

Local and Global Variability in Developing Human T-Cell Repertoires

Giulio Isacchini  ^{1,2,*} Valentin Quiniou  ^{3,4} Pierre Barennes  ^{3,4} Vanessa Mhanna  ^{3,4} Hélène Vantomme  ^{3,4}
 Paul Sty,  ³ Encarnita Mariotti-Ferrandiz,  ^{3,†} David Klatzmann  ^{3,4,‡} Aleksandra M. Walczak,  ^{2,†,‡}
 Thierry Mora,  ^{2,†,‡} and Armita Nourmohammad  ^{1,5,6,7,8,†,‡}

¹*Max Planck Institute for Dynamics and Self-organization, Am Faßberg 17, 37077 Göttingen, Germany*

²*Laboratoire de physique de l'école normale supérieure (PSL University), CNRS, Sorbonne Université, and Université de Paris, 75005 Paris, France*

³*Sorbonne Université, INSERM, Immunology-Immunopathology-Immunotherapy (i3), F-75005 Paris, France*

⁴*AP-HP, Hôpital Pitié-Salpêtrière, Biotherapy (CIC-BTi), F-75651 Paris, France*

⁵*Department of Physics, University of Washington, 3910 15th Avenue Northeast, Seattle, Washington 98195, USA*

⁶*Paul G. Allen School of Computer Science and Engineering, University of Washington,*

85 E Stevens Way NE, Seattle, Washington 98195, USA

⁷*Department of Applied Mathematics, University of Washington, 4182 W Stevens Way NE, Seattle, Washington 98105, USA*

⁸*Fred Hutchinson Cancer Center, 1241 Eastlake Ave E, Seattle, Washington 98102, USA*



(Received 25 July 2023; accepted 16 February 2024; published 8 March 2024)

The adaptive immune response relies on T cells that combine phenotypic specialization with diversity of T-cell receptors (TCRs) to recognize a wide range of pathogens. TCRs are acquired and selected during T-cell maturation in the thymus. Characterizing TCR repertoires across individuals and T-cell maturation stages is important for better understanding adaptive immune responses and for developing new diagnostics and therapies. Analyzing a dataset of human TCR repertoires from thymocyte subsets, we find that the variability between individuals generated during the TCR V(D)J recombination is maintained through all stages of T-cell maturation and differentiation. The interindividual variability of repertoires of the same cell type is of comparable magnitude to the variability across cell types within the same individual. To zoom in on smaller scales than whole repertoires, we defined a distance measuring the relative overlap of locally similar sequences in repertoires. We find that the whole repertoire models correctly predict local similarity networks, suggesting a lack of forbidden T-cell receptor sequences. The local measure correlates well with distances calculated using whole repertoire traits and carries information about cell types.

DOI: [10.1103/PRXLIFE.2.013011](https://doi.org/10.1103/PRXLIFE.2.013011)

I. INTRODUCTION

The T-cell adaptive immune response leverages various cell subsets. CD8+ cells take on mostly a cytotoxic role, i.e., killing infected cells. CD4+ cells differentiate into two subsets, regulatory (Treg) and conventional (Tconv) cells. Tconv cells acquire effector helper function and help coordinate the immune response upon activation in the periphery. Tregs modulate the immune response by down-regulating the activity and response of different cells in the immune system.

To perform their distinct functions, CD8+ and CD4+ cells bind to different families of the major histocompatibility complex (MHC) molecule (class I and II, respectively). Therefore, the selective pressures exerted on their associated T-cell receptor (TCR) repertoires are believed to be markedly different. Indeed, previous studies [1–4] have described differential properties of the receptors associated with the CD4+ and CD8+ repertoires, reporting significant yet limited statistical differences, mostly related to the V and J gene usage. Similarly, small statistical differences have been reported in the repertoires of conventional (Tconv) versus regulatory (Treg) CD4+ T cells [4,5]. However, despite a recent study comparing the repertoires of different thymic subsets in mice [6], it is still unclear how thymic selection shapes the repertoires of distinct T-cell subsets, and how these selection forces vary across individuals.

To recognize the large variety of possible antigens, T cells express a broad diversity of TCRs generated by the random rearrangement of their α and β chains. Because this process is stochastic, it may generate receptors with undesirable properties, which must be vetted [7]. During their development in the thymus, T cells undergo a selection process that promotes receptors with good affinity to the MHC, to

*Present address: Peter Debye Institute for Soft Matter Physics, Leipzig University, Leipzig, Germany.

†Co-senior author.

‡These authors contributed equally to this work. Correspondence should be addressed to: aleksandra.walczak@phys.ens.fr, thierry.mora@phys.ens.fr, and armita@uw.edu

make sure that they have the minimum necessary recognition properties (*positive selection*). At the same time, receptors with a too strong affinity to self-peptides presented by the MHC are less likely to be released into the periphery, to limit an immune response against the self (*negative selection*). Since the MHC is highly polymorphic [8], this process is expected to be at least in part specific to each individual. Understanding how these processes shape the TCR repertoire and characterizing the heterogeneity within and between individuals is key to a better understanding of the fundamental aspects underlying autoimmunity as well as immune disorders.

High-throughput sequencing of TCR repertoires [9–15] has made it possible to study differences between different individuals as well as T-cell subsets, giving insight into their respective selection processes.

Because TCR repertoires are so diverse, typical samples still show substantial variability, even when they are from the same subset and same individual [16]. Comparisons between repertoires must be done at an aggregate level, by identifying statistical features that discriminate between them. A general strategy for this task is to define generative models of TCR sequences for each subset, and compare them using measures of divergence between their distributions [4,17]. We refer to this approach as *global* repertoire comparison. A previous study [17] of bulk TCR repertoires from blood samples showed only moderate variability across individuals, largely driven by differences in the statistics of the recombination process [18,19] rather than individual-specific selection, which contrasts with the idea that differences in MHC alleles drives the intrapopulation variation of the repertoires. It remains unclear to what extent the population variability in the recombination process impacts our ability to detect the differences between T-cell subsets due to functional selection.

Model-based comparisons, even when they rely on powerful neural networks, are usually dominated by some key statistics of the data, such as germline gene usage or CDR3 length and position-dependent amino-acid usage [4,17]. A possible concern is that they can miss important differences in repertoire properties, such as the depletion or enrichment of particular sequence motifs, which would strongly affect the structure of the repertoires locally in sequence space but may only have a limited impact on their global features as captured by the generative models. Previous studies have used local network measures of similarity between close-by sequences to detect immune stimuli or phenotypes shared by TCRs [20–26], suggesting that local differences may be important for discriminating between repertoires beyond global statistical features. We refer to such approaches as *local* repertoire comparisons.

We applied both global and local comparison methods to TCR repertoire data obtained from the thymi of nine organ donors, sorted into functional T-cell subsets at different levels of thymic maturation. By comparing these different subsets using a combination of model-based and local network analyses, we study how the T-cell repertoire evolves across the various stages of thymic development and quantify the heterogeneity between individuals.

II. RESULTS

A. Generative models of individual T-cell subsets

We analyze high-throughput TCR β repertoire sequence data of purified T-cell subsets from different stages of maturation [see Fig. 1(a) for a schematic of the different stages]. We analyzed purified samples of CD4+CD8+CD3+ (double positive, DP), CD3+CD8+(CD8+), CD3+CD4+CD25+ (Treg), and CD3+CD4+CD25- (Tconv) from thymic samples from $n = 9$ individuals. Details on sorting strategies and high-throughput sequencing steps are described in Ref. [27], where part of this dataset was presented. DP cells represent an early stage of development after successful recombination and selection of a receptor, prior to commitment to any functional fate. The statistics of their repertoire should closely follow that of the recombination process. During selection, cells with higher affinity to one of the two major classes of MHC differentiate into either CD8+ or CD4+ cells (Treg and Tconv).

We used the IGoR software [28] to infer a VDJ recombination model (called P_{gen}) for each of the $n = 9$ individuals from the unproductive TCR sequences pooled from all four subsets. Unproductive sequences are fossil records of unsuccessful recombination events on the second chromosome and are thus believed to be free of selection effects that impact productive sequences. To accurately account for selection, we then trained a SONIA model [17] on productive sequences for each subset, P_{DP} , P_{CD8} , P_{Tconv} , P_{Treg} , and each individual (9 \times 4 models in total). This model is built on top of the VDJ recombination model, $P = Q \times P_{\text{gen}}$, where Q is a subset-specific selection factor dependent on receptor sequence features, including V-, J- gene usages, and junction length and amino acid composition. As previously done in Ref. [4], we can characterize the diversity of repertoires by evaluating the entropy of the inferred models and estimate the similarities between the repertoires by computing the Jensen-Shannon divergence of the models [see Fig. 1(b) for a summary of the analysis strategy].

B. Thymic selection reduces sequence diversity by amplifying biases of VDJ recombination

We first quantified how the effective diversity of the repertoire changed throughout thymic development, by computing the Shannon entropy of each inferred model [Fig. 1(c), top]. This entropy is a measure of the diversity of TCRs at different developmental stages and can be thought of as the number of distinct TCR sequences. It is different from the raw number or clonality of TCR, both of which are subject to sampling biases. Following VDJ recombination, when receptors are well described by the P_{gen} distribution, the entropy is the highest at around 36 bits, corresponding to an effective sequence space of size $2^{36} \approx 10^{11}$. Entropy is significantly reduced at the DP stage (34 bits $\approx 10^{10}$ sequences), and even further at the single-positive stages of CD8+, Tconv, and Treg (31 bits $\approx 10^9$ sequences).

We asked whether this reduction of diversity results from existing biases in the VDJ recombination process, or is independent of them. The bottom part of Fig. 1(c) shows that entropy reduction is accompanied by an increasing correlation

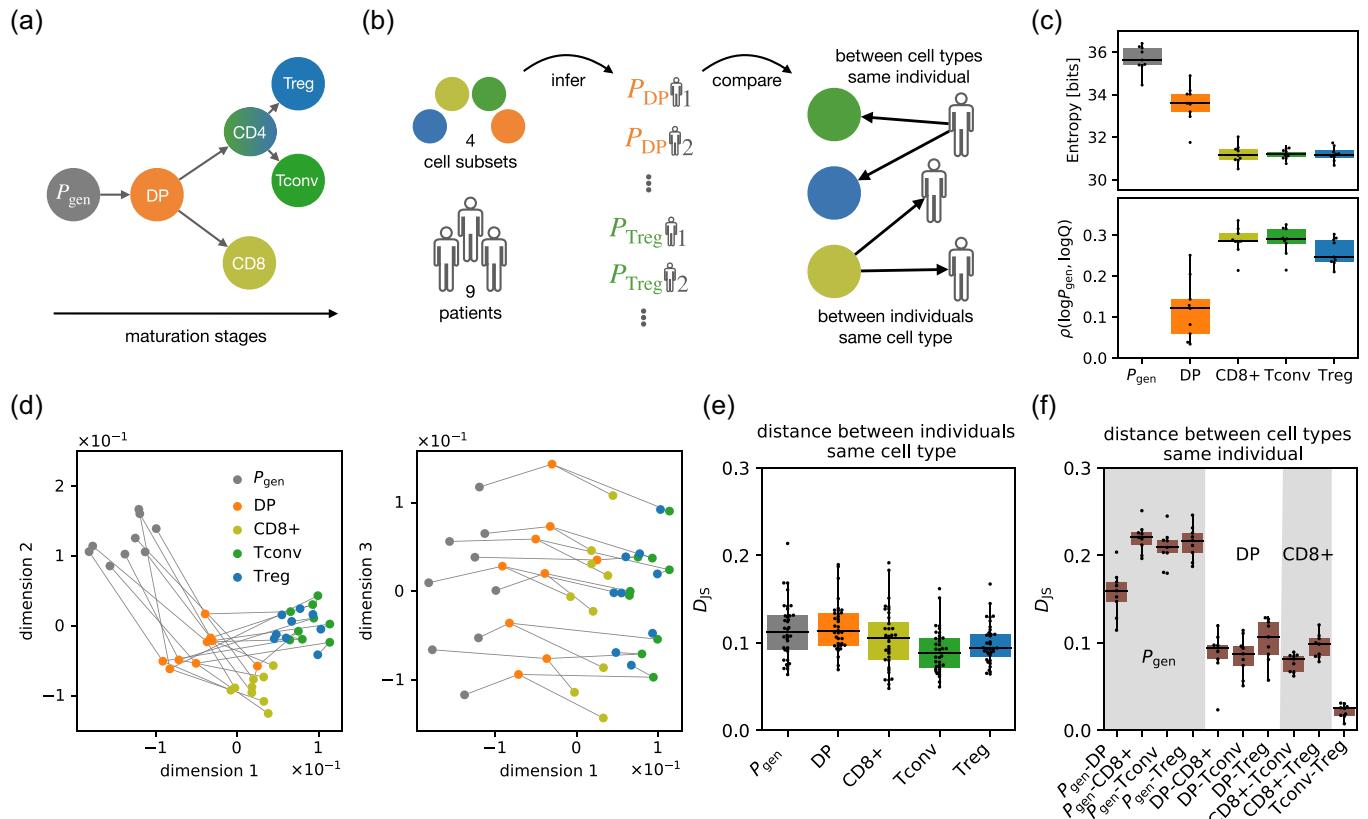


FIG. 1. Comparison of global repertoire features across individuals and thymic development stages. (a) Schematic of the thymic development stages at which repertoires were sequenced. P_{gen} denotes the output of the raw VDJ recombination process. Sequences are initially selected into the DP pool, and further differentiate into single-positive phenotypes: CD8+, and conventional and regulatory CD4+. (b) Analysis workflow. A model is inferred for each individual and cell subset, in addition to the P_{gen} model inferred from unproductive sequences. Models are then compared across cell types and individuals. (c) Entropy (top) and Pearson correlation between the logarithms of the generation probability and selection factor (bottom), as a function of the maturation stage. (d) Low-dimensional projection of distances between repertoires. Dimensionality reduction was performed using multidimensional scaling of the distance matrix defined by the Jensen-Shannon divergence D_{JS} between the inferred model distributions. Gray lines connect repertoires from the same individual. (e) D_{JS} between repertoires of the same cell type but from different individuals. (f) D_{JS} between repertoires from the same individual but of different cell types.

between P_{gen} and Q as maturation progresses, meaning that selection reinforces heterogeneities already present at generation, and thus reduces diversity through the “rich get richer” effect. This observation is consistent with reports of a similar correlation between P_{gen} and Q in fully matured TCR sampled from blood [29]. The results in Fig. 1(c) show the progression of this correlation during thymic selection, with onset as early as the DP stage.

C. Global comparison of repertoires between individuals and thymic maturation stages

We then estimated the similarity between repertoires at various stages of thymic development and in different individuals, by computing the Jensen-Shannon divergence D_{JS} (an information-theoretic measure of distance between distributions) between the inferred models. The resulting distance matrix [Fig. S1(A)] displays a complex structure, which is better interpreted by projecting repertoires into a distance-preserving low-dimensional map using multidimensional scaling (MDS) [30]. We find that three embedding

dimensions are sufficient to describe the main properties of the distance matrix [Fig. S1(B)]. To effectively visualize these embeddings in two dimensions, we rotated these three main axes in Fig. 1(d) such that the third dimension best aligns with the identity of individuals, which is tracked across subsets by gray lines. We observe a progression of thymic differentiation along dimensions one and two, with clusters corresponding to each stage, while dimension three delineates individuals. Notably, the VJ gene features alone cannot separate the intraindividual cell subsets [Fig. S6(B)]. These results suggest that individual heterogeneity imprinted by initial differences in VDJ recombination remains frozen throughout the maturation process. This finding is in agreement with previous analyses on unsorted repertoires [17], which showed that the strongest determinant of repertoire variability among individuals is the variation in VDJ recombination statistics for each person.

An alternative way to represent similarities between repertoires is to find a common encoding space for *sequences* (by using neural networks, as discussed in the Methods section) instead of whole repertoires, and then directly compare the lo-

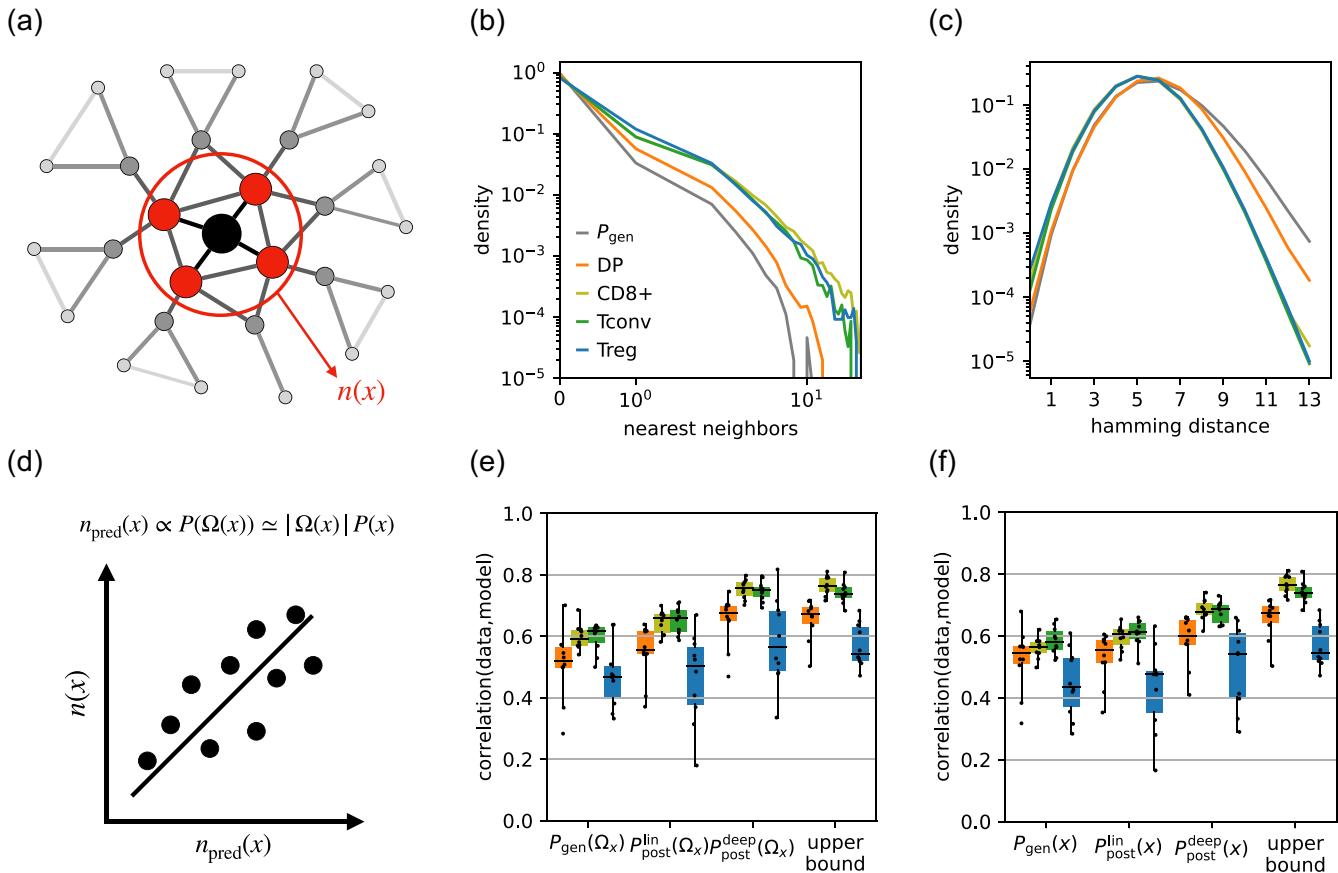


FIG. 2. Local properties of a repertoire. (a) Nearest neighbors of a sequence x are defined as sequences that differ by only one amino-acid substitution. (b) The probability that a sequence has exactly n nearest neighbors with the same VJ genes is shown for different repertoire subsets. As T cells experience selection the average number of nearest neighbors increases. The distribution is well reproduced by a synthetic repertoire sampled from the model P_{post} ; see Fig. S3. (c) Probability that two randomly sampled sequences with the same VJ gene combination have Hamming distance d_H between their amino-acid sequences is shown for different repertoire subsets. During thymic development, the average sequence distance decreases. The color code is similar to (b). (d) We quantify the accuracy of our model by computing the Pearson correlation between the predicted and the observed number of neighbors over the ten most probable VJ combinations. (e) The Pearson correlation between the true value for the number of nearest neighbors and the estimated values based on the model $n_{\text{pred}}(x)$, for P_{gen} , $P^{\text{lin}}_{\text{post}}$, and $P^{\text{deep}}_{\text{post}}$ models. The upper bound is computed by Poisson bootstrap resampling using the observed number of nearest neighbors as the rate parameter. (f) Similar to (e) but for an approximate predicted number of neighbors $\tilde{n}_{\text{pred}}(x)$, for which the model probabilities are evaluated using only the center sequence, instead of all the sequences in a local neighborhood. This approximation relies on the smoothness of the probability landscape. The prediction of the estimators remains significantly correlated with the observed number of neighbors in the data.

cation of different receptor sequences from different datasets (Fig. S2). While sequence variability is large and no clear separation of subsets is visible at the individual sequence level, we find the average location of receptor sequences in the encoding space to follow the expectations from the differentiation process (Fig. S2), consistent with the global repertoire analysis in Fig. 1(d).

In Figs. 1(e) and 1(f) we observe that variability of the same cell type between individuals is quantitatively comparable to that between cell types within a single individual. A more detailed quantification of differences across individuals with fixed cell types [Fig. 1(e)] shows that individual differences are stable across stages of maturation. Differences across cell types (with fixed individuals), shown in Fig. 1(f), follow the known hierarchy of thymic development, with the recombination model P_{gen} furthest from all subsets, but closer

to DP, and single-positive stages being equidistant from DP and from each other, with the exception of conventional and regulatory CD4+ cells, which are very similar.

D. Repertoires become concentrated as they mature in the thymus

While the previous analysis gives a general bird's-eye view of differences between repertoires, it is not informative about differences at the local sequence level. Following Ref. [27], we asked whether the sequence neighborhood of TCR sequences carried signatures of thymic selection. We define as “neighbors” sequences that have the same V and J genes, junction length, and differ by at most one amino acid (Hamming distance of 1); see Fig. 2(a). The larger the overall number of neighbors, the more “concentrated” is the

repertoire. Figure 2(b) shows the distribution of the number of neighbors of a random TCR from the repertoire, at different stages of T-cell development. The mean number of neighbors increases during thymic development, indicating that the repertoire gets increasingly concentrated around some preferred sequences. In other words, selection during thymic development amplifies neighboring sequences within certain regions of the sequence space, while depleting others, which results in repertoires with reduced overall diversity but with local agglomeration of sequences.

An alternative way to measure repertoire concentration is to look at the distribution of distances between any two receptor sequences [Fig. 2(c)]. Here, we consider the Hamming distance (number of unmatched amino acids in the junction) between sequences with the same V and J gene and junction length. Receptors tend to have on average lower Hamming distances following selection, for both CD4+ and CD8+ cells [Fig. 2(b)]. Together, these observations are consistent with the results of Ref. [27] on the CD8+ T cells obtained from the same datasets, and extend them to CD4+ T cells.

Although we observed that the diversity in repertoires declines with maturation [Fig. 1(c)], this global pruning of sequences did not necessarily imply that the pairwise distances between the receptors in a subset should decay with maturation. We asked whether this local behavior could be reproduced by our generative models. We pooled data from all patients for each subset to collect enough sequences to be able to train a soNNia (neural-network based) model [4] for the selection factor Q for each maturation stage (see Methods section). We find that both the distributions of the number of neighbors and the Hamming distances are well reproduced by the models (Fig. S3), indicating that these local neighborhood differences are well captured despite the models being trained on global features of the repertoires.

We then asked whether the model could predict the number of neighbors $n(x)$ of a particular sequence x , and not just their distribution. Assuming that each sequence is drawn at random independently of others, the number of neighbors should be distributed according to a Poisson distribution of mean [20,21]

$$n_{\text{pred}}(x) = \frac{N_{\text{VJ}}}{P(V, J)} \sum_{y \in \Omega(x)} P(y) \simeq \frac{N_{\text{VJ}} |\Omega(x)|}{P(V, J)} P(x), \quad (1)$$

where $P(x)$ is the model distribution over sequences x , $P(V, J)$ is the probability of picking a particular VJ pair estimated from the model, N_{VJ} is the observed number of sequences with that VJ pair, and $\Omega(x)$ is the set of all potential (not necessarily extant) neighbors of x . The approximation in the second equality relies on the assumption that the probability landscape is smooth, so that the probability of a neighboring sequence is on average the same as the focal sequence.

The accuracy of models can be evaluated by calculating the correlation between $n(x)$ and $n_{\text{pred}}(x)$ across all sequences with >3 neighbors [Fig. 2(d)]. Figure 2(e) compares the performance of three models across cell subsets: the naked recombination model P_{gen} , the linear (SONIA), and the neural-network (soNNia) selection models. The performance of the soNNia model reaches the Spearman correlation of $\rho \sim 0.75$ for CD8+ and Tconv subsets. The Treg and DP datasets have

smaller sizes and contain fewer nearest neighbors on average (Fig. S7), resulting in a nosier comparison to the model, which reduces the absolute correlation. The performance of the soNNia model compares well with the upper limit on predictability due to experimental noise, which we estimated using Poisson bootstrap resampling from data [Fig. 2(e)]. The smooth landscape approximation only moderately degrades predictability [Fig. 2(f)], while showing much faster computation times (by a factor of 19 times the mean junction length), thanks to the fact that it does not require computing the probability of each neighboring sequence. These observations carry over to the prediction of the number of second neighbors (i.e., sequences with at most two amino-acid differences) and beyond (Fig. S4).

Overall, these results show that the local properties of individual repertoires are well captured by the model and that the probability landscape of finding receptors sequences is relatively smooth as a function of sequence distance. If they were forbidden regions in the space of receptor sequences, we would expect a depletion of observed neighbors relative to the model in these regions. To test for this possibility more directly, we asked how well the model predicted the number of neighbors of sequences with high sequence probabilities (Fig. S8), and found good agreement, indicating that such forbidden regions either do not exist or are rare, when focusing on beta chains only.

E. Local differences between repertoires

Next, we asked how the neighborhood structures differed between cell types and individuals at the level of single sequences. For a given sequence x , we want to measure differences in the local network structure of nearest neighbors across different repertoires. We quantify these differences by defining a distance based on the spearman correlation between the number of neighbors of each sequence in the two datasets:

$$D_L(D_1, D_2) = \sqrt{1 - \rho[n_{D_1}(x), n_{D_2}(x)]}, \quad (2)$$

where $n_{D_1}(x)$ and $n_{D_2}(x)$ are the number of neighbors of sequence x in repertoires D_1 and D_2 .

Compared to the Jensen-Shannon divergence D_{JS} , this distance is based on the local sequence information only, and in principle is model independent. However, its accuracy strongly depends on sequencing depth, because it relies on counting the number of neighbors for each sequence, see Fig. S5(A). To overcome that difficulty, we used the estimators n_{pred} evaluated with the inferred models of each repertoire to compute the $D_L(D_1, D_2)$.

When applied to all possible pairs of repertoires, the resulting local distance D_L correlates well with the global distance between repertoires D_{JS} (Spearman $\rho = 0.68$); see Fig. 3(a) for results using $n_{\text{pred}}(x)$, and Fig. S5(C) using $n(x)$. Figures 3(b) and 3(c) show this local distance between individuals and between cell types. These statistics of local distances closely mirror those of global distance shown in Figs. 1(e) and 1(f). This means that differences in the local sequence structure of the repertoire follow global differences captured by the models, and that the global models fully capture the local statistics.

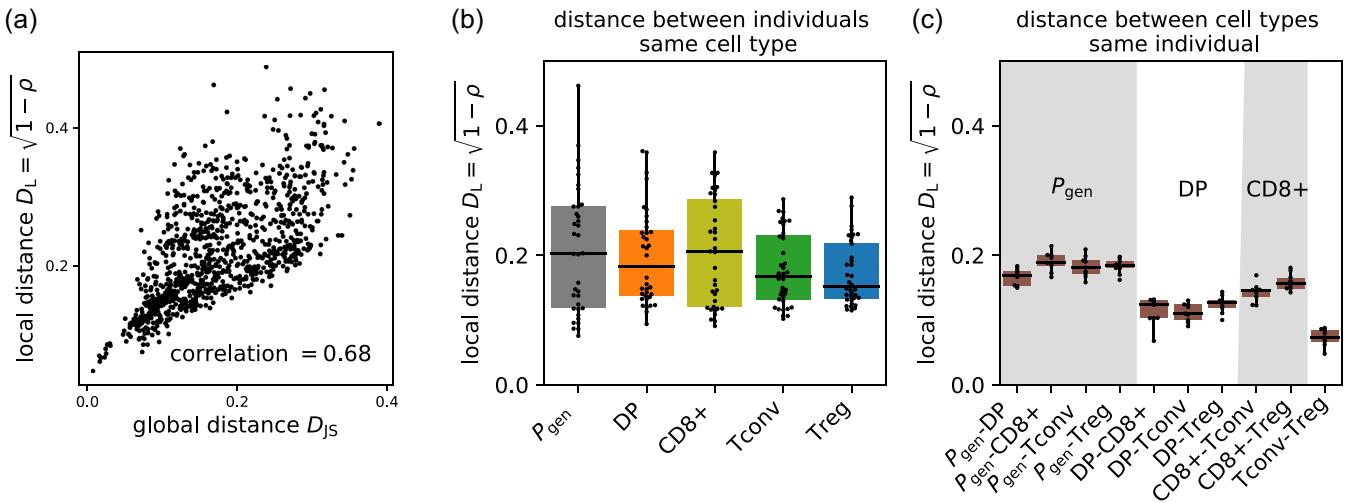


FIG. 3. Comparison of local repertoire features across individuals and thymic development stages. (a) Local distance D_L and global distance D_{JS} (Fig. 1) are significantly correlated. (b) Local distance D_L between repertoires of the same cell type but from different individuals. (c) Local distance D_L between repertoires from the same individual but of different cell types.

III. DISCUSSION

Previous work [1–4] has characterized global statistical differences between the T cell receptor β chain repertoires of fully mature CD4+ and CD8+ T cells. Here, we tracked how these differences emerge during thymic development. We find that the diversity of the repertoire shrinks as a result of thymic selection. This reduction occurs through the concentration of the repertoire around particular regions of the sequence space, consistent with the findings of Ref. [27] on CD8+ repertoires from the same dataset. These favored sequences are typically likely to be generated by VDJ recombination (high P_{gen}) even before thymic selection, suggesting that the recombination process has evolved to produce sequences that are likely to survive thymic selection. This idea was first proposed in Ref. [29] based on the analysis of peripheral repertoires. Our analysis shows how the process unfolds and amplifies from initial receptor recombination (P_{gen}) to the DP stage, and then further to each SP stage.

Network analyses of repertoires have been successfully applied in a variety of contexts [20,22,31–35]. Following Ref. [21], we showed that diversity reduction is accompanied with the concentration of the network around the high-degree nodes. Remarkably, our models can accurately predict the changes in these local network structures. Our best-achieving model soNNia uses a nonlinear artificial neural network architecture, allowing it to capture complex interactions. However, we expect the high dimensionality of the sequence space to limit that potential, in particular, if the repertoire landscape is irregular or “rugged,” with deep valleys and hills. The usual view of thymic selection is that it should deplete specificities to self-antigens. Since antigen-specific TCRs form clusters in sequence space [23,25], negative selection could have translated in the elimination of entire clusters corresponding to these forbidden self-antigen-specific TCRs. This in turn should have created valleys and shaped a rugged landscape in the sequence space. The success of the model in the face of this potential issue is supported by our observation that the

landscape is mostly smooth, as demonstrated by the ability to approximate the probability of a sequence by that of its neighbors. These findings relate to the generic structure of the sequence space and do not rule out the existence of less frequent but significant peaks or valleys. As demonstrated in previous studies [22,23,25,36], antigen presentation does lead to the formation of dense clusters of similar sequences. Our current research is based solely on the beta chain of the T-cell receptor, and it is important to note that neighboring receptors may exhibit different specificities due to variations in their alpha chains [23,25,36]. Consequently, the smoothness in sequence space that we have observed may not necessarily apply to data involving paired chains. Exploring this aspect represents a promising avenue for future research.

The elimination of specific TCRs is a complex phenomenon, likely multifactorial, integrating not just specificities but also other parameters that control the level of T-cell activation. Understanding how these observations can be reconciled with the classical view of negative selection [6] remains an interesting direction for investigation.

We further asked whether the local neighborhood of a sequence carries information about its cell type. We defined a distance between repertoires based on the similarity of their neighborhood structure, which correlates well with the global, model-based Jensen-Shannon divergence, and recapitulates the hierarchy between the different stages of thymic development. While our neighbor-based distance is in principle model free, in practice we could only evaluate it reliably using our data-driven generative models because of sampling issues. Nonetheless, our results suggest possible ways to use local network information to compare repertoire subsets, and to study their dependence on cell type, health condition, or age.

Local network structures can be used to detect responding clones during an infection, by looking for sequences with more neighbors than expected in a single repertoire [22]. Our results on comparing repertoires with a local neighborhood distance suggest that these local differences could also be

used to identify sequences that are particularly enriched in one repertoire versus another. This could allow us to define sequences that are characteristic of particular repertoires, and use them to better understand the relationship between cell subset and TCR, with potential applications for diagnosis and phenotyping.

IV. METHODS

A. Multidimensional scaling

The objective of multidimensional scaling is to find a lower-dimensional representation of data that preserves the similarity between samples in the dataset. Given D data points and precomputed distances d_{ij} between points, it infers data coordinates (x_1, \dots, x_D) with $x_i \in \mathbb{R}^N$ and N small by minimizing an objective function called stress:

$$\text{Stress}(x_1, \dots, x_D) = \sqrt{\sum_{i \neq j} (d_{ij} - \|x_i - x_j\|_2)^2}. \quad (3)$$

B. Inference of a representation space for TCR sequences

In order to better visualize the differences between repertoire subsets from different stages of maturation in the thymus, we develop a method to map the receptor sequences in a representation space that carries information about selection. We build a feed-forward neural network that outputs the selection factor for each subset and has a hidden layer of dimension 2; see Fig. S2(A) for a sketch of the network architecture. We infer selection factors by maximizing the joint objective

$$\mathcal{L}(\theta) = \sum_{t \in \mathcal{T}} \mathbb{E}_t[-E_\theta^t] - \log \mathbb{E}_{\mathcal{G}}[e^{-E_\theta^t}], \quad (4)$$

where \mathcal{G} is the set of sequences generated by the P_{gen} model, θ the model parameters, E_θ the energy that the model assigns to the sequence, and, with abuse of notation, we identify with \mathcal{T} the cell types of the thymic samples and the corresponding dataset. It is important to clarify that this architecture is not an autoencoder because we are not trying to reconstruct the distribution of sequences, but we are only interested in characterizing the selection factor. For this reason, the representation space will carry information only about selection. We implement the model using the KERAS software [37] and infer its parameters using the RMSprop stochastic gradient descent algorithm [38]. After inference, we map all sequences to the two-dimensional hidden space. As it can be seen in Fig. S2(B), we do not observe any clear separation between cell types in this representation space. On the other hand, the averages of the distributions are organized in a clear one-dimensional subspace that follows increasing selection.

C. Smoothness of P_{post}

If we assume that the postselection probability $P_{\text{post}} = P_{\text{gen}}Q$ does not vary significantly in a local region of sequence space, we can ask whether our estimators can predict the number of observed neighbors $N^{\text{obs}}(x)$ of a receptor sequence defined by higher cutoffs in the Hamming distance.

In Fig. S4(A) we show that the estimators $P_{\text{post}}^{\text{lin}}$ (linear selection model), $P_{\text{post}}^{\text{deep}}$ (deep selection model), and P_{gen} (preselection generation model) perform reasonably well also in this regime. Surprisingly, there is no loss in performance for the threshold at Hamming distance 2. For higher cutoffs the average performance decreases and the standard deviation increases. The estimators remain, however, significantly correlated with $N^{\text{obs}}(x)$.

The previous result is consistent with the idea that P_{post} is smooth in the neighborhood of a given sequence. We are then motivated to push this assumption even further to define alternative estimators for $N^{\text{obs}}(x)$. The first approximation we can perform is to assume that all neighbors have similar probability, as it is done in the main text. This estimator is computationally more efficient as it requires only a single evaluation of the probability.

Alternatively, we can assume that selection factors do not vary considerably within a neighborhood [see Figs. 2(e), 2(f), and S4] and approximate $P_{\text{post}}(nn_x) \simeq P_{\text{gen}}(nn_x)Q(x)$ for the postselection models. Since $P_{\text{gen}}(nn_x)$ can be efficiently estimated via dynamic programming [39], the resulting estimator turns out to be more efficient than the exact one, yet not as much as $P_{\text{post}}(x)$. We compare the two approximations in Fig. S4(B) for $P_{\text{post}}^{\text{deep}}$. Their performance is approximately comparable to the $P_{\text{post}}^{\text{lin}}(nn_x)$ for the smallest cutoff value. They generically perform better than $P_{\text{gen}}(nn_x)$ for all cutoff values that we tested. As expected the estimator that assumes smoothness only in selection Q slightly outperforms the one that assumes smoothness of the whole probability P_{post} .

In conclusion, we find that the three best choices for estimating the number of nearest neighbors are $P_{\text{post}}^{\text{deep}}(x)$, $P_{\text{gen}}(nn_x)Q_{\text{post}}^{\text{deep}}(x)$, and $P_{\text{post}}^{\text{deep}}(nn_x)$. The three estimators have increasing performance but are also more computationally expensive: the choice of which one to use will thus depend on the specific application and the amount of available data. Since these three estimators perform better at higher cutoffs in Hamming distance than the original $P_{\text{gen}}(nn_x)$ [20] at the smallest cutoff, we expect that integrating information from higher cutoffs in Hamming will increase the statistical power of enrichment analyses.

D. Global distance between repertoires D_{JS}

The Jensen-Shannon divergence D_{JS} is a symmetric measure between two probability distributions. It can be used to quantify the difference between any two repertoires defined by the postselection probabilities $P_{\text{post}}^r = Q^r P_{\text{gen}}^r$ and $P_{\text{post}}^{r'} = Q^{r'} P_{\text{gen}}^{r'}$:

$$D_{\text{JS}}(r, r') = \frac{1}{2} \left\langle \log_2 \frac{2P_{\text{post}}^r}{P_{\text{post}}^r + P_{\text{post}}^{r'}} \right\rangle_r \quad (5)$$

$$+ \frac{1}{2} \left\langle \log_2 \frac{2P_{\text{post}}^{r'}}{P_{\text{post}}^r + P_{\text{post}}^{r'}} \right\rangle_{r'}, \quad (6)$$

where $\langle \cdot \rangle_r$ denotes the expectation value with respect to P_{post}^r . In practice, we estimate $\langle \cdot \rangle_r$ as an empirical average by sampling 20 000 sequences from P_{post}^r .

E. Local distance between repertoires D_L

The local distance D_L between repertoires, as defined in Eq. (2), can be directly evaluated on sequences shared between repertoires by computing the correlation between the observed number of neighbors in the respective datasets. In Fig. S5(A), we evaluate the D_L between independent subsamples from the same dataset. We show that for typical repertoire sizes that are found in the literature (10^4 – 10^5 unique receptor sequences), the estimator is too noisy to resolve the observed differences between repertoires [Figs. 3(b) and 3(c)], which can be as small as $D_L \sim 0.1$ bits.

In order to overcome this difficulty we evaluate the estimator using the expected number of neighbors n_{pred} computed by the models. We find that D_L converges quickly as a function of the number of sequences evaluated; see Fig. S5(B). In the main analysis we evaluate D_L with 200 sequences sampled from each dataset.

As sequencing technologies improve, we expect the model-free estimator to successfully be used to compare the local structure of different repertoires. As a proof of concept, we compare the model-free version of the D_L estimator to the global distance D_{JS} for the biggest datasets present in our data (six datasets with at least 80 000 unique amino-acid sequences) and we are able to reproduce the results of Fig. 3(a), see Fig. S5(C).

The code used to reproduce the figures can be found in Ref. [40]. All supplementary figures can be found in Ref [41].

ACKNOWLEDGMENTS

We thank Thomas Dupic for discussions during the development of the neighbor analysis. The study was supported by the European Research Council Grant COG 724208 (A.M.W., T.M., and G.I.), the European Research Council Grant TriPoD ERC Advanced 322856 (D.K.), the Agence Nationale de la Recherche Grant ANR-19-CE45-0018 “RESP-REP” (A.M.W., and T.M.), the Agence Nationale de la Recherche Grant ANR-16-RHUS-0001 “RHU iMAP” (D.K.), the Agence Nationale de la Recherche Grant ANR-11-IDEX-0004-02 “LabEx Transimmunom” (D.K.), the Deutsche Forschungsgemeinschaft Grant CRC 1310 on Predictability in Evolution (A.M.W., A.N., T.M., and G.I.), the National Institutes of Health Award R35 GM142795 (A.N.), the National Science Foundation award (CAREER award to AN; Grant No: 2045054), the Max Planck Research Group funding through the Max Planck Society (A.N., and G.I.), the Horizon 2020 Award 825821 for the iReceptorPlus project (D.K., and E.M.F.), and the European Research Area Network on Cardiovascular Diseases Grant ANR-18-ECVD-0001 “AIR-MI” (E.M.F.). This work benefited from equipment and services from the iGenSeq core facility at Institut du Cerveau (ICM, Paris, France) for all the data production. This work is also supported, in part, through the Departments of Physics, and the College of Arts and Sciences at the University of Washington.

- [1] R. Emerson, A. Sherwood, C. Desmarais, S. Malhotra, D. Phippard, and H. Robins, Estimating the ratio of CD4+ to CD8+ T cells using high-throughput sequence data, *J. Immunol. Methods* **391**, 14 (2013).
- [2] H. M. Li, T. Hiroi, Y. Zhang, A. Shi, G. Chen, S. De, E. J. Metter, R. Wood, H. William, A. Sharov, J. D. Milner, K. G. Becker, M. Zhan, and N.-p. Weng, Tcr β repertoire of CD4+ and CD8+ T cells is distinct in richness, distribution, and CDR3 amino acid composition, *J. Leukoc. Biol* **99**, 505 (2016).
- [3] J. A. Carter, J. B. Preall, K. Grigaityte, S. J. Goldfless, E. Jeffery, A. W. Briggs, F. Vigneault, and G. S. Atwal, Single T cell sequencing demonstrates the functional role of $\alpha\beta$ TCR pairing in cell lineage and antigen specificity, *Front. Immunol.* **10**, 1516 (2019).
- [4] G. Isacchini, A. M. Walczak, T. Mora, and A. Nourmohammad, Deep generative selection models of T and B cell receptor repertoires with soNNia, *Proc. Natl. Acad. Sci. USA* **118**, e2023141118 (2021).
- [5] K. A. Lagattuta, J. B. Kang, A. Nathan, K. E. Pauken, A. H. Jonsson, D. A. Rao, A. H. Sharpe, K. Ishigaki, and S. Raychaudhuri, Repertoire analyses reveal T cell antigen receptor sequence features that influence T cell fate, *Nat. Immunol.* **23**, 446 (2022).
- [6] F. Camaglia, A. Ryvkin, E. Greenstein, S. Reich-Zeliger, B. Chain, T. Mora, A. M. Walczak, and N. Friedman, Quantifying changes in the T cell receptor repertoire during thymic development, *eLife* **12**, e81622 (2023).
- [7] A. J. Yates, Theories and quantification of thymic selection, *Front. Immunol.* **5**, 13 (2014).
- [8] K. Murphy, K. Murphy, P. Travers, M. Walport, C. Janeway, C. Mauri, and M. Ehrenstein, *Janeway’s Immunobiology*, Janeway’s Immunobiology No. 978, No. 0-4129 (Garland Science, New York, 2008).
- [9] X.-L. Hou, L. Wang, Y.-L. Ding, Q. Xie, and H.-Y. Diao, Current status and recent advances of next generation sequencing techniques in immunological repertoire, *Genes Immun.* **17**, 153 (2016).
- [10] G. Georgiou, G. C. Ippolito, J. Beausang, C. E. Busse, H. Wardemann, and S. R. Quake, The promise and challenge of high-throughput sequencing of the antibody repertoire, *Nat. Biotechnol.* **32**, 158 (2014).
- [11] D. A. Bolotin, S. Poslavsky, I. Mitrophanov, M. Shugay, I. Z. Mamedov, E. V. Putintseva, and D. M. Chudakov, MiXCR: Software for comprehensive adaptive immunity profiling, *Nat. Methods* **12**, 380 (2015).
- [12] J. R. McDaniel, B. J. DeKosky, H. Tanno, A. D. Ellington, and G. Georgiou, Ultra-high-throughput sequencing of the immune receptor repertoire from millions of lymphocytes, *Nat. Protoc.* **11**, 429 (2016).
- [13] B. J. DeKosky, G. C. Ippolito, R. P. Deschner, J. J. Lavinder, Y. Wine, B. M. Rawlings, N. Varadarajan, C. Giesecke, T. Dörner, S. F. Andrews, P. C. Wilson, S. P. Hunicke-Smith, C. G. Willson, A. D. Ellington, and G. Georgiou, High-throughput sequencing of the paired human immunoglobulin heavy and light chain repertoire, *Nat. Biotechnol.* **31**, 166 (2013).
- [14] M. A. Turchaninova, O. V. Britanova, D. A. Bolotin, M. Shugay, E. V. Putintseva, D. B. Staroverov, G. Sharonov, D. Shcherbo, I. V. Zvyagin, I. Z. Mamedov, C. Linnemann, T. N.

Schumacher, and D. M. Chudakov, Pairing of T-cell receptor chains via emulsion PCR, *Eur. J. Immunol.* **43**, 2507 (2013).

[15] B. J. Dekosky, T. Kojima, A. Rodin, W. Charab, G. C. Ippolito, A. D. Ellington, and G. Georgiou, In-depth determination and analysis of the human paired heavy- and light-chain antibody repertoire, *Nat. Med.* **21**, 1 (2014).

[16] M. B. Koraichi, S. Ferri, A. M. Walczak, and T. Mora, Inferring the T cell repertoire dynamics of healthy individuals, *Proc. Natl. Acad. Sci. USA* **120**, e2207516120 (2023).

[17] Z. Sethna, G. Isacchini, T. Dupic, T. Mora, A. M. Walczak, and Y. Elhanati, Population variability in the generation and selection of T-cell repertoires, *PLOS Comput. Biol.* **16**, e1008394 (2020).

[18] S. Tonegawa, Somatic generation of antibody diversity, *Nature (London)* **302**, 575 (1983).

[19] M. M. Davis and P. J. Bjorkman, T-cell antigen receptor genes and T-cell recognition, *Nature (London)* **334**, 395 (1988).

[20] M. V. Pogorelyy, A. A. Minervina, M. Shugay, D. M. Chudakov, Y. B. Lebedev, T. Mora, and A. M. Walczak, Detecting T-cell receptors involved in immune responses from single repertoire snapshots, *PLoS Biol.* **17**, e3000314 (2019).

[21] P.-G. Ritvo, A. Saadawi, P. Barennes, V. Quiniou, W. Chaara, K. El Soufi, B. Bonnet, A. Six, M. Shugay, E. Mariotti-Ferrandiz, and D. Klatzmann, High-resolution repertoire analysis reveals a major bystander activation of TFH and TFR cells, *Proc. Natl. Acad. Sci. USA* **115**, 9604 (2018).

[22] M. V. Pogorelyy and M. Shugay, A framework for annotation of antigen specificities in high-throughput T-cell repertoire sequencing studies, *Front. Immunol.* **10**, 2159 (2019).

[23] P. Dash, A. J. Fiore-Gartland, T. Hertz, G. C. Wang, S. Sharma, A. Souquette, J. C. Crawford, E. B. Clemens, T. H. O. Nguyen, K. Kedzierska, N. L. La Gruta, P. Bradley, and P. G. Thomas, Quantifiable predictive features define epitope-specific T cell receptor repertoires, *Nature (London)* **547**, 89 (2017).

[24] K. Mayer-Blackwell, S. Schattgen, L. Cohen-Lavi, J. C. Crawford, A. Souquette, J. A. Gaever, T. Hertz, P. G. Thomas, P. Bradley, and A. Fiore-Gartland, TCR meta-clonotypes for biomarker discovery with *tcrlist3* enabled identification of public, HLA-restricted clusters of SARS-CoV-2 TCRs, *eLife* **10**, e68605 (2021).

[25] J. Glanville, H. Huang, A. Nau, O. Hatton, L. E. Wagar, F. Rubelt, X. Ji, A. Han, S. M. Kram, C. Pettus, N. Haas, C. S. L. Arlehamn, A. Sette, S. D. Boyd, T. J. Scriba, O. M. Martinez, and M. M. Davis, Identifying specificity groups in the T cell receptor repertoire, *Nature (London)* **547**, 94 (2017).

[26] H. Huang, C. Wang, F. Rubelt, T. J. Scriba, and M. M. Davis, Analyzing the mycobacterium tuberculosis immune response by T-cell receptor clustering with GLIPH2 and genome-wide antigen screening, *Nat. Biotechnol.* **38**, 1194 (2020).

[27] V. Quiniou, P. Barennes, V. Mhanna, P. Stys, H. Vantomme, Z. Zhou, F. Martina, N. Coatnoan, M. Barbie, H.-P. Pham, B. Clémenceau, H. Vie, M. Shugay, A. Six, B. Brandao, R. Mallone, E. Mariotti-Ferrandiz, and D. Klatzmann, Human thymopoiesis produces polyspecific CD8+ α/β T cells responding to multiple viral antigens, *eLife* **12**, e81274 (2023).

[28] Q. Marcou, T. Mora, and A. M. Walczak, High-throughput immune repertoire analysis with IGoR, *Nat. Commun.* **9**, 561 (2018).

[29] Y. Elhanati, Jr., A. Murugan, C. G. Callan, T. Mora, and A. M. Walczak, Quantifying selection in immune receptor repertoires, *Proc. Natl. Acad. Sci. USA* **111**, 9875 (2014).

[30] I. Borg and P. J. F. Groenen, *Modern Multidimensional Scaling - Theory and Application*, Series in Statistics (Springer, New York, 1997).

[31] R. Ben-Hamo and S. Efroni, The whole-organism heavy chain B cell repertoire from zebrafish self-organizes into distinct network features, *BMC Syst. Biol.* **5**, 27 (2011).

[32] R. J. M. Bashford-Rogers, A. L. Palsler, B. J. Huntly, R. Rance, G. S. Vassiliou, G. A. Follows, and P. Kellam, Network properties derived from deep sequencing of human B-cell receptor repertoires delineate B-cell populations, *Genome Res.* **23**, 1874 (2013).

[33] A. Madi, A. Poran, E. Shifrut, S. Reich-Zeliger, E. Greenstein, I. Zaretsky, T. Arnon, F. V. Laethem, A. Singer, J. Lu, P. D. Sun, I. R. Cohen, and N. Friedman, T cell receptor repertoires of mice and humans are clustered in similarity networks around conserved public CDR3 sequences, *eLife* **6**, e22057 (2017).

[34] E. Miho, R. Roškar, V. Greiff, and S. T. Reddy, Large-scale network analysis reveals the sequence space architecture of antibody repertoires, *Nat. Commun.* **10**, 1321 (2019).

[35] T. Ronel, M. Harries, K. Wicks, T. Oakes, H. Singleton, R. Dearman, G. Maxwell, and B. Chain, The clonal structure and dynamics of the human T cell response to an organic chemical hapten, *eLife* **10**, e54747 (2021).

[36] A. Mayer, Jr. and C. G. Callan, Measures of epitope binding degeneracy from T cell receptor repertoires, *Proc. Natl. Acad. Sci. USA* **120**, e2213264120 (2023).

[37] F. Chollet, Keras, 2015, <https://keras.io/>.

[38] T. Tieleman and G. Hinton, Lecture 6.5-rmsprop: Divide the Gradient by a Running Average of Its Recent Magnitude, COURSERA: Neural Networks for Machine Learning **4**, 26 (2012).

[39] Z. Sethna, Jr., Y. Elhanati, C. G. Callan, A. M. Walczak, and T. Mora, OLGA: Fast computation of generation probabilities of B- and T-cell receptor amino acid sequences and motifs, *Bioinformatics* **35**, 2974 (2019).

[40] https://github.com/statbiophys/Global_and_Local_Variability_2023.

[41] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PRXLIFE.2.013011> for additional data analysis and figures.