Immunological Diversity with Similarity

Rohit Arora, PhD¹, Harry M. Burke², and Ramy Arnaout, MD, DPhil^{1,3*}

¹Division of Clinical Pathology, Department of Pathology and ²Division of Clinical Informatics,
Department of Medicine, Beth Israel Deaconess Medical Center, Boston, MA 02215. ³Harvard
Medical School, 25 Shattuck St, Boston, MA 02115

*To whom correspondence should be addressed at rarnaout@gmail.com

7 Abstract

8 A diverse immune repertoire is considered a hallmark of good health, but measuring diversity
 9 requires a framework that incorporates not only sequences' relative frequencies but also their
 10 functional similarity to each other. Using experimentally measured dissociation constants from
 11 over 1,300 antibody-antigen and T-cell receptor (TCR)-peptide pairs, we developed a
 12 framework for functional immunological diversity based on binding and applied it to nearly 400
 13 high-throughput antibody and TCR repertoires to reveal patterns in immunological
 14 memory, infection, vaccination, and aging. We show that functional diversity adds information
 15 that is not captured by raw diversity, revealing signatures of e.g. clonal selection, and that unlike
 16 raw diversity, functional diversity is a robust measure that does not require correction for
 17 sampling error. Finally, we show that according to functional diversity, unlike raw diversity,
 18 individuals' repertoires overlap substantially, indicating a definable ceiling for the functional
 19 diversity of human adaptive immunity. Similarity redefines diversity in complex systems.

20 Immune repertoires are famously diverse. Collectively, a person's $\sim 10^{12}$ B and T cells express
 21 many millions of unique recombined antibody and TCR genes as part of millions of clonal
 22 lineages, more unique sequence than in the entire germline genome¹. At the sequence level,
 23 the repertoires of any two people overlap by only a fraction of a percent, indicating still higher
 24 diversity in the population^{2,3}. Yet repertoires are formed from V, D, and J gene segments that
 25 almost all people share and that are expressed at similar frequencies across individuals, and
 26 most repertoires are shaped by similar antigenic exposures and a consequent need to
 27 recognize and bind similar targets^{4,5}. Squaring the diversity that is seen with the similarity that
 28 must exist is a major goal of high-throughput immunology.

29 This goal has relevance for disease stratification and clinical management across a range of
 30 conditions. B- and T-cell diversity fall with age, as specific exposures expand a few select
 31 lineages at the expense of others⁶. Chronic infection appears to have a similar effect, impairing
 32 vaccination¹. Low B-cell diversity is associated with physiological frailty, a syndrome seen
 33 alongside conditions that have traditionally been considered to be unrelated to adaptive
 34 immunity (e.g., atherosclerotic cardiovascular disease), independent of chronological age⁷. In
 35 cancer, a rise in sequence-level T-cell diversity is thought to predict a successful response to
 36 immune-checkpoint inhibitors, drugs that make tumors more visible to the immune system⁸.

37 Two key contributors to any intuitive measure of diversity are frequency and similarity, but
 38 traditionally only frequency has factored into standard measures of immunological diversity. The
 39 simplest measure, species richness—the number of different sequences, lineages, or clones in
 40 a sample—ignores frequency entirely. Yet in a practical sense, a repertoire with a single
 41 dominant (e.g., leukemic) clone should be considered less diverse than a repertoire that has the
 42 same number of clones but no dominant clone. To incorporate frequency into measurements of
 43 diversity, there exist a family of measures—species richness (as a limiting case), Shannon
 44 entropy, the Simpson index (and related Gini coefficient), and the Berger-Parker index, among
 45 others—that differ in how much weight each measure places on frequency: i.e., in how much
 46 more a large clone adds to the total diversity than a small one⁹. Mathematically, this family is
 47 derived from the single master equation

$$48 \quad (1) \quad {}^qD = (\sum_i p_i^q)^{1/(1-q)}$$

49 in which the different measures correspond to differences in the “frequency parameter” q (e.g.
 50 $q=0$ corresponds to species richness; $q=1$ to Shannon entropy), qD is read “ D - q ” (“diversity of

51 order q ”), and i indexes the different species present in the sample. No single diversity measure
 52 is “best:” the different measures are complementary single-parameter summaries of the
 53 underlying species-frequency distribution¹⁰. qD for immunological diversity is subject to
 54 considerable sampling error, but robust methods exist for correcting for this error and are
 55 becoming standard in the field^{11,12}.

56 In the present work our goal was to incorporate and assess the utility of the second key
 57 contributor: species similarity. A repertoire made up of all-different sequences is intuitively more
 58 diverse than a repertoire that has the same number of sequences, present at the same
 59 frequencies as in the first repertoire, but all drawn from the same lineage or clone. Although our
 60 focus here is immunological diversity, similarity is a key contributor to diversity in many complex
 61 systems, including metagenomics (species vs. operational taxonomic units), machine learning
 62 (similarities within training data vs. the size of the training set), and transcriptomics (cell-to-cell
 63 variation vs. cell types; see Discussion). In the immune-repertoire literature, this point is
 64 sometimes addressed indirectly by grouping sequences together before measuring diversity, for
 65 example by clustering reads, collapsing clones, or binning by V(D)J segment usage^{1,5,13–15}.
 66 However, grouping usually imposes a binary threshold—in or out—on what is by nature a
 67 continuous and overlapping relationship among sequences and their encoded proteins.
 68 Grouping also usually zeros out or ignores any diversity that might exist within groups. It is
 69 unclear what is lost by ignoring similarity, or what might be gained from a more complete
 70 synthesis of diversity with similarity. Therefore we sought to develop and explore a continuous
 71 framework for immunological diversity incorporating similarity, and to test its utility in situations
 72 of clinical and research interest.

73 Results

74 **Framework.** We measured diversity-with-similarity on high-throughput B- and T-cell repertoires
 75 using a robust mathematical framework initially proposed for studying diversity in ecology and
 76 environmental settings¹⁶:

$$77 \quad (2) \quad {}^qD_s = \left(\sum_i p_i S_i^{q-1} \right)^{1/(1-q)}$$

78 where

$$79 \quad (3) \quad S_i = \sum_j p_j S_{ij}$$

80 Here the key innovation relative to Eq. 1 is a similarity matrix S , whose entries S_{ij} quantify how
 81 similar each pair of species (sequences, clonotypes, etc.) is. This framework provides “with-
 82 similarity” counterparts for species richness and the frequency-weighted diversity measures:
 83 species richness with similarity (0D_s , which places a very small weight on frequency, and ${}^{\infty}D_s$,
 84 which like 0D ignores frequency altogether), the exponential form of entropy with similarity (1D_s ,
 85 henceforth simply “entropy with similarity,” and likewise for other named measures), and so on.
 86 In qD_s notation, q as before is the frequency parameter, D_s denotes diversity-with-similarity, and
 87 D without the subscript means diversity without similarity, which we refer to as “raw diversity”
 88 (see Methods for more on notation).

89 Constructing the similarity matrix necessitates a choice of similarity measure. (Note the
 90 difference between similarity measures and diversity measures: a similarity measure is used to
 91 build the similarity matrix, which then is used to calculate diversity measures.) The choice of
 92 similarity measure is up to the investigator and depends on the biological feature(s) of interest.
 93 For example to study somatic hypermutation, one might use the Hamming distance. Just as
 94 different weightings provide complementary information about rare vs. frequent species—for
 95 example, the number of new thymic emigrants (species richness; 0D or ${}^{\infty}D_s$) vs. large leukemic
 96 clones (Berger-Parker index; ${}^{\infty}D$)—different similarity measures are expected to reveal different
 97 systems-level features of repertoires’ sequence-level configuration. Also as with raw diversity
 98 measures, expressing results for the new diversity-with-similarity measures as effective
 99 numbers, also known as number equivalents^{9,17–19}, as opposed to as bits or nats (for 1D_s) or as
 100 various fractions (for ${}^{>1}D_s$), makes it possible to compare them to each other, regardless of
 101 weighting or similarity measure, on a single intuitive scale (Box 1).

102

103 **Similarity measure.** We were interested in the single most fundamental mechanistic feature of
 104 antibodies and TCRs: binding to specific targets (Fig. 1). Therefore for our similarity measure,
 105 we chose a proxy for binding affinity that follows from the empirically observed changes in
 106 dissociation constant (K_d) associated with amino-acid substitution in antibody and TCR CDRs²⁰.
 107 We found that on average, a single amino-acid substitution at an antibody-antigen or TCR-
 108 peptide binding surface lowers affinity by 4-5 fold (geometric mean), with a long tail
 109 corresponding to rare orders-of-magnitude effects (Fig. 2a). We focused on CDR3, the third
 110 complementarity determining region, of IgH and TCR β , since this is the single most important
 111 contributor to binding specificity²¹; however, our approach can be applied to other regions (or
 112 indeed other protein families). Because the relationship between sequence and specificity

remains non-predictive and therefore complex, for any given sequence pair the similarity imputed from the observed distribution will be approximate; however, averaged over the many millions of pairs in the average high-throughput repertoire, it was expected to be a reasonably accurate first-pass repertoire-level view of immunological diversity with binding similarity. Using this similarity measure, diversity-with-similarity is interpreted as the effective number of sequences in a repertoire if the sequences were equally common and had no binding overlap with each other (Box 1), or equivalently, the number of equally common non-overlapping binding targets that a repertoire can recognize. We therefore refer to this choice of diversity-with-similarity as “functional diversity” (Fig. 1). Functional diversity can be interpreted in the context of a “shape space”²² that contains all possible CDR3 binding targets, with nearby targets having similar three-dimensional shapes and conformations (Fig. 1a). Each CDR3 binds a (possibly overlapping) subset of targets; together, a repertoire's CDR3s cover some part of shape space (Fig. 1b). Functional diversity measures the size of this region, controlling for similarity and overlap in binding among different CDR3s (Fig. 1c).

127

Validity. We first established that our similarity measure behaved sensibly by testing that closely related sequences scored close to 1 and unrelated sequences scored close to zero (in a continuous manner, following from an assumption of multiplicative independence and the data in Fig. 2a; see Methods) (Fig. 2b). We next established that it resulted in intuitive values for functional similarity by testing against expectations on simple *in silico* repertoires. In a representative test, we constructed four repertoires with 34 unique sequences each and 752 sequences total (Fig. 2c-d). In each repertoire, a few sequences were common (larger circles) while most were rare (smaller circles), representing the long-tailed frequency distribution seen in real repertoires^{3,23}. Importantly, the species-frequency distribution for all four repertoires was identical, meaning that raw diversity was also identical across the repertoires, for all frequency weightings. The only difference between the repertoires was in the pairwise similarity among sequences.

For the first repertoire in this illustrative example (Fig. 2d, top row), we chose closely related sequences from a single real-world CDR3 clone. We expected that species richness with similarity—functional species richness—would be close to 1. (We used 0D_s here; sD_s performed similarly.) We observed a value of 1.5; the extra 0.5 reflected sequence diversity within the clone. For the second repertoire (Fig. 2d, second row), we swapped out half the unique CDR3s with CDR3s from a different, unrelated real-world clone. As expected, we observed a rough

146 doubling of functional diversity, to 2.4. For the third repertoire (Fig. 2d, third row), we replaced
 147 all the sequences with 34 randomly chosen real-world CDR3s. We expected a functional
 148 diversity that was much higher than in the first two repertoires but less than 34 because of the
 149 inherent sequence similarities that make a CDR3 a CDR3, and, consistent with this expectation,
 150 observed a value of 22. For the final repertoire (Fig 2d, bottom row), we replaced the CDR3s
 151 with random amino-acid sequences (controlling for length), expecting a functional similarity of
 152 nearly 34, and this again was observed ($^0D_s=32$). In contrast to these differences in functional
 153 diversity, raw diversity was indistinguishably 34 for all four repertoires. In every example,
 154 functional diversity fit an intuitive sense of what diversity should mean (Fig. 2c), while raw
 155 diversity failed to detect the defining differences among these repertoires. The agreement
 156 between expectation and observation supports the validity of our functional-diversity framework
 157 for immune repertoires.

158 **Robustness.** Sampling error—the “missing-species” problem²⁴—is known to be a major
 159 potential confounder when measuring raw diversity, necessitating large (e.g.) blood volumes
 160 and/or post-hoc statistical correction for measurements on the sample to reflect repertoire
 161 diversity in the individual as a whole¹¹. A practical feature of functional similarity is that the
 162 values are smaller than those of raw diversity (reflecting clustering of similar sequences; Fig. 1-
 163 2). The effective coverage is therefore greater, meaning that less information about the
 164 functional diversity of an individual’s overall repertoire is lost upon sampling than is the case for
 165 raw diversity. This observation suggested that functional diversity is more robust to sampling
 166 error, possibly even making it accurate enough to use without statistical correction, and thus
 167 useful for the sample sizes typically available for sequencing (10,000-1 million cells).

168 To test this possibility, we systematically downsampled from a representative TCR β repertoire
 169 and two representative IgH (IgG) repertoires, one prepared from mRNA and one from genomic
 170 DNA, each with $\sim 10^6$ unique sequences, and compared raw vs. functional diversity on the
 171 subsamples to those of the full sample (Fig. 3). (We tested mRNA and DNA separately to test
 172 the possibility of lower diversity from mRNA than DNA, which was expected since
 173 transcriptionally less-active cells may be less likely to be sampled.) For TCR β , we found that
 174 functional species richness saturated at a sample size of $\sim 30,000$ sequences and functional
 175 entropy at $\sim 10,000$ sequences (Fig. 3b, first and third columns). Functional diversity for higher
 176 frequency-parameter values (q) saturated with even fewer sequences. For IgH from mRNA,
 177 functional species richness did not saturate but did plateau, with a final increase of ≤ 2 percent
 178 per order of magnitude. Assuming that each unique sequence corresponds to a cell and 10^{11} B

179 cells in the body, this final measured rate of increase means that the individual's total functional
180 species richness is no more than 50 percent higher than the value measured on the sample
181 (Fig. 3, middle row). This is the maximum expected sampling error. For IgH from DNA,
182 functional species richness had begun to plateau at the full sample size, resulting in the value
183 for the individual being no more than three times as much as in the sample (maximum three-fold
184 error). Meanwhile, functional entropy saturated at 30,000 cells for IgH from mRNA and 300,000
185 cells for IgH from DNA. This behavior was in marked contrast to that of raw diversity, which did
186 not saturate or plateau for species richness (Fig. 3, white symbols), consistent with previous
187 reports and illustrating the need for statistical correction¹¹.

188 We then asked whether the robustness of functional diversity can be expected to generalize for
189 any IgH or TCR repertoire. We reasoned that a “meta-repertoire” comprising sequences drawn
190 uniformly (i.e., without regard to frequency) from many individuals will be more diverse, by any
191 measure, than any single repertoire (which will have fewer sequences, and in which the same
192 sequence may appear multiple times, resulting in lower diversity for $q \geq 0$). Downsampling from a
193 meta-repertoire therefore provides an upper bound or worst-case scenario for the sampling
194 requirements for any single repertoire. To build meta-repertoires, we pooled and then uniqued
195 CDR3s from 114 different IgH repertoires from 79 individuals including Americans of African,
196 European, and Hispanic descent^{5,15,25} to build an IgH meta-repertoire of roughly 36 million
197 unique sequences—as many as or more than ever observed or currently estimated to be in a
198 typical individual—and similarly for CDR3s from TCR β repertoires from 69 healthy individuals
199 (of mostly European but some Asian descent)²⁶ to build a TCR β meta-repertoire of 10 million
200 unique sequences, and downsampled from each of these meta-repertoires as above (Fig. 3,
201 large circles). We chose healthy individuals to avoid any down-biasing of diversity that might
202 come with sampling many related sequences e.g. from expanded somatically hypermutated
203 clones (for IgH). We found that functional diversity plateaued for all q , saturated for $q \geq 1$ and
204 reflected overall diversity to within a few percent from sample sizes of 50,000 for TCR β and IgH
205 RNA and 100,000 for IgH DNA for $q=0$, and 30,000 for TCR β and IgH RNA and 300,000 for IgH
206 DNA for $q \geq 1$ (Fig. 3, large colored circles). Together, these results confirmed that functional
207 diversity measured on samples is an accurate measure of overall functional diversity in the
208 individual, at conventional sample sizes.

209 **Raw and functional diversity.** We measured raw and functional diversity on 141 healthy
210 human subjects (Fig. 4). For IgH, we found a (geometric) mean functional species richness (${}^{\circ}D_s$)
211 of 677 (range, 487-916) from mRNA and 2,205 (range, 2,042-2,485) from DNA, suggesting that

on average, the human antibody repertoire is capable of recognizing the equivalent of no more than a few thousand unique non-overlapping heavy-chain CDR3 binding targets. (As above, lower diversity from mRNA was not unexpected, since inactive cells, which produce less IgH mRNA than active cells, may be underrepresented.) For TCR β the mean functional diversity was 140 targets (range, 115-167). Functional diversity can be thought of as clustering similar sequences together, although functional clusters can overlap and sequences can belong to multiple clusters. An indication of the average size of these clusters can be obtained by taking the ratio of raw to functional diversity measures. For species richness, we found that IgH typically had on average hundreds of sequences per cluster, while TCR β had thousands. Thus by both functional diversity and average functional-cluster size, IgH CDR3 repertoires are roughly 5-10 times as diverse as TCR β (for small q). Repertoires with higher raw diversity might be expected to be more functionally diverse, but we found no consistent trend across all repertoire types (measured by correlation coefficient). Thus, functional diversity generally complements raw diversity, adding information not captured by raw diversity alone.

Naïve vs. memory B cells. We next sought to investigate what the complementary information provided by functional diversity might add to our understanding of adaptive immunity. We began by investigating two widely studied B-cell subsets, naïve (IgM) and memory cells (predominantly IgG). Previous studies have shown that naïve repertoires have higher raw diversity than memory repertoires⁵. This is at least superficially consistent with the fact that only a subset of naïve cells are selected to enter into the memory compartment. However, in a functional sense there is a case to be made that memory/IgG repertoires should be considered more diverse, since somatic hypermutation differentiates and thereby distances memory cells from naïve cells, and indeed from each other. Using well-characterized publicly available repertoires from DNA from three healthy human subjects, we confirmed that by raw species richness, naïve (CD27⁻ IgM⁺) B-cell repertoires are ~10 times as diverse as memory (CD27⁺ IgM⁻) repertoires⁷ (Fig. 5a-b). Yet by functional species richness, we found that memory repertoires were at least as diverse as naïve (Figs. 5a-b). Comparing raw and functional diversity for 34 IgM and 32 IgG repertoires from mRNA (repertoires with less than 100,000 total sequences were discarded) from 28 additional healthy individuals from a separate dataset showed a similar pattern as for the three DNA repertoires: in all but a few outliers, IgM had higher raw diversity but IgG had higher functional diversity (Fig. 5c-d). For raw diversity, the IgM:IgG ratio rose from ~3:1 at $q=0$ to peak at 10:1 around $q=1$, due to a large fraction of rare IgG sequences (Fig. 5d). This effect was more pronounced for the naïve:memory ratio in the DNA dataset (Fig. 5b). For functional

245 diversity, the absence of a peak in the IgM:IgG ratio suggests that these many rare sequences
246 must nonetheless be similar to others in the repertoire, possibly because they are clonally
247 related (Fig. 5b,d).

248 Cytomegalovirus (CMV) exposure. CMV is a herpesvirus to which half of the adult population
249 has been exposed and results in life-threatening opportunistic infections in newborns, transplant
250 recipients, and immunocompromised individuals²⁷. In most healthy individuals, it causes a
251 chronic infection marked by clonal expansion of both B cells and T cells and a consequent fall in
252 raw diversity, an effect also seen during aging (see below)^{4,28}. We measured raw and functional
253 TCR β CDR3 diversity for 120 individuals: 69 CMV-seronegative and 51 CMV-seropositive
254 subjects aged 19-35²⁶, the narrow age range helping control for any age-related effects. There
255 was a clear trend toward lower diversity in the CMV-seropositive group relative to the CMV-
256 seronegative group by both raw and functional diversity, for all weighting parameters (Fig. 6a).
257 However, combining raw with functional diversity facilitated identification of two subgroups
258 among the subjects with known CMV status: subjects with a high raw Berger-Parker Index (${}^{\infty}D$)
259 were almost always CMV-seronegative (Fig. 6b), whereas subjects with low functional Berger-
260 Parker Index (${}^{\infty}D_s$) were almost always CMV-seropositive (Fig. 6c). The reverse—low ${}^{\infty}D$ or high
261 ${}^{\infty}D_s$ —did not distinguish between the two groups. Using both measures gave a better indication
262 of CMV status than did either one alone (Fig. 6d). The conclusion is that CMV is unlikely in the
263 absence of large clones/expanded lineages, as has been reported, but is likely only if the large
264 clones/expanded lineages that are present exhibit high similarity to other clones/lineages in the
265 repertoire, or else are indeed very large (Fig. 6e). Again, the addition of functional diversity
266 offers insight that raw diversity alone does not.

267 Flu vaccination. Vaccination with a seasonal trivalent influenza vaccine (TIV) triggers clonal
268 expansion in B cells. Previous work on five vaccinees showed likely flu-specific memory IgG
269 lineages emerging by day 7 post-vaccination¹⁵. We found that combining raw and functional
270 diversity reveals a signature of clonal expansion and selection without the need for lineage
271 tracking (Fig. 7). We measured raw and functional diversity for IgM and IgG at day 0 (pre-
272 administration) and day 7 from all 14 vaccinees in Vollmers' dataset. We found that for most
273 subjects, for IgG, raw species richness rose from day 0 to day 7 while functional species
274 richness fell (Fig. 7a-b). This means that even as the number of sequences increased, many of
275 the new sequences were similar to each other (or to existing sequences), and they tended to
276 replace different-looking sequences. Meanwhile, there was no obvious pattern in IgM (Fig. 7c).
277 Together, these results are what we would expect from clonal expansion and selection in a

memory response, and thus represent a repertoire-level signature of these phenomena. Interestingly, in most cases, raw and functional entropy both fell (Fig. 7a-b, right panels). This suggests that most of the new sequences at day 7 were rare, while at the same time a subset of sequences and functional clusters grew. Thus overall, the addition of functional diversity reveals a key feature of clonal dynamics, which is not evident from raw diversity alone.

Aging. To explore the effect of age, we measured raw and functional diversity for TCR β CDR3 repertoires from 41 healthy individuals aged 6-90 years old²⁹ (Fig. 8). We found that raw diversity falls with age regardless of frequency parameter; a fall in only raw species richness had been reported previously²⁹. Functional diversity also fell, regardless of frequency parameter. However, for species richness, four septuagenarians bucked the trend (Fig. 8, arrows): even as their raw species richness was unremarkable relative to that of other individuals of similar age, their functional species richness was similar to that of children. Only one of these four had a high likelihood of being CMV-negative; the probability that all four were CMV-negative was low. We therefore consider CMV unlikely as an explanation for their high functional species richness. Unlike their peers, these four individuals appear to have retained functional diversity among their rarest sequences. (An alternative hypothesis is that these four saw a rise in functional species richness from a lower level earlier in their adulthood, but we consider this unlikely given the overall downward trend across individuals.) We considered but excluded PCR/sequencing artifacts as the cause, as we expected such artifacts would have led to larger raw species richness, which was not observed. Thus, functional diversity can identify for further study individuals who are unremarkable by raw diversity alone.

Discussion

Diversity both affects function and reflects it. In the adaptive immune system, the defining tradeoff is breadth vs. depth: a repertoire must be sufficiently diverse to contain sequences that can recognize a given target and lead to useful clones, but not so diverse that cells that express such sequences are too rare to encounter their target on biologically meaningful timescales^{30,31}. To monitor immunological diversity, either diagnostically or therapeutically, it must be measured, and to measure it, must be defined. The literature increasingly recognizes that any reasonable definition of immunological diversity must account for differences in species frequency. Here we add that such a definition must also account for species' pairwise similarity, and present a continuous framework showing that binding similarity, which leads to what we call functional diversity, provides useful insights into repertoire function.

310 Generically, pairwise similarity S_{ij} can be seen as governed by a tunable parameter that helps
311 define the similarity matrix S , analogous to how q is a tunable parameter that governs the effect
312 of differences in frequency³². In our study, the similarity matrix is defined by the average single
313 amino-acid change in K_d , an average based on over 1,300 independent measurements and
314 multiplicative independence. Our source data is not systematic, but to our knowledge is the best
315 available. While the present study is to our knowledge the first of its kind, it follows a long
316 tradition of attempts to estimate the number of binding targets that can be recognized by the
317 adaptive immune system. In these previous studies, typically a sample of B or T cells was
318 diluted until binding to/protection from a given target was abolished, using whatever thresholds
319 the various investigators deemed appropriate. If the limiting frequency for binding/protection was
320 found to be, for example, 1:3,000, the conclusion was that the repertoire could recognize 3,000
321 different targets. This conclusion was based on the assumption that on average, all targets
322 behave the same as the one under study. Such studies gave a wide range of functional
323 diversities: 100 to 100,000 targets for various T-cell populations and for B cells after antigen
324 exposure^{33,34}. This range may reflect real differences in the frequencies of cells that are specific
325 for different targets, variability in stringency or experimental setup, or some combination.
326 Interestingly, this “how-many-can-fit” logic seems not to have been employed when testing so-
327 called natural antibodies, which bind many targets at low affinity^{35–37}. For example, one in five
328 natural antibodies bind insulin³⁸, but this is not taken to mean that the repertoire could recognize
329 only five targets, because of presumed overlapping specificity of these antibodies for other
330 targets. Meanwhile, theoretical studies have suggested a need for $\leq 10,000$ binding targets, and
331 fewer for T cells than B cells, because of major histocompatibility complex (MHC) restriction^{31,39}.

332 In the present study, for raw diversity, we found that a typical repertoire contains on the order of
333 10 million unique CDR3s, well above the upper end of the range for previous estimates of the
334 number of binding targets. The present findings are in line with other recent estimates that were
335 likewise based on a combination of deep sequencing and statistical correction^{5,11,29}. On
336 average, the observed raw species richnesses mean that each of 100-100,000 putative binding
337 targets can be bound by 100-100,000 unique CDR3s ($10^7/10^5$ - $10^7/10^2$). From a medical
338 perspective, such redundancy is good for treatment, because it supports the prevailing view that
339 there are many ways to design an antibody- or TCR-based drug that will recognize a given
340 target, but potentially a complicating factor for attempts to diagnose specific diseases based on
341 repertoire sequence, because it suggests that signatures of exposure to a given target may
342 often (though not necessarily always) be quite variable.

343 A key new finding in the present work is that functional diversity is much lower than raw
344 diversity: repertoires contain only a few hundred functional clusters for TCR β CDR3s and at
345 most a few thousand for IgH. The fact that functional diversity is based on K_d s suggests that
346 functional diversity should correlate with the number of structurally unique, non-overlapping
347 target clusters that CDR3s can recognize (Fig. 1). Yet the present measurements of functional
348 species richness lie at the low end of the range of past estimates. We propose two not-mutually-
349 exclusive explanations. First, our measurements are limited to CDR3s; variability in the rest of
350 the antibody or TCR protein likely add to the total number of potential binding targets. This
351 possibility is testable by extending our method to more or indeed all of the antibody or TCR
352 sequence. Second, functional diversity may be providing a less detailed description of shape
353 space than limiting-dilution studies: i.e., functional diversity may be a coarse graining of the
354 target-binding landscape^{40,41} (Fig. 9). A pair of sequences may be similar enough to lie near
355 each other in shape space, but only one may bind a given target above a threshold level of
356 specificity in a binding study. In short, binding studies may be counting peaks while functional
357 diversity of CDR3s counts mountains. If true, our results suggest that the landscape of TCR β
358 CDR3 binding is more clustered than that of IgH, such that there are on average several times
359 as many functional IgH clusters as TCR β . This prediction is testable through large-scale
360 systematic binding assays to measure K_d , or by measuring binding as a binary outcome at
361 multiple stringency thresholds⁴². To our knowledge this is the first attempt a quantitative
362 summary of this landscape using data from large-scale binding studies and high-throughput
363 repertoire sequencing.

364 Why is CDR3 functional diversity higher for IgH than for TCR β ? We hypothesize that it is for the
365 same three reasons that there is more sequence diversity for IgH than TCR β . First, humans
366 have 23 D_H gene segments vs. only 2 D _{β} segments (a 10-fold difference), and D is the largest
367 germline contributor to CDR3. V and J segments tend to directly contribute little more than the
368 canonical starts and ends of CDR3s, and there are roughly similar numbers of V-J combinations
369 in IgH as TCR β (49 $\times$ 6=294 and 48 $\times$ 13=624, respectively; i.e. a twofold difference in the
370 opposite direction, or an aggregate five-fold difference in favor of IgH). Second is somatic
371 hypermutation, which diversifies IgH but not TCR β . And third, IgH CDR3s are longer than TCR β
372 CDR3s, allowing for a larger number of possible sequences. Further analysis will be needed to
373 test these possible contributions.

374 We have shown how functional diversity complements raw diversity to offer insight into the
375 difference between naïve and memory repertoires, to aid in identification of disease states, and

376 to illustrate clonal selection and other repertoire dynamics. We hope these examples will
377 encourage others to use and/or expand our framework to investigate repertoire dynamics in
378 other conditions, in other subsets, in the other chains (TCR α and IgL), and in other model
379 systems such as zebrafish²³ and mouse^{3,14}. We draw attention to the fascinating difference
380 between the number of unique sequences, which ran into the millions in most of the repertoires
381 we investigated, and the much smaller numbers of what we call functional clusters (the effective
382 numbers of functional diversity). The result is a “functional degeneracy” among sequences that
383 are organized into functional clusters. Characterizing these clusters is an interesting topic for
384 future work.

385 Another key finding is the perspective that functional similarity offers on diversity across
386 individuals and on similarities and differences between people. To show that functional diversity
387 is robust to sampling, we generated “meta-repertoires” by pooling sequences from scores and in
388 some cases over a hundred people, including people from different ethnic backgrounds.
389 Surprisingly, and somewhat unexpectedly given the low sequence overlap among individuals,
390 the functional diversity of these meta-repertoires never exceeded the functional diversity of any
391 given repertoire by more than a few fold; moreover, the functional diversity trended toward
392 saturating in samples of just a million sequences (Fig. 3). Together, these findings indicate that
393 any two individuals share a majority of their functional clusters, in stark contrast to the
394 vanishingly small fraction of sequences they share. Further, these findings imply that the
395 functional diversity of the entire population is only a few hundred clusters for TCR β CDR3s and
396 a few thousand for IgH, and moreover that these clusters can be sampled exhaustively by
397 sequencing fewer than 20 individuals. It will be fascinating to test this finding by generalizing
398 beyond the ~100 individuals studied in this experiment, including with additional ethnically and
399 geographically diverse populations, to further examine our prediction that, contrary to
400 conventional wisdom, the functional limits of the adaptive immune system for the entire human
401 species are in a practical sense both finite and within reach.

402 Finally, the focus of this study was binding similarity, but we expect that the utility of the
403 diversity-with-similarity framework will extend to other facets of immunology (using other
404 similarity measures to study, e.g., somatic hypermutation) and to other fields, most readily
405 metagenomics, sociology, oncology, and cellular cartography^{43–47}. We hope this study will serve
406 as a template for the methodology and value of incorporating similarity into the study of these
407 and other complex systems.

408 **Acknowledgements**

409 This research was supported by grants from the National Institutes of Health (NIAID
410 K08AI11495801) and the American Heart Association (15GPSPG23830004), contracts
411 HHSN268201500003I, N01-HC-95159, N01-HC-95160, N01-HC-95161, N01-HC-95162, N01-
412 HC-95163, N01-HC-95164, N01-HC-95165, N01-HC-95166, N01-HC-95167, N01-HC-95168
413 and N01-HC-95169 from the National Heart, Lung, and Blood Institute, and by grants UL1-TR-
414 000040, UL1-TR-001079, and UL1-TR-001420 from NCATS, as well as support from the
415 Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by
416 National Science Foundation grant number ACI-1548562. The authors thank the investigators,
417 staff, and participants of the MESA study for their valuable contributions (a full list of
418 participating MESA investigators and institutions can be found at <http://www.mesa-nhlbi.org>), as
419 well as the Research Computing Group at the High Performance Computing Cluster at Harvard
420 Medical School. Finally, the authors would like to thank Drs. Mohammed Al-Quraishi and
421 Jeremy Gunawardena for helpful conversations and comments on the manuscript.

422 **Methods**

423 **High-throughput repertoires.** We obtained 391 quantitative high-throughput IgH and TCR β
424 repertoires from 202 human subjects. These included IgH from naïve and memory B cells from
425 DNA ($n=3$ individuals)⁵; TCR β chains from DNA from healthy subjects known to be serologically
426 negative for cytomegalovirus (CMV) ($n=69$ individuals)²⁶ and from healthy subjects whose CMV
427 serostatus was unknown ($n=41$ individuals)²⁹; pooled barcoded IgG and IgM heavy chains from
428 mRNA from healthy subjects before and seven days after administration of one of two influenza
429 vaccines ($n=28$ individuals)¹⁵; quantitative pooled TCR β chains from DNA for subjects who were
430 otherwise healthy but serologically CMV positive ($n=51$ individuals)²⁶; and IgH chains from DNA
431 for subjects enrolled in the Multi-Ethnic study of Atherosclerosis (MESA; $n=41$ individuals)²⁵.
432 CDR3 annotation was performed using our in-house pipeline as previously reported¹⁴ and
433 standard tools (e.g. IMGT). Details for obtaining these datasets are available from the primary
434 publications referenced above.

435 **Similarity measures.** A functional measure of similarity between polypeptides is how well they
436 bind the same target. We were interested in similarity as a function of the number of amino acid
437 substitutions (i.e., as a function of edit distance). The effect of substitutions on binding is
438 complex and depends on the position and identity of the specific amino acids involved; many

439 substitutions may have little or no effect, while a few may abolish binding entirely⁴⁸. When
440 comprehensive data are available, detailed statistical models can offer reasonable predictions of
441 the effect of specific amino-acid substitutions^{49–51}. However, this type of data does not yet exist
442 across entire antibody and TCR repertoires, and so simpler models are required. These models
443 are not expected to precisely predict the effects of specific substitutions, but should accurately
444 reflect the effects of substitutions when averaged over many pairs of proteins, such as the
445 millions of pairs in megacell-scale repertoires.

446 To develop a model for our similarity measure, s , we downloaded SKEMPI 2.0, which is to our
447 knowledge the largest and best-curated database of experimentally measured effects of amino-
448 acid substitution on protein-protein binding²⁰. Each entry includes a Protein Data Bank (PDB)
449 identifier⁵², the type of structural region⁵³ that contains the substitution(s), one or more PDB
450 coordinates, and (in nearly all cases) the dissociation constant (K_d) of each member of the pair
451 (referred to in the database and Fig. 2a as “wild type” and “mutant”). We extracted entries for all
452 single amino-acid substitutions for which K_d for both wild type and mutant were recorded, and
453 considered only entries that involved binding between antibody and antigen ($n=797$) or TCR and
454 peptide/MHC ($n=531$; total $n=1,328$). Although amino-acid substitutions anywhere in a protein
455 may affect binding, substitutions at the core of the binding interface are more likely to affect
456 binding than substitutions elsewhere⁵³. Therefore we split the data into core ($n=584$) and non-
457 core ($n=744$) groups and analyzed the effect of substitution binding, measured as
458 $|\log_{10}(K_{dmut}/K_{dwt})|$, separately for each group.

459 As expected, the probability distributions for the two groups differed substantially from each
460 other (Mann-Whitney U p -value 2.0×10^{-33}). Substitution of a core residue had a 13-fold
461 (geometric) mean effect on binding, consistent with prior reports⁵⁴, while substitution of a non-
462 core residue had a 4-fold effect. Both probability distributions were long-tailed, and were
463 reasonably well described by exponential probability-density functions (i.e., of the form ke^{-kx}).
464 We confirmed that the distributions for antibody-antigen core residues ($n=352$) and TCR-
465 peptide/MHC core residues ($n=232$) were similar to each other, that the distributions for
466 antibody-antigen non-core residues ($n=445$) and TCR-peptide/MHC non-core residues ($n=299$)
467 were also similar to each other, that within each of the antibody-antigen and TCR-peptide/MHC
468 subgroups the distributions for core and non-core residues were different, and that these results
469 held separately for human and non-human (nearly all of which were mouse) sequences (using
470 the Structural Antibody Database⁵⁵ and the Structural TCR Database⁵⁶ to assign species), all
471 using Mann-Whitney U and visualized as histograms. Using PyMol v2.2.0⁵⁷, we next manually

reviewed nine structures containing substitutions in human IgH or TCR β CDR3s (1BD2, 1OGA, 3BN9, 3QDJ, 3SE8, 3SE9, 4I77, 5C6T, 5E9D) and found that to a good approximation, a constant fraction 0.15 ± 0.05 of CDR3 amino acids consist of core residues, with no obvious difference between chain types. To estimate the effect of a single amino-acid substitution in a CDR3 in our datasets, we therefore combined core and non-core distributions with a weighting of 0.15:0.85.

The resulting distribution was again long-tailed, with most substitutions having small effects and a few having effects of many orders of magnitude (Fig. 2a). There were small spikes in the tail for substitutions with ≥ 60 -fold effects, i.e. $|\log_{10}(K_{\text{dmut}}/K_{\text{dwt}})| \geq 1.8$. A review of sources cited by SKEMPI suggested that these spikes likely reflect ascertainment bias: selective experimentation on amino acids with unusually strong effects (e.g. refs 58,59). To counteract such bias, we therefore built a high-confidence dataset using 1.8 as the cutoff. Ascertainment bias in two- and three-amino-acid substitutions is expected to follow the square and cube of the bias in single-amino-acid substitutions, respectively, precluding rigorous conclusions from being drawn from independence testing. However, comparison with those groups was broadly consistent with either multiplicative ($s=c^m$, where s =similarity; c =the cost of binding, i.e., $1/(\text{fold effect})$; m =edit distance) or additive ($s=c/m$ for $m \geq 1$) independence. Because additive effects result in higher pairwise similarities and therefore smaller repertoire diversities than multiplicative effects, the multiplicative-independence model is more conservative for studying the effects of similarity on diversity. We therefore chose the multiplicative model for further analysis.

To determine the similarity between two CDR3s with edit distance m , we sampled independently from the high-confidence dataset m times, and multiplied the costs together. We confirmed that on average, the results of this stochastic sampling were the same as deterministic calculation of $s=c^m$ with $c \approx 0.55$. We performed sensitivity analysis based on lower-confidence cutoffs (down to $c=0.48$) and alternative assumptions (up to $c=0.60$). This resulted in somewhat higher or lower diversity values, but qualitative patterns were robust to these perturbations.

Diversity measures. We calculated qD according to Eq. (1) as previously described^{9,11} and qD_s according to Leinster and Cobbold¹⁶ following Eqs. (2)-(3). We corrected qD for sampling error using Recon (github.com/ArnaoutLab/Recon; default settings) as previously described¹¹. We note Hill's framework⁹ (Eq. (1)) has inspired several methods for incorporating similarity into diversity measurements, each of which retains useful features of Hill's framework^{16,32,60,61}. Two

of the new frameworks were introduced with explicit discussion of how to decompose population-level diversity into within- and between-group components^{16,60}. Each of these has advantages and disadvantages over the other⁶². We chose Leinster and Cobbold's framework here because we found it easier to apply and interpret. For readability, we made two minor changes to the notation, from ${}^qD^Z$ to qD_s and from $\mathbf{Zp}_i$ to $\mathbf{S}_i$.

Use of this framework raised two issues that we addressed. First, its $q=0$ measure, 0D_s , depends on frequency albeit to a very small extent, unlike the Hill framework's $q=0$ measure, 0D , which is species richness (and is independent of frequency). Therefore, as a more direct comparison to species richness, we calculated 0D_s both with frequency information and without it (i.e., setting the frequencies of each of the n species to $1/n$). We refer to the latter as 0D_s ("D-null"). Second, it has been shown that this framework can result in unreasonably low diversity values when most of the off-diagonal entries of the similarity matrix are far from zero, resulting in an insensitivity to q ^{67,69}. We expected most of our off diagonals to be close to zero, since our similarity measures directly or indirectly involve exponential decay, which generates small values, but confirmed that most of the off diagonals in our similarity matrices were indeed close to zero by plotting histograms. Consequently, our measures were sensitive to q , as desired and expected (Figs. 5-8).

Robustness analyses. For robustness analyses, IgH and TCR β were analyzed separately. The upper-bound/worst-case scenario for IgH was evaluated by constructing a "meta-repertoire" by combining IgG sequences of subjects before vaccination ($n=28$ individuals)¹⁵, sequences from memory cells from healthy subjects from public database ($n=3$ individuals)⁵, and sequences from subjects enrolled in MESA study ($n=41$ individuals)²⁵, and sampling from this meta-repertoire without regard to the frequency of sequences. We chose IgG/memory sequences where possible because those sets exhibited higher functional diversity than naïve sets, and we were interested in maximizing diversity. We ignored the frequency of sequences for the same reason: uniform frequency maximizes diversity, other things equal. For TCR β , we constructed a meta-repertoire by combining sequences from CMV seronegative individuals ($n=69$ individuals)²⁶ and again sampling at uniform frequency. We chose CMV seronegative individuals for the same reason as we chose memory/IgG sequences above: seronegative individuals exhibited higher diversity. For both IgH and TCR β , including all sequences lowered diversities slightly. The representative samples were from subject D3 for IgH (from DNA), subject SRR960344 for IgH (from mRNA), and subject Keck0070 for TCR β (CMV seronegative). CDR3 sequences were sampled proportional to their frequency in the repertoire.

537 Figure titles and legends

538 **Figure 1: Functional diversity.** (a) Each dot represents a binding target (e.g. an epitope) with
 539 a different shape. Nearby targets have similar shapes (inset). Targets form clusters of similarity.
 540 (b) Each colored region represents the targets that can be bound by one of six unique CDR3
 541 sequences in a representative repertoire; this repertoire has a raw species richness of 6.
 542 Together, the colored regions cover the part of shape space that can be bound by the
 543 repertoire. (The unbound region might include, e.g., self-antigens.) Note the substantial overlap
 544 in binding targets for the orange, yellow, green, and blue CDR3s. This overlap reflects binding
 545 similarity among these CDR3s. (c) Because of this similarity, the repertoire covers only the
 546 region denoted by the four identical non-overlapping squares. The functional species richness of
 547 this repertoire is therefore 4: this repertoire has the same species richness as a repertoire
 548 comprising four CDR3s that have zero overlap in binding specificity. (d) Schematic
 549 representation of repertoires without incorporating frequency or similarity (top), with frequency
 550 only (middle), and with frequency and similarity (bottom).

551 **Figure 2: Validity.** (a) Single amino-acid mutations in antibody and TCR molecules have a
 552 range of effects on affinity, as measured by change in dissociation constant, K_d (gray). This was
 553 well fit by a simple exponential (black line), providing parameterization for the similarity metric.
 554 (b) CDR3s with high sequence identity have high similarity, while different CDR3s have low
 555 similarity. Shown are two clones, represented by red and white subnetworks, each composed of
 556 17 unique CDR3 sequences drawn from clonotypes of two real IgH repertoires. Node size
 557 corresponds to the frequency of each sequence; edges connect pairs of sequences that differ at
 558 a single amino-acid position. (c)-(d) Functional similarity agrees with an intuitive sense of
 559 repertoire diversity. Each of the four repertoires in (c) has the same number of unique
 560 sequences, present at the same frequencies; as a result, they all have identical raw diversity
 561 (for every value of q) despite their obvious quantitative and qualitative differences. In contrast,
 562 functional diversity increases with the number of, and increasing difference between,
 563 repertoires' constituent sequences. Node size denotes sequence frequency. Shades denote
 564 different clonotypes. Comparing the third and fourth rows, note that even when two repertoires
 565 have the same number and frequencies of unique sequences, the repertoire whose sequences
 566 are more different from each other (random peptides) has the higher functional diversity.

567 **Figure 3: Robustness.** Results for raw and functional species richness ($q=0$; 0D and 0D_s) and
 568 raw and functional entropy ($q=1$; 1D and 1D_s). Raw, white shapes; functional, colored shapes.

569 Large symbols give an upper bound/worst-case scenario based on sampling meta-repertoires;
 570 small symbols give results for a representative sample from DNA (circles) and mRNA
 571 (triangles). First row: sample diversity is plotted as the effective number of sequences. For $q=0$,
 572 functional diversity plateaus for TCR β and IgH RNA and trends toward a plateau for IgH DNA at
 573 the tested sample sizes; all three plateau for $q=1$; raw diversity does not plateau for either $q=0$
 574 or $q=1$. Second row: discovery rate is the probability that the next sampled sequence will add to
 575 the diversity. For example, for the IgH DNA representative sample, for $q=0$ raw diversity, at a
 576 sample size of 1 million sequences there is still a probability of ~ 0.5 (a 50 percent chance) that
 577 the next sequence to be sampled will be one that has not yet been seen and will therefore add
 578 to the diversity. Third row: maximum error is the maximum fraction by which the diversity in the
 579 sample can underestimate the diversity in the individual from whom the sample was taken.
 580 Horizontal dashed lines indicate the threshold for two-fold error. For example, for the worst-case
 581 scenario for TCR β , $q=0$ functional diversity measured on a sample of 10,000 sequences will be
 582 no more than a two-fold underestimate of diversity in the individual as a whole; in other words,
 583 the sample value will be at least 50 percent of the overall value.

584 **Figure 4: Diversity in individuals.** Raw (black lines; left vertical axis) and functional (colored
 585 bars; right vertical axis) species richness ($q=0$) for 179 CDR3 repertoires from healthy
 586 individuals representing (a) IgH from mRNA (all isotypes: IgA, IgG, IgM, IgD, and IgE), (b) IgM
 587 and (c) IgG from mRNA from the subjects in (a), (d) IgH from DNA (all isotypes), (e) naïve IgH
 588 from DNA, (f) memory IgH from DNA, and (g) TCR β from DNA. See Methods for references.
 589 Matched pairs of symbols below the horizontal axis denote replicates. Note the difference in the
 590 scale for functional diversity between IgH and TCR β . Note also a general lack of correlation
 591 between raw and functional species richness, except in (c).

592 **Figure 5: Naïve vs. memory.** (a) Diversity profiles for naïve (red) and memory (black)
 593 compartments from three deeply sequenced individuals. A diversity profile is a way to show
 594 diversity across a range frequency-weighting parameter values at once. By raw diversity (left),
 595 the naïve compartment is more diverse across the range of weightings. By functional diversity
 596 (right), this distinction disappears. In (b), this disappearance is highlighted by plotting the ratio of
 597 naïve:memory diversity for raw diversity (red) and functional diversity (black). According to
 598 functional diversity, the naïve compartment is no more diverse, and indeed sometimes
 599 somewhat less diverse, than the memory compartment. This reversal is even more prominent in
 600 comparisons of repertoires from an additional 28 healthy subjects (c,d).

Figure 6: Infection. (a) Diversity profiles showing effective number of species as a function of weighting parameter q for diversity without similarity (left) and diversity with similarity (right) showing a trend toward lower diversity in CMV-seropositive individuals (red) relative to CMV-seronegative individuals (black), especially for large q . (b) Raw Berger-Parker Index ($q=\infty$), which measures the largest clones, showing that high diversity—an absence of large clones—is rare in CMV-seropositive individuals. (c) Functional Berger-Parker Index, showing that low diversity—the presence of large clones with similarity to other clones in the repertoire—is rare in CMV-seronegative individuals. (d) Combining raw and functional Berger-Parker Indices (first principal component of PCA, which explains 72 percent of variance) illustrates both of the trends in (c): for the third of subjects beyond the cutoffs indicated by the horizontal dashed lines CMV serological status is assigned with an accuracy of 95 percent. (e) Schematic representation of the three classes revealed by combining diversity with and without similarity. Each circle is a clone; each collection of clones is a representative repertoire. Top: subjects without large clones are almost always CMV seronegative. Bottom: subjects with large clones that are similar to other clones in the sample (shown in red) are almost always CMV seropositive. Middle: repertoires with large clones that are not similar to other clones in the repertoire may be either CMV seropositive or CMV seronegative. Receiver-operator characteristic (ROC) analysis gave an area under the curve (AUC) of 0.79.

Figure 7: Vaccination. Raw and functional diversity together reveal clonal expansion and selection without needing lineage analysis. (a) In the IgG compartment, raw species richness rises while functional species richness falls in most vaccinees (left). Meanwhile raw and functional entropy both fall (right). The difference vs. species richness suggests most new sequences at day 7 are rare. (b) Meanwhile, the IgM compartment changes less by these measures.

Figure 8: Aging. Raw and functional species richness ($q=\infty$) for TCR β CDR3 repertoires from 41 healthy individuals. Arrows denote four septuagenarians who bucked the trend of lower functional species richness with age. Note that for each individual, the raw species richness is ~10-fold higher than previously reported (Britanova 2014), likely because the method we used to correct for missing species (Recon) is more sensitive than the method used in the previous report (Fisher).

Figure 9: Binding landscape. (a) Schematic of the target-binding landscape. The gray distribution represents CDR3 sequences that bind a given target. Sequences are ordered by

633 their similarity to each other. (In reality, similarity is a multidimensional property that makes it
634 impossible to order sequences in a single dimension as shown here; this is done for illustrative
635 purposes only.) The height at each sequence denotes the affinity with which it binds the target
636 (vertical axis; measured e.g. by K_d). Many more sequences bind the target at low affinity than at
637 high affinity, resulting in a “mountain-and-peak” appearance. This schematic is useful for
638 interpreting functional diversity as described in this study, and the raw diversity estimates based
639 on previous binding studies, as described in the Discussion. (Note that in this schematic, the
640 raw diversity as measured simply corresponds to the total number of sequences along, i.e. the
641 width of, the horizontal axis.) At high affinity, very few sequences bind a given target. At medium
642 affinity, more sequences bind, and can be binned into two small clusters, represented by the
643 small circles. At low affinity, many sequences bind, and can be binned into a single large
644 cluster, represented by the large circle. In **(b)-(d)**, many targets are shown. Each color
645 corresponds to a different target; nearby targets are structurally similar. As in (a), each colored
646 area denotes the sequences that bind a given target, as a function of binding affinity (vertical
647 axis). Experiments usually detect the highest-affinity sequences: the peaks of the landscape
648 (above the horizontal dotted line). The narrower the peak when it crosses the experimental
649 threshold, the rarer specific sequences are, and the larger the number of targets that the
650 repertoire will be estimated to bind. (For example, if 100,000 sequences are shown across the
651 horizontal axis in each plot, and only one crosses the experimental threshold for a given target,
652 the frequency of sequences specific for that target is 1:100,000, and the conclusion will be that
653 there must be 100,000 such targets that the repertoire can bind. If 100 cross the experimental
654 threshold, the conclusion will be that the repertoire can bind only 1,000 targets.) Functional
655 diversity measures the overall contours of the landscape. Conceptually, this can be thought of
656 as measuring the number and size of the “mountains” at a lower affinity threshold (horizontal
657 solid lines). The differences in functional diversity between (b) memory IgH, (c) TCR β , and (d)
658 naïve IgH correspond to different landscapes. The raw species richness of memory IgH and
659 TCR β are comparable, represented here by the same width of all the plots. In addition, a similar
660 number of sequences per target cross the experimental threshold, so estimates of the total
661 number of targets that the repertoires can bind will also be comparable. However, less low-
662 affinity overlap between the targets of the IgH sequences in (b) gives it higher functional species
663 richness than the TCR β repertoire in (c): here, six functional clusters (white circles) vs. three.
664 (The sizes of the clusters are related to frequency-weighted functional diversity measures, i.e.
665 larger q .) The sequences in the naïve IgH repertoire in (d) have only low affinity for the six
666 colored targets, and many recognize more than one target (overlapping colored areas). Note the

667 lower experimental threshold (horizontal dotted line), consistent with the ~10% or more of
 668 antibodies that recognize a target and the high degree of cross-reactivity in studies of natural
 669 antibodies^{39,40}. The functional diversity threshold controlled by the average effect on K_d of a
 670 single-amino-acid change in CDR3. If the effect were larger—or if it were amplified by e.g.
 671 raising it to a power when building the similarity matrix—the threshold would be higher, and vice
 672 versa.

673 Tables and Figures

Box 1: Glossary

Diversity Any of several measures of the number of species present in a population, possibly accounting for the relative frequencies of, or pairwise similarity between, species (e.g. species richness, Shannon entropy, Gini coefficient, Berger-Parker index)

Frequency The relative number of each species in the population under study, which is taken into account by most measures of diversity but not species richness (which is a simple count)

Similarity An investigator-defined measure of how alike or different a given pair of species are to each other (e.g. the Hamming distance between each pair of sequences)

Effective numbers A method for comparing different diversity measures in the same units, i.e. numbers of species. Consider two repertoires of 100 clones each. In the first repertoire, one clone is large and accounts for 91 percent of all cells (e.g. a leukemic clone); the other 99 clones are small and account for the remaining 9 percent. In the second repertoire, all 100 clones are equally common, each accounting for 1 percent of cells. The Shannon entropies of the two repertoires are 1.0 bit and 6.6 bits. Entropy is converted to an effective number— 1D —by exponentiation: the effective number of clones in the first repertoire is $2^{1.0}=2$, while in the second repertoire it is $2^{6.6}=100$. Thus per entropy, the first repertoire can be thought of as “effectively” consisting of just two clones: the 99 rare clones collectively count the same as the one large clone. In other words, the first repertoire has the same effective diversity as a repertoire that consists of just two clones that are equally common. The second repertoire already consists of clones that are equally common, so the effective number of clones in this repertoire, $2^{6.6}=100$, is the same as its species richness. Diversity with similarity is interpreted in a similar fashion: a repertoire with a qD_s of n species has the same effective diversity as a repertoire with n species that are equally common (as above), with the additional constraint that these species are now also completely unrelated to/dissimilar from each other.

674

Figure 1

675

Figure 2

678

Figure 3

679

Figure 4

Figure 5

684

Figure 6

685

686

Figure 7

687

688

Figure 8

689

690

Figure 9

691

References

1. Jiang, N. *et al.* Lineage structure of the human antibody repertoire in response to influenza vaccination. *Sci Transl Med* **5**, 171ra19 (2013).
2. Robins, H. S. *et al.* Overlap and Effective Size of the Human CD8+ T Cell Receptor Repertoire. *Science Translational Medicine* **2**, 47ra64-47ra64 (2010).
3. Arnaout, R. *et al.* High-Resolution Description of Antibody Heavy-Chain Repertoires in Humans. *PLOS ONE* **6**, e22365 (2011).
4. de Bourcy, C. F. A. *et al.* Phylogenetic analysis of the human antibody repertoire reveals quantitative signatures of immune senescence and aging. *Proc Natl Acad Sci U S A* **114**, 1105–1110 (2017).
5. DeWitt, W. S. *et al.* A Public Database of Memory and Naive B-Cell Receptor Sequences. *PLoS ONE* **11**, e0160853 (2016).
6. Messaoudi, I., Lemaoult, J., Guevara-Patino, J. A., Metzner, B. M. & Nikolich-Zugich, J. Age-related CD8 T cell clonal expansions constrict CD8 T cell repertoire and have the potential to impair immune defense. *J. Exp. Med.* **200**, 1347–1358 (2004).
7. Gibson, K. L. *et al.* B-cell diversity decreases in old age and is correlated with poor health status. *Aging Cell* **8**, 18–25 (2009).
8. Hopkins, A. C. *et al.* T cell receptor repertoire features associated with survival in immunotherapy-treated pancreatic ductal adenocarcinoma. *JCI Insight* **3**, (2018).
9. Hill, M. O. Diversity and Evenness: A Unifying Notation and Its Consequences. *Ecology* **54**, 427–432 (1973).

10. Morris, E. K. *et al.* Choosing and using diversity indices: insights for ecological applications from the German Biodiversity Exploratories. *Ecol Evol* **4**, 3514–3524 (2014).
11. Kaplinsky, J. & Arnaout, R. Robust estimates of overall immune-repertoire diversity from high-throughput measurements on samples. *Nat Commun* **7**, 11881 (2016).
12. Greiff, V. *et al.* A bioinformatic framework for immune repertoire diversity profiling enables detection of immunological status. *Genome Med* **7**, 49 (2015).
13. Ju, C.-H. *et al.* Plasmablast antibody repertoires in elderly influenza vaccine responders exhibit restricted diversity but increased breadth of binding across influenza strains. *Clin. Immunol.* **193**, 70–79 (2018).
14. Kaplinsky, J. *et al.* Antibody repertoire deep sequencing reveals antigen-independent selection in maturing B cells. *Proc. Natl. Acad. Sci. U.S.A.* **111**, E2622-2629 (2014).
15. Vollmers, C., Sit, R. V., Weinstein, J. A., Dekker, C. L. & Quake, S. R. Genetic measurement of memory B-cell recall using antibody repertoire sequencing. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 13463–13468 (2013).
16. Leinster, T. & Cobbold, C. A. Measuring diversity: the importance of species similarity. *Ecology* **93**, 477–489 (2012).
17. Macarthur, R. H. Patterns of Species Diversity. *Biological Reviews* **40**, 510–533 (1965).
18. Jost, L. Partitioning diversity into independent alpha and beta components. *Ecology* **88**, 2427–2439 (2007).
19. Marion, Z. H., Fordyce, J. A. & Fitzpatrick, B. M. Extending the Concept of Diversity Partitioning to Characterize Phenotypic Complexity. *Am. Nat.* **186**, 348–361 (2015).

20. Jankauskaite, J., Jiménez-García, B., Dapkunas, J., Fernández-Recio, J. & Moal, I. H. SKEMPI 2.0: An updated benchmark of changes in protein-protein binding energy, kinetics and thermodynamics upon mutation. *Bioinformatics* (2018).
doi:10.1093/bioinformatics/bty635
21. Xu, J. L. & Davis, M. M. Diversity in the CDR3 region of V(H) is sufficient for most antibody specificities. *Immunity* **13**, 37–45 (2000).
22. Perelson, A. S. & Oster, G. F. Theoretical studies of clonal selection: minimal antibody repertoire size and reliability of self-non-self discrimination. *J. Theor. Biol.* **81**, 645–670 (1979).
23. Weinstein, J. A., Jiang, N., White, R. A., Fisher, D. S. & Quake, S. R. High-throughput sequencing of the zebrafish antibody repertoire. *Science* **324**, 807–810 (2009).
24. Bunge, J. & Fitzpatrick, M. Estimating the Number of Species: A Review. *Journal of the American Statistical Association* **88**, 364–373 (1993).
25. Bild, D. E. *et al.* Multi-Ethnic Study of Atherosclerosis: objectives and design. *Am. J. Epidemiol.* **156**, 871–881 (2002).
26. Emerson, R. O. *et al.* Immunosequencing identifies signatures of cytomegalovirus exposure history and HLA-mediated effects on the T cell repertoire. *Nat. Genet.* **49**, 659–665 (2017).
27. Emery, V. Investigation of CMV disease in immunocompromised patients. *J Clin Pathol* **54**, 84–88 (2001).
28. Qi, Q. *et al.* Diversity and clonal selection in the human T-cell repertoire. *Proc Natl Acad Sci U S A* **111**, 13139–13144 (2014).

29. Britanova, O. V. *et al.* Age-related decrease in TCR repertoire diversity measured with deep and normalized sequence profiling. *J. Immunol.* **192**, 2689–2698 (2014).
30. Schober, K., Buchholz, V. R. & Busch, D. H. TCR repertoire evolution during maintenance of CMV-specific T-cell populations. *Immunol. Rev.* **283**, 113–128 (2018).
31. Zarnitsyna, V. I., Evavold, B. D., Schoettle, L. N., Blattman, J. N. & Antia, R. Estimating the diversity, completeness, and cross-reactivity of the T cell repertoire. *Front Immunol* **4**, 485 (2013).
32. Chao, A. *et al.* An attribute-diversity approach to functional diversity, functional beta diversity, and related (dis)similarity measures. *Ecological Monographs* (2018).
doi:10.1002/ecm.1343
33. Bachmann, M. F., Kündig, T. M., Kalberer, C. P., Hengartner, H. & Zinkernagel, R. M. How many specific B cells are needed to protect against a virus? *J. Immunol.* **152**, 4235–4241 (1994).
34. Obar, J. J., Khanna, K. M. & Lefrançois, L. Endogenous naive CD8+ T cell precursor frequency regulates primary and memory responses to infection. *Immunity* **28**, 859–869 (2008).
35. Frank, S. A. *Immunology and Evolution of Infectious Disease*. (Princeton University Press, 2002).
36. Holodick, N. E., Rodríguez-Zhurbenko, N. & Hernández, A. M. Defining Natural Antibodies. *Front Immunol* **8**, (2017).
37. Notkins, A. L. Polyreactivity of antibody molecules. *Trends Immunol.* **25**, 174–179 (2004).

38. Chen, Z. J. *et al.* Polyreactive antigen-binding B cells are the predominant cell type in the newborn B cell repertoire. *Eur. J. Immunol.* **28**, 989–994 (1998).
39. Langman, R. E. & Cohn, M. The E-T (elephant-tadpole) paradox necessitates the concept of a unit of B-cell function: the protection. *Mol. Immunol.* **24**, 675–697 (1987).
40. Fairlie-Clarke, K. J., Shuker, D. M. & Graham, A. L. Why do adaptive immune responses cross-react? *Evol Appl* **2**, 122–131 (2009).
41. Smith, D. J., Forrest, S., Hightower, R. R. & Perelson, A. S. Deriving Shape Space Parameters from Immunological Data. *Journal of Theoretical Biology* **189**, 141–150 (1997).
42. Mora, T. & Walczak, A. Quantifying lymphocyte receptor diversity. *arXiv:1604.00487 [q-bio]* (2016).
43. Almendro, V. *et al.* Genetic and phenotypic diversity in breast tumor metastases. *Cancer Res* **74**, 1338–1348 (2014).
44. Heindl, A., Lan, C., Rodrigues, D. N., Koelble, K. & Yuan, Y. Similarity and diversity of the tumor microenvironment in multiple metastases: critical implications for overall and progression-free survival of high-grade serous ovarian cancer. *Oncotarget* **7**, 71123–71135 (2016).
45. Koopmans, R. & Schaeffer, M. *De-composing diversity: In-group size and out-group entropy and their relationship to neighbourhood cohesion*. (WZB Berlin Social Science Center, 2013).
46. Li, K., Bihan, M., Yooseph, S. & Methé, B. A. Analyses of the Microbial Diversity across the Human Microbiome. *PLOS ONE* **7**, e32118 (2012).

47. Taraska, J. W. Cell biology of the future: Nanometer-scale cellular cartography. *J Cell Biol* **211**, 211–214 (2015).
48. Lunzer, M., Golding, G. B. & Dean, A. M. Pervasive cryptic epistasis in molecular evolution. *PLoS Genet.* **6**, e1001162 (2010).
49. Hopf, T. A. *et al.* Mutation effects predicted from sequence co-variation. *Nat. Biotechnol.* **35**, 128–135 (2017).
50. Salinas, V. H. & Ranganathan, R. Coevolution-based inference of amino acid interactions underlying protein function. *Elife* **7**, (2018).
51. Lee, J. *et al.* Surface Sites for Engineering Allosteric Control in Proteins. *Science* **322**, 438–442 (2008).
52. Berman, H. M. *et al.* The Protein Data Bank. *Nucleic Acids Res.* **28**, 235–242 (2000).
53. Levy, E. D. A simple definition of structural regions in proteins and its use in analyzing interface evolution. *J. Mol. Biol.* **403**, 660–670 (2010).
54. Whittaker, J. *et al.* Alanine scanning mutagenesis of a type 1 insulin-like growth factor receptor ligand binding site. *J. Biol. Chem.* **276**, 43980–43986 (2001).
55. Dunbar, J. *et al.* SAbDab: the structural antibody database. *Nucleic Acids Res.* **42**, D1140–1146 (2014).
56. Leem, J., de Oliveira, S. H. P., Krawczyk, K. & Deane, C. M. STCRDab: the structural T-cell receptor database. *Nucleic Acids Res.* **46**, D406–D412 (2018).
57. PyMOL The PyMOL Molecular Graphics System, Version 2.0, Schrödinger, LLC.

58. Pons, J., Rajpal, A. & Kirsch, J. F. Energetic analysis of an antigen/antibody interface: alanine scanning mutagenesis and double mutant cycles on the HyHEL-10/lysozyme interaction. *Protein Sci.* **8**, 958–968 (1999).
59. Taylor, M. G., Rajpal, A. & Kirsch, J. F. Kinetic epitope mapping of the chicken lysozyme.HyHEL-10 Fab complex: delineation of docking trajectories. *Protein Sci.* **7**, 1857–1867 (1998).
60. Chiu, C.-H. & Chao, A. Distance-Based Functional Diversity Measures and Their Decomposition: A Framework Based on Hill Numbers. *PLOS ONE* **9**, e100014 (2014).
61. Scheiner, S. M. A metric of biodiversity that integrates abundance, phylogeny, and function. *Oikos* **121**, 1191–1202 (2012).
62. Botta-Dukát, Z. The generalized replication principle and the partitioning of functional diversity into independent alpha and beta components. *Ecography* **41**, 40–50 (2018).