: A survey of problem characteristics and applications. *Pattern Recognition* (2018) 77:329–53. doi: 10.1016/j.patcog.2017.10.009

11. Uriot T. Learning with sets in multiple instance regression applied to remote sensing. *arXiv preprint arXiv:1903.07745* (2019). doi: 10.48550/arXiv.1903.07745

12. Greiff V, Weber CR, Palme J, Bodenhofer U, Miho E, Menzel U, et al. Learning the high-dimensional immunogenomic features that predict public and private antibody repertoires. *J Immunol* (2017) 199:2985–97. doi: 10.4049/jimmunol.1700594

13. Elhanati Y, Sethna Z, Callan CGJr., Mora T, Walczak AM. Predicting the spectrum of TCR repertoire sharing with a data-driven model of recombination. *Immunol Rev* (2018) 284:167–79. doi: 10.1111/imr.12665

14. Wucherpfennig KW, Allen PM, Celada F, Cohen IR, De Boer R, Garcia KC, et al. Polyspecificity of T cell and b cell receptor recognition. *Semin Immunol* (2007) 19 (4):216–24.

15. Mora T, Walczak AM. How many different clonotypes do immune repertoires contain? *Curr Opin Syst Biol* (2019) 18:104–10. doi: 10.1016/j.coisb.2019.10.001

16. Sethna Z, Elhanati Y, Callan CGJr., Walczak AM, Mora T. OLGA: fast computation of generation probabilities of b-and T-cell receptor amino acid sequences and motifs. *Bioinformatics* (2019) 35:2974–81. doi: 10.1093/bioinformatics/btz035

17. DeWitt WSIII, Smith A, Schoch G, Hansen JA, Matsen FAIV, Bradley P. Human t cell receptor occurrence patterns encode immune history, genetic background, and receptor specificity. *Elife* (2018) 7:e38358. doi: 10.7554/eLife.38358.043

18. Christophersen A, Ráki M, Bergseng E, Lundin KE, Jahnsen J, Sollid LM, et al. Tetramer-visualized gluten-specific CD4+ T cells in blood as a potential diagnostic marker for coeliac disease without oral gluten challenge. *United Eur Gastroenterol J* (2014) 2:268–78. doi: 10.1177/2050640614540154

19. Dvorkin S, Levi R, Louzoun Y. Autoencoder based local T cell repertoire density can be used to classify samples and T cell receptors. *PloS Comput Biol* (2021) 17:e1009225. doi: 10.1371/journal.pcbi.1009225

20. Katayama Y, Kobayashi TJ. MotifBoost: k-mer based data-efficient immune repertoire classification method. *bioRxiv* (2021). doi: 10.1101/2021.09.28.462258

21. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. *Adv Neural Inf Process Syst* (2017), 30.

22. Benichou JI, van Heijst JW, Glanville J, Louzoun Y. Converging evolution leads to near maximal junction diversity through parallel mechanisms in b and T cell receptors. *Phys Biol* (2017) 14:045003. doi: 10.1088/1478-3975/aa7366

23. Tickotsky-Moskovitz N, Louzoun Y, Dvorkin S, Rotkopf A, Kuperman AA, Efroni S. CDR3 and V genes show distinct reconstitution patterns in T cell repertoire post-allogeneic bone marrow transplantation. *Immunogenetics* (2021) 73:163–73. doi: 10.1007/s00251-020-01200-7

24. Gordin M, Philip H, Zilberberg A, Gidoni M, Margalit R, Clouser C, et al. Breast cancer is marked by specific, public T-cell receptor CDR3 regions shared by mice and humans. *PloS Comput Biol* (2021) 17:e1008486. doi: 10.1371/journal.pcbi.1008486

25. Snir T, Efroni S. T Cell repertoire sequencing as a cancer's liquid biopsy can we decode what the immune system is coding? *Curr Opin Syst Biol* (2020) 24:135–41. doi: 10.1016/j.coisb.2020.10.009

26. Bradley AP. The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern Recognition* (1997) 30:1145–59. doi: 10.1016/S0031-3203(96)00142-2

27. Nolan S, Vignali M, Klinger M, Dines JN, Kaplan IM, Svejnoha E, et al. A large-scale database of t-cell receptor beta (tcrβ) sequences and binding associations from natural and synthetic exposure to sars-cov-2. *Res Square* (2020). doi: 10.21203/rs.3.rs-51964/v1

28. Ling CX, Huang J, Zhang H. (2003). AUC: a better measure than accuracy in comparing learning algorithms, in: *Conference of the Canadian Society for Computational Studies of Intelligence*, vol 706. (Springer). pp. 329–41.

29. Fu H, Li C, Liu X, Gao J, Celikyilmaz A, Carin L. Cyclical annealing schedule: A simple approach to mitigating kl vanishing. *arXiv preprint arXiv:1903.10145* (2019). doi: 10.18653/v1/N19-1021

30. Kipf TN, Welling M. (2016). Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907.

31. Huisman W, Hageman L, Leboux DA, Khmelevskaya A, Efimov GA, Roex MC, et al. Public T-cell receptors (TCRs) revisited by analysis of the magnitude of identical and highly-similar TCRs in virus-specific T-cell repertoires of healthy individuals. *bioRxiv* (2021). doi: 10.1101/2021.11.29.470325

32. Glanville J, Huang H, Nau A, Hatton O, Wagar LE, Rubelt F, et al. Identifying specificity groups in the T cell receptor repertoire. *Nature* (2017) 547:94–8. doi: 10.1038/nature22976

33. Dash P, Fiore-Gartland AJ, Hertz T, Wang GC, Sharma S, Souquette A, et al. Quantifiable predictive features define epitope-specific T cell receptor repertoires. *Nature* (2017) 547:89–93. doi: 10.1038/nature22383

34. Springer I, Besser H, Tickotsky-Moskovitz N, Dvorkin S, Louzoun Y. Prediction of specific TCR-peptide binding from large dictionaries of TCR-peptide pairs. *Front Immunol* (2020) 11:1803. doi: 10.3389/fimmu.2020.01803

35. Jurtz VI, Jessen LE, Bentzen AK, Jespersen MC, Mahajan S, Vita R, et al. NetTCR: sequence-based prediction of TCR binding to peptide-MHC complexes using convolutional neural networks. *bioRxiv* (2018), 433706. doi: 10.1101/433706

36. Moris P, De Pauw J, Postovskaya A, Ogunjimi B, Laukens K, Meysman P. Treating biomolecular interaction as an image classification problem–a case study on T-cell receptor-epitope recognition prediction. *bioRxiv* (2019).

37. Fischer DS, Wu Y, Schubert B, Theis FJ. Predicting antigen specificity of single T cells based on TCR CDR 3 regions. *Mol Syst Biol* (2020) 16:e9416. doi: 10.15252/msb.20199416

38. Gielis S, Moris P, Bittremieux W, De Neuter N, Ogunjimi B, Laukens K, et al. TCRex: detection of enriched T cell epitope specificity in full T cell receptor sequence repertoires. *Frontiers Immunol* (2019) 10, 2820. doi: 10.1101/373472

39. Springer I, Tickotsky N, Louzoun Y. Contribution of t cell receptor alpha and beta cdr3, mhc typing, v and j genes to peptide binding prediction. *Front Immunol* (2021) 12. doi: 10.3389/fimmu.2021.664514

40. Montemurro A, Schuster V, Povlsen HR, Bentzen AK, Jurtz V, Chronister WD, et al. NetTCR-2.0 enables accurate prediction of tcr-peptide binding by using paired TCRα and β sequence data. *Commun Biol* (2021) 4:1–13. doi: 10.1038/s42003-021-02610-3

41. Tong Y, Wang J, Zheng T, Zhang X, Xiao X, Zhu X, et al. SETE: sequence-based ensemble learning approach for TCR epitope binding prediction. *Comput Biol Chem* (2020) 87:107281. doi: 10.1016/j.compbiolchem.2020.107281

42. Beshnova D, Ye J, Onabolu O, Moon B, Zheng W, Fu Y-X, et al. *De novo* prediction of cancer-associated T cell receptors for noninvasive cancer detection. *Sci Trans Med* (2020) 12:eaaz3738. doi: 10.1126/scitranslmed.aaz3738

43. Jokinen E, Huuhtanen J, Mustjoki S, Heinonen M, Lähdesmäki H. Predicting recognition between T cell receptors and epitopes with TCRGP. *PloS Comput Biol* (2021) 17:e1008814. doi: 10.1371/journal.pcbi.1008814

44. De Neuter N, Bittremieux W, Beirnaert C, Cuypers B, Mrzic A, Moris P, et al. On the feasibility of mining CD8+ T cell receptor patterns underlying immunogenic peptide recognition. *Immunogenetics* (2018) 70:159–68. doi: 10.1007/s00251-017-1023-5

45. Glazer N, Akerman O, Louzoun Y. Naive and memory T cells TCR-HLA binding prediction. *Oxford Open Immunol* (2022). doi: 10.1093/oxfimm/iqac001

46. Zhang H, Lund O, Nielsen M. The pickpocket method for predicting binding specificities for receptors based on receptor pocket similarities: application to mhc-peptide binding. *Bioinformatics* (2009) 25:1293–9. doi: 10.1093/bioinformatics/btp137

47. Liu G, Li D, Li Z, Qiu S, Li W, Chao CC, et al. Pssmhcpan: a novel PSSM-based software for predicting class I peptide-hla binding affinity. *Giga Science* (2017) 6(5): gix017. doi: 10.1093/gigascience/gix017

48. Andreatta M, Nielsen M. Gapped sequence alignment using artificial neural networks: application to the mhc class i system. *Bioinformatics* (2016) 32:511–7. doi: 10.1093/bioinformatics/btv639

49. Mei S, Li F, Xiang D, Ayala R, Faridi P, Webb GI, et al. Anthem: a user customised tool for fast and accurate prediction of binding between peptides and HLA class I molecules. *Briefings Bioinf* (2021). doi: 10.1093/bib/bbaa415

50. O'Donnell TJ, Rubinsteyn A, Bonsack M, Riemer AB, Laserson U, Hammerbacher J. MHCflurry: open-source class I MHC binding affinity prediction. *Cell Syst* (2018) 7:129–32. doi: 10.1016/j.cels.2018.05.014

51. Reynisson B, Alvarez B, Paul S, Peters B, Nielsen M. NetMHCpan-4.1 and NetMHCIIpan-4.0: improved predictions of MHC antigen presentation by concurrent motif deconvolution and integration of MS MHC eluted ligand data. *Nucleic Acids Res* (2020) 48:W449–54. doi: 10.1093/nar/gkaa379

52. Liberman G, Vider-Shalit T, Louzoun Y. Kernel multi label vector optimization (kmlvo): A unified multi-label classification formalism. In: Nicosia G, Pardalos P, editors. *Learning and intelligent optimization.* Berlin, Heidelberg: Springer Berlin Heidelberg (2013). p. 131–7. doi: 10.1007/978-3-642-44973-4\s\do5(1)5

53. Vider-Shalit T, Louzoun Y. MHC-I prediction using a combination of T cell epitopes and MHC-I binding peptides. *J Immunol Methods* (2011) 374:43–6. doi: 10.1016/j.jim.2010.09.037. High-throughput methods for immunology: Machine learning and automation.

54. Ginodi I, Vider-Shalit T, Tsaban L, Louzoun Y. Precise score for the prediction of peptides cleaved by the proteasome. *Bioinformatics* (2008) 24:477–83. doi: 10.1093/bioinformatics/btm616

55. Vider-Shalit T, Louzoun Y. Mhc-i prediction using a combination of t cell epitopes and mhc-i binding peptides. *J Immunol Methods* (2011) 374:43–6. doi: 10.1016/j.jim.2010.09.037

56. Liberman G, Vider-Shalit T, Louzoun Y. (2013). Kernel multi label vector optimization (kmlvo): a unified multi-label classification formalism, in: *International Conference on Learning and Intelligent Optimization*, . pp. 131–7. Springer.

10.3389/fimmu.2022.1031011

Appendix

In the paper, we have shown that repertoires can be classified using bayesian and machine learning tools on the content of the TCR repertoire. However, in the Appendix we want to prove that there does not exist an easier, more superficial method of distinguishing between positive and negative repertoires. In Supplementary Images 1, 2, we show that different general attributes of the repertoires are the same with positive and negative repertoires, and they cannot be differentiated using this attributes.