

RESEARCH ARTICLE

Learning the statistics and landscape of somatic mutation-induced insertions and deletions in antibodies

Cosimo Lupo¹, Natanael Spisak¹, Aleksandra M. Walczak^{1*}, Thierry Mora^{1*}

Laboratoire de physique de l'École normale supérieure, CNRS, PSL University, Sorbonne Université, and Université de Paris, Paris, France

* These authors contributed equally to this work.

* Current address: Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Roma I, Rome, Italy

* awalczak@phys.ens.fr (AMW); tmora@phys.ens.fr (TM)

OPEN ACCESS

Citation: Lupo C, Spisak N, Walczak AM, Mora T (2022) Learning the statistics and landscape of somatic mutation-induced insertions and deletions in antibodies. *PLoS Comput Biol* 18(6): e1010167. <https://doi.org/10.1371/journal.pcbi.1010167>

Editor: Andrew J. Yates, Columbia University Medical Center, Columbia University Irving Medical Center, UNITED STATES

Received: December 15, 2021

Accepted: May 5, 2022

Published: June 2, 2022

Copyright: © 2022 Lupo et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All the data analyzed in this paper has been previously published and can be accessed from original publications. The software here presented is freely available at <https://github.com/statbiophys/indie>.

Funding: This study was supported by the European Research Council COG 724208 (to AMW), <https://erc.europa.eu/>. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Abstract

Affinity maturation is crucial for improving the binding affinity of antibodies to antigens. This process is mainly driven by point substitutions caused by somatic hypermutations of the immunoglobulin gene. It also includes deletions and insertions of genomic material known as indels. While the landscape of point substitutions has been extensively studied, a detailed statistical description of indels is still lacking. Here we present a probabilistic inference tool to learn the statistics of indels from repertoire sequencing data, which overcomes the pitfalls and biases of standard annotation methods. The model includes antibody-specific maturation ages to account for variable mutational loads in the repertoire. After validation on synthetic data, we applied our tool to a large dataset of human immunoglobulin heavy chains. The inferred model allows us to identify universal statistical features of indels in heavy chains. We report distinct insertion and deletion hotspots, and show that the distribution of lengths of indels follows a geometric distribution, which puts constraints on future mechanistic models of the hypermutation process.

Author summary

Affinity maturation of B cell receptors is an important mechanism by which our body designs neutralizing antibodies to defend us against pathogens, including broadly neutralizing antibodies, which target a wide range of variants of the same pathogen. Such antibodies often contain key insertions and deletions to the germline gene, or “indels”, which are caused by somatic hypermutations. However, the mechanism, frequency and role of these indels are still elusive. We designed a computational method based on a probabilistic framework to infer the characteristics of this mutational process from high-throughput antibody sequencing experiments. Applied to human data, our approach provides a comprehensive quantitative description of insertions and deletions, opening avenues for better understanding the process of affinity maturation and the design of vaccines for eliciting a broad antibody response.

Competing interests: The authors have declared that no competing interests exist.

Introduction

The extraordinary ability of the immune system to detect and neutralize pathogens is partly assured by B cells, with the production of huge amounts of diverse immunoglobulins (Ig). Their initial diversity, provided by stochastic V(D)J recombination [1–8], is further increased by affinity maturation, a mutation-selection Darwinian process that takes place primarily in germinal centers (as well as extrafollicularly prior to germinal center formation [9]), where cells with the highest affinity to the pathogen compete for survival and proliferation [10–12].

During affinity maturation, cells diversify their immunoglobulin receptors through somatic hypermutations (SHM), which are initiated by the activation-induced cytidine deaminase (AID) enzyme [13] targeting immunoglobulin-coding genes. Its main effect is to cause point substitutions with a much higher rate than classical somatic mutations, namely $\sim 10^{-3}$ substitutions per base-pair per cell division [14, 15]. Since affinity maturation implies many proliferation and selection rounds, this process can lead to Ig with up to 30–40% of amino-acid substitutions with respect to the initial V(D)J rearrangement. Statistics of SHM have been extensively studied [5, 7, 16–22], mainly thanks to high-throughput repertoire sequencing [2, 3, 6, 8], although the link to the molecular details of AID functioning remains elusive.

But point mutations are not the only modifications found in matured Ig receptors. Deletions and insertions of one or multiple base pairs, collectively referred to as *indels*, have been detected in the V and J segments, albeit at a much lower rate than substitutions. Indels have been reported both in both heavy and light Ig chains [23–32], with a frequency that grows with the substitution rate during affinity maturation [23, 24, 28, 29]. Indels occur non-uniformly along the V germline templates, with a preference towards complementary-determining regions (CDRs) with respect to framework regions (FWRs); these ‘hotspots’ often coincide with those for point mutations [24, 32–34].

Indels are known to be critical for the increase of antibody repertoire diversity, enhancing the immune response in the presence of viral and bacterial pathogens as e. g. HIV-1 and influenza [35–39] and in response to vaccination [32]. Several works have shown a beneficial impact of indels for binding affinity and neutralization activity *in vitro* and *in vivo* [32, 40, 41], opening new evolutionary pathways beyond point mutations. A particularly striking example is given by anti-HIV-1 broadly-neutralizing antibodies (bnAbs), many of which have a high indel load [42–44]. In particular, bnAb CH31 was shown to owe its neutralization breadth to a particular long insertion in its lineage [39].

Yet, little is known about the biological mechanisms behind the occurrence of indels. Observations show that inserted segments have a high homology with their flanking regions (either on the 3′ or 5′ sides) [29, 31–33], possibly suggesting duplications rather than non-templated, random insertions. A proposed mechanism [45, 46] involves template regions with repeats and palindromic sequences, where the polymerase can “slip” during replication and “jump” to a nearby homologous element on the templated strand, producing a loop in either the daughter or the templated strand. If this unpaired loop is not properly repaired by the specific DNA repair mechanisms, according to its location on the daughter or the templated strand, it will be propagated as an insertion or a deletion, respectively. This mechanism would also explain co-localization of indels and SHM, since they could likely occur independently in the same susceptible regions (hot-spots due to repeats and palindromes) or also as a product of the repair mechanisms targeting existing replication errors.

Deep high-throughput sequencing of Ig repertoires now allows us to gain insight into the mechanisms of indels, and to collect statistics about their frequency, length distributions, and positional preference. By contrast, previous studies [33, 34] have relied on relatively small datasets and therefore small number of indel events due to their rarity. In addition, most studies

have focused on productive sequences, for which it is difficult to decouple intrinsic preferences from selection effects. Here, we exploit a recently published large immunoglobulin heavy chain (IgH) repertoire dataset comprising sequences from 9 healthy donors [8]. We developed an inference methodology to provide robust indel statistics and to capture their universal features in the face of natural individual variability.

Our approach uses a probabilistic framework to overcome the issue of annotation errors. Classically, indels are identified by looking at the best-scoring annotation. Such deterministic annotation is sometimes wrong, and we will show that this leads to systematic biases. Our probabilistic method, which considers multiple indel and point-mutation scenarios weighted by their probabilities, allows us to remove these biases, in a similar way to what was done for inferring the statistics of VDJ recombination [5, 7, 47]. A second key ingredient of our approach is to assume that each sequence may have a different mutation rate, since repertoire data combines B cells from various stages of affinity maturation. We first tested and validated our method on synthetically produced data. We then applied it to Ig sequences with a frameshift in their CDR3 region, which are believed to be nonproductive and thus free of selection biases. This allows us to characterize in great detail the statistics and intrinsic preferences of indel generation.

Results

Insertion and deletion events scale with the mutation rate

We conduct our analysis on a recently published high-throughput IgH repertoire dataset, obtained from the blood of 9 healthy donors [8]. IgM and IgG expressing cells were isolated and analyzed separately. Using raw sequence data, we further segregated IgH sequences into productive (P) and nonproductive (NP) sequences, depending on whether their CDR3 had a frameshift or not. Since hypermutation indels are rare (see below) compared to VDJ recombination frameshifts, we assumed that nonproductive sequences were already faulty at generation. Cells with NP sequences owe their survival to a successful IgH rearrangement on the second chromosome, meaning that the NP sequences themselves are free of selection effects. Most productive sequences undergo various selection processes, which bias the statistics of their features [5]. Here we mainly restrict our analysis to nonproductive sequences to capture the biochemistry of the hypermutation machinery. We obtained 421, 185 IgG and 459, 165 IgM nonproductive sequences from 9 donors.

To get preliminary insights into indel statistics, we first annotated each sequence deterministically using IgBLAST [48]. We estimated the rates per base pair of SHM point mutations, deletions and insertions, in both IgM and IgG compartments, separately for each individual and each of the productive and nonproductive subsets (Fig 1A). Mutation rates for productive sequences are consistent with previous estimates [33], with more mutations in IgG than in IgM. Some individual variability is present, suggesting variations in the degree of antigen exposure across individuals. Nonproductive sequences have higher rates of SHM than productive ones, especially for insertions and deletions (~10-fold difference for IgM and a bit less for IgG), suggesting that mutations are mainly deleterious.

Fig 1B shows the distribution of the number of indels in the IgG subset across donors. For comparison, a Poisson-distribution fit agrees well with the data, but underestimates sequences with many indels. Point mutations and indels both accumulate during affinity maturation and are mediated by AID [24, 28, 29, 32–34]. Indels are mostly present in highly point-mutated sequences. Both types of mutations have been reported to co-localize along the sequence, preferably in the CDRs [33, 34]. They have also been shown to co-localize temporally, as deduced from the analysis of antibody phylogenetic trees [32, 39]. Consistent with this picture, we

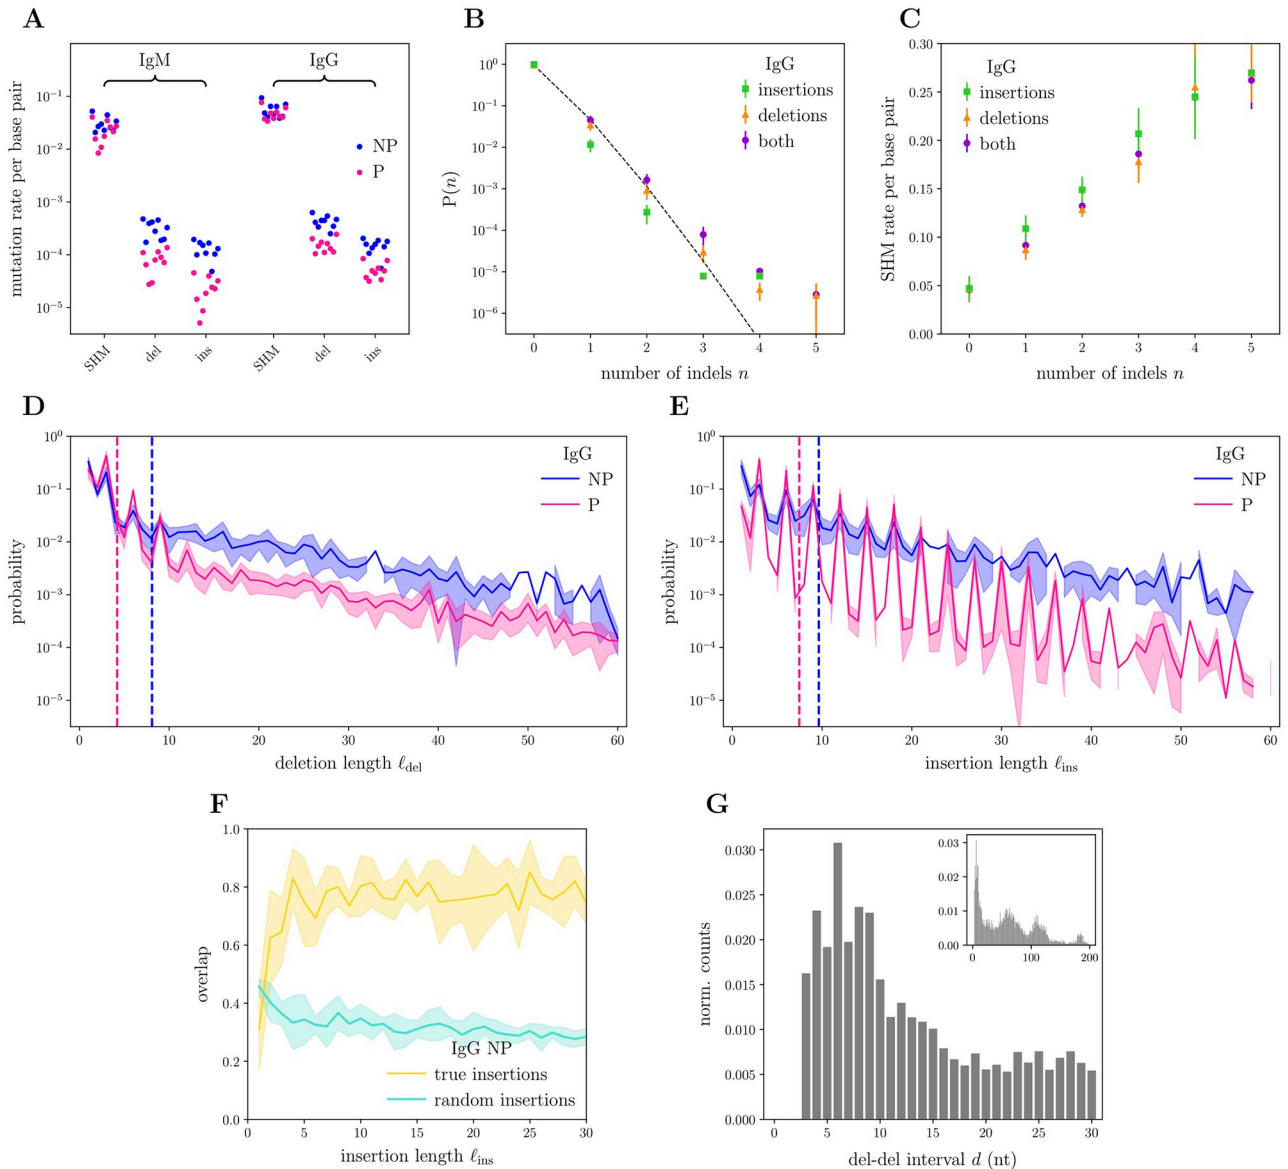


Fig 1. Indel statistics in real IgG sequences from deterministic annotation. (A) Somatic hyper-mutations (SHM), deletion (del) and insertion (ins) absolute rates per base pair for the 9 Briney donors [8]. Segregation into productive (P) and nonproductive (NP) sequences coded by colors; data points from IgM sequences are also reported, for comparison. (B) Fraction of IgG sequences (productive and nonproductive) with exactly n indels, either separated by type (del or ins) or cumulated; mean and standard deviation over the 9 donors represented by the markers and bars, respectively. The dashed black line represents the closest Poisson trend for cumulated indels of both types. (C) Fraction of SHM (in nucleotides) along IgG sequences (productive and nonproductive) conditioned on having exactly n indels, again either cumulated or separated by type; averages and one standard deviation bars over the 9 donors. (D) Length profiles for deletions, after NP-P segregation. Average and standard deviation over the 9 donors represented by the solid line and the shaded area around, respectively; corresponding average lengths are represented by the two vertical dashed lines. (E) Same as (D), but for insertions. (F) Overlap between inserted base pairs and same-length flanking regions on either 3' or 5' side along the sequence (larger is considered); average and one standard deviation over the 9 donors given by solid line and shaded area around, respectively. For comparison, overlap between random insertions and flanking regions is also reported (details in the main text). (G) Distribution of distances between deletions along the same sequence, measured as the number of base pairs in between, for IgG sequences (productive and nonproductive) hosting exactly two deletions and no insertions (data pooled together from the 9 donors); inset shows the full distribution, main plot focuses on the short-distance region.

<https://doi.org/10.1371/journal.pcbi.1010167.g001>

observe that the average point-mutation rate per base pair in IgG grows linearly with the number of indel events (Fig 1C).

Length and composition of inserted and deleted segments

Deletion and insertion length profiles, which show the distributions of the number of deleted or inserted base pairs, are shown in Fig 1D and 1E. Signatures of selection can be seen through an increased preference for multiples of 3 in productive *versus* nonproductive sequences, suggesting that indels that induce a frameshift are deleterious. The effect is only present at short deletion lengths, but holds for long insertion lengths. Note that a weaker 3-fold periodicity is also present in nonproductive sequences, despite them being in principle free of selection effects. This could come from sequences that were previously productive but were frameshifted relatively recently due to an indel in the CDR3. Long insertions are favored in productive sequences in comparison to long deletions, probably because of stability and functionality constraints [33]. By contrast, in nonproductive sequences deletions and insertions follow similar statistics, suggesting a common mechanism. Both insertions and deletions are approximately geometrically distributed, at least in the tail that describes large inserted or deleted segments. This geometric law puts strong constraints on the type of biophysical mechanisms by which indels may be created.

A proposed mechanism for indels is the *polymerase slippage model* [45, 46]. In germline regions that have repeats and palindromic sequences, the polymerase can “slip” during replication and “jump” to the closeby homologous element on the templated strand, resulting in an insertion or duplication depending on the location of the polymerase on the daughter or templated strand. In this scenario, insertions would be the result of duplications, and would thus be expected to be homologous to their flanking regions. The analysis of inserted segments in nonproductive sequence does reveal a higher than expected overlap with the best matching of the two immediately flanking sequences (Fig 1F). Previous similar observations were made on small numbers of productive sequences [33], for which effects of selection may confound the effect. Note that overlap is substantially different from 1 (≈ 0.8), suggesting 20% additional errors in the duplication, which is much higher than the mean point-substitution rate per base pair (around 5–6%). Also, we find a weak preference for the 5' end as a duplication source (S1 Fig). However, we found no significant overlap between deleted segments and their flanking regions (S2 Fig). This suggests either that the slip-and-jump mechanism does not favor homologous regions for deletions, or that other mechanisms are responsible for deletions, such as double-strand breaks [34, 49].

Indels co-localize with each other and with point mutations even in absence of selection

Since point mutations tend to co-localize in hotspot regions, and indels are associated with point mutations, we wondered if indels co-localized as well. Fig 1G shows the distribution of distances between two consecutive deletion events along the same sequence. To rule out possible confounding factors, we focused on sequences with exactly two deletion events and no insertions. Since such sequences are rare, we pooled the IgG repertoires (P and NP) of all individuals for this analysis. The full distribution (inset) is highly non-uniform, consistent with the observation that indels concentrate on CDRs [33, 34]. In addition, we observe a striking depletion of deletion pairs separated by less than 3 base pairs. This depletion is an artifact of the alignment software, which merges deletion events if they are too close to each other. This effect implies that the indels involving ≈ 50 –60 base pairs reported in Fig 1D and 1E may result

from such mergers, which standard alignment tools cannot disentangle. Our probabilistic framework will allow us to address this issue in the following sections.

We then studied the positional preference of indels along the V segment. To establish a universal profile across all V genes, we used the gapped absolute indexing provided by IMGT [50], and pooled the sequences of all 9 donors. The resulting profiles for nonproductive and productive sequences are shown in Fig 2A and 2B. Due to alignment gaps and finite read lengths, many IMGT positions are only observed in a fraction of sequences (black line). The indel profiles are highly non-uniform, with generally higher rates in the CDRs than in the Framework Regions (FWR), with the exception of a deletion hotspot the end of FWR3. We note that these trends, which are consistent with previous reports on productive sequences [33, 34], are found also in nonproductive sequences, suggesting that they cannot be attributed to selective or functional effects.

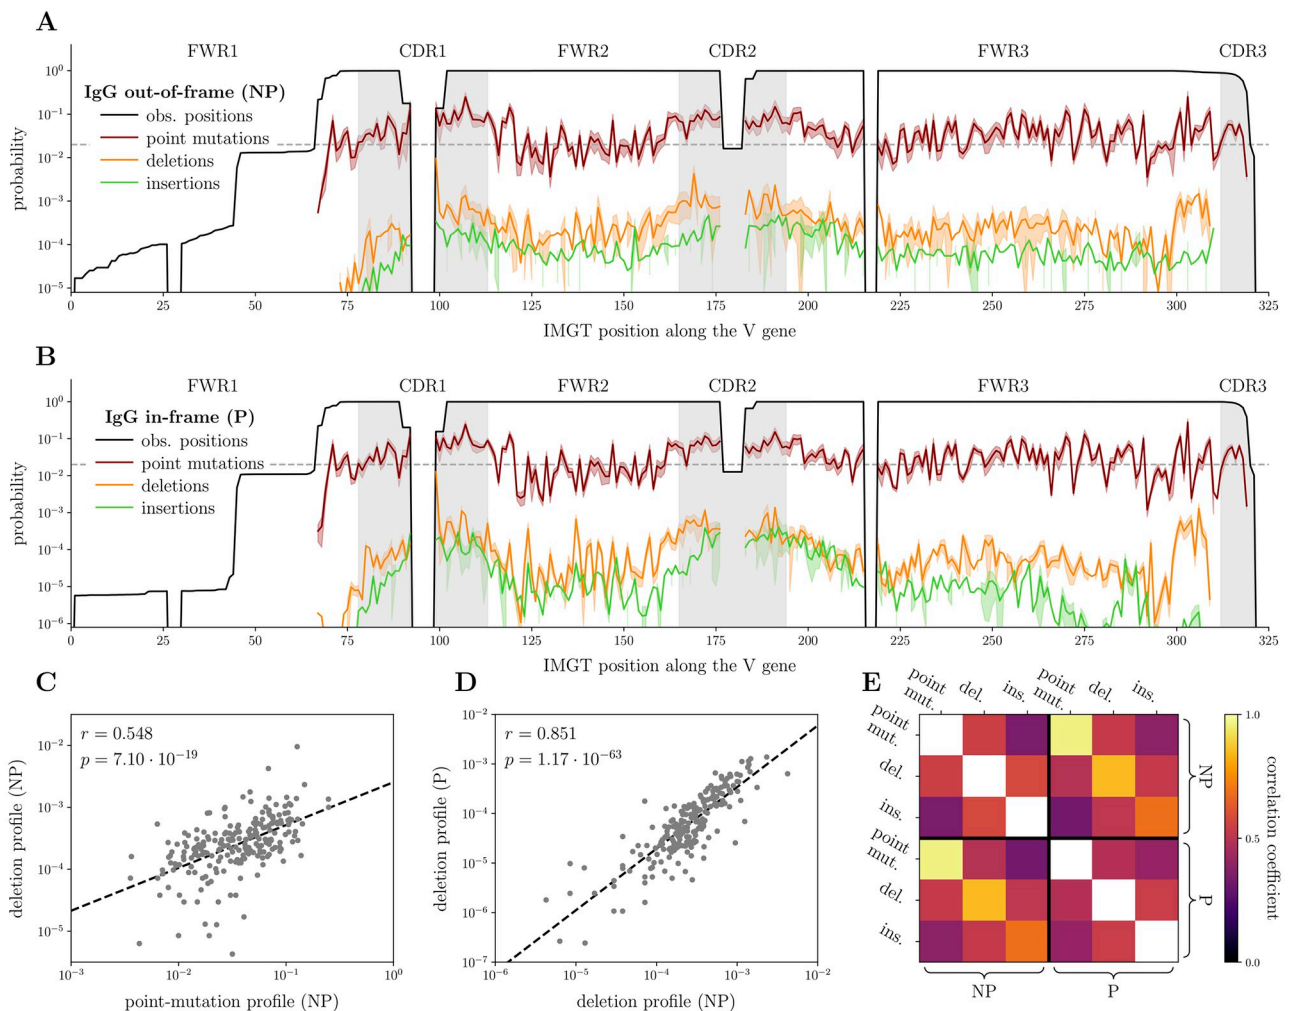


Fig 2. Indel profiles. (A) Mutation profiles for out-of-frame (NP) sequences, pooling together all the 9 Briney donors [8] and all the V templates (with shaded area representing individual variability), using the absolute IMGT indexing of positions [50]. Due to alignment gaps and finite read lengths, many IMGT positions are only observed in a fraction of sequences (black line). Mutation profiles are rescaled to account for this, and are masked for IMGT positions appearing in less than 2% of sequences (dashed horizontal line). Indel rates past IMGT position 310 are also not reported due to ambiguities with junctional diversity. (B) Same as (A), for in-frame (P) sequences. (C) Scatter plot between the SHM profile and the deletion profile for NP sequences; linear regression line is reported, together with numerical results (correlation coefficient r and p value). (D) Same as (C), here comparing deletion profiles for NP sequences versus P sequences. (E) Correlation between all the mutation profiles, computed either on NP or on P sequences.

<https://doi.org/10.1371/journal.pcbi.1010167.g002>

The different mutation profiles (point mutations, insertions, deletions, productive versus nonproductive) are correlated (Fig 2C–2E), meaning that indel and point-mutation hotspots are largely shared. Deletion profiles are well correlated with both point-mutation and insertion profiles. These correlations are larger than between insertions and point mutations, suggesting a possibly different mechanism for the two types of indels. In addition and consistently with this picture, deletions and point mutations show an enrichment of co-localization (within a few base pairs) relative to what would be expected from their shared positional preferences, while insertions and point mutations do not (S3 Fig). These observations are consistent with the leading model that deletions are caused by double-strand breaks [34, 49].

Profiles in productive versus nonproductive sets are very similar, suggesting that the positional effects of functional selection on indels and point-mutation preferences is weak, and that most of the previously reported biases in indel positions [33] are due to the positional preference of the involved enzymes rather than to selection.

A probabilistic alignment algorithm

While deterministic annotations give us a good picture of indel statistics, they may introduce systematic biases that could confound the analysis, as e.g. the spurious gap in the inter-deletion interval distribution (Fig 1G), or the over-estimation of long indel events. To give just an example of possible mistakes by deterministic annotation algorithms, a long indel observed in a sequence may have in fact been caused by several shorter indel events at the same site or close by. To evaluate the probability of the end product, one should consider all possible such decompositions into two or more indel events, and sum their probabilities. To address this issue, we designed a probabilistic approach to annotate and analyze hypermutated V segments, similarly to what was previously proposed to infer the statistics of V(D)J recombination from B- and T-cell receptor repertoires [5, 47] and implemented in the IGoR tool [7].

We first define a generative model of point mutations, insertions and deletions, with unknown parameters that describe the main statistics of the mutation process: distribution of SHM rates across sequences, mean ratios of deletion and insertion rates to the point mutation rate, and insertion and deletion length profiles. These parameters are then learned from the data using a maximum-likelihood estimator, which is implemented in practice through an expectation-maximization algorithm. In this framework, the likelihood of a particular outcome of the SHM process is obtained by summing the probabilities of all possible scenarios of germline V templates, point mutations, insertions, and deletions that would yield that sequence. In principle, this means that all alignments to germline V segments should be taken into account, and not just the best-scoring one as in traditional methods. However, many unlikely alignment scenarios can be pruned out of the list before computation (see Methods). In practice, this implies that only one plausible germline V gene is typically considered. Once the parameters of the model are learned, they can be analyzed in their own right, or the model may be used to generate synthetic datasets to test hypotheses. Fig 3A summarizes the general workflow of the method.

The generative model is defined so as to be simple enough to be tractable, while recapitulating the main features of the SHM process. Its parameters are directly interpretable in terms of basic SHM statistics. Each sequence has a distinct point-mutation rate μ per base pair, corresponding to its maturation age, drawn from a distribution $P(\mu)$. Motivated by the linear relation between indel and point-mutation rates (Fig 1C), we assume that the rates of insertion and deletion events are given by $\beta_{\text{ins}} \times \mu$ and $\beta_{\text{del}} \times \mu$, where β_{ins} and β_{del} are sequence-independent factors. Then insertions, deletions, and point mutations are drawn randomly with these three rates, uniformly along the sequence. The lengths of each insertion and deletion event are

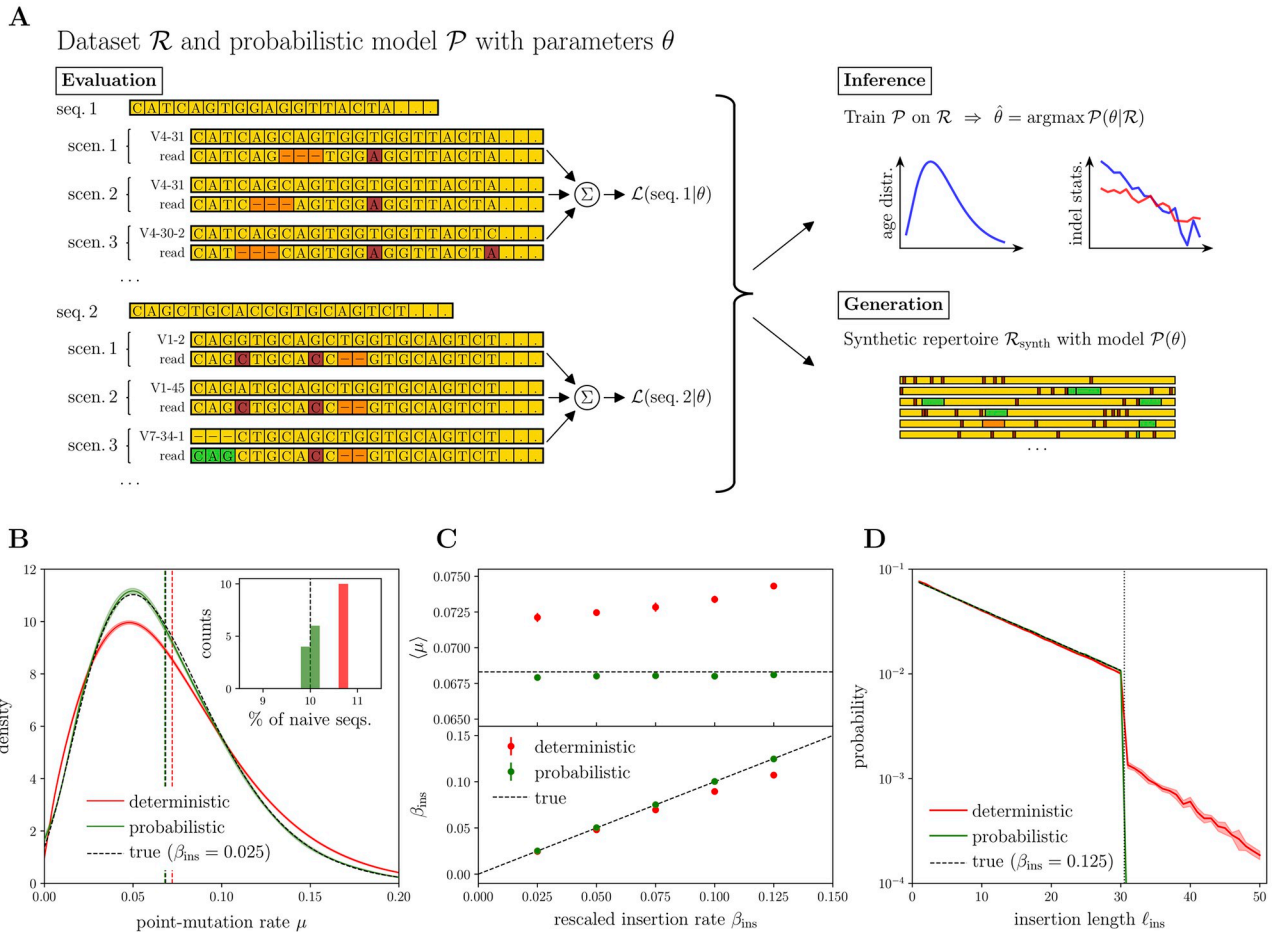


Fig 3. Workflow of the probabilistic framework and validation on synthetic data. (A) Given the set of parameters θ for the probabilistic model \mathcal{P} , for each sequence a set of alignment scenarios (given by template, read and SHM-indel pattern) is considered and the cumulated alignment likelihood $\mathcal{L}(\text{seq.}|\theta)$ from all of them is computed. Indel statistics and repertoire-wide age distribution from \mathcal{R} can be extracted by using a given model (evaluation mode) or by training a new model on \mathcal{R} itself (inference mode). Finally, a synthetic repertoire of arbitrary size can be generated, according to a given input model $\mathcal{P}(\theta)$ (generation mode). (B) Repertoire-wide point-mutation rate distribution $P(\mu)$; mean and standard deviation over $\mathcal{N} = 10$ independent realizations of a 100,000 synthetic sequences repertoire represented as a solid line with the shaded area around, for both deterministic and probabilistic estimates. In the inset, the \mathcal{N} estimations of the fraction of truly naive sequences; color coding is the same as the main plot. (C) Performance of the probabilistic alignment software vs deterministic one when increasing the density of indels; estimation of the average point-mutation rate in the upper panel, and of the ratio of insertion to point-mutation rates β_{ins} in the lower panel (the deletion rate β_{del} has a very similar behavior, shown in S4 Fig). Average values plus one standard deviation error bars are obtained over the \mathcal{N} independent synthetic repertoires. (D) Insertion length profiles for the largest value of β_{ins} considered in panel (C); mean and one standard deviation over the \mathcal{N} independent realizations shown. Deletion length profile is again completely analogous, shown in S4 Fig. The vertical dotted line signals the maximal size $\ell = 30$ used for generating insertions and deletions, then correctly identified by the probabilistic approach.

<https://doi.org/10.1371/journal.pcbi.1010167.g003>

drawn from two distributions $P_{\text{ins}}(\ell_{\text{ins}})$ and $P_{\text{del}}(\ell_{\text{del}})$. The set of parameters to be inferred is collectively called $\theta = (P, \beta_{\text{ins}}, \beta_{\text{del}}, P_{\text{ins}}, P_{\text{del}})$, and corresponds to the basic statistics of the SHM process. To calculate the sum of probabilities over all plausible scenarios efficiently, we use a forward algorithm that avoids computing the sum explicitly. Mathematical details of the probabilistic model and of the likelihood maximization are reported in the Methods section.

Validation of the probabilistic approach on synthetic data

The ability of our inference procedure to recover the true indel statistics can be tested by using large enough synthetic datasets. To this aim, we generated $\mathcal{N} = 10$ independent repertoires of

100,000 synthetic sequences each, restricted to the V segment, using realistic parameters similar to those obtained by the analysis of standard annotation (Fig 1). Naive sequences were first generated by drawing V templates according to their frequencies in the data. Then, point and indel mutations were added according to a generative SHM model belonging to the class described above. We assumed a mixture of 10% naive cells, defined as cells that have not gone through the hypermutation process at all, $\mu = 0$, and 90% experienced cells with μ distributed according to a shifted Gamma distribution mimicking a realistic distribution for point mutations (Gamma with mode $\mu = 0.07$, standard deviation = 0.04, and shift -0.02 ; see Methods for details). The lengths of insertions and deletions followed geometric distributions with characteristic scale $\ell = 15$ and a maximum value of 30 base pairs.

On each of these synthetic datasets, we compare two methods for estimating the parameters of the model. The first method is the one used in the previous sections of the paper: synthetic sequences are aligned to germline V segments using a deterministic tool; then, from these alignments, the statistics $P(\mu)$, β_{ins} , etc, are directly computed from their frequencies. The second method is the probabilistic inference approach outlined in the previous section.

In Fig 3B we report results for a realistic rescaled indel rate, $\beta_{\text{del}} = \beta_{\text{ins}} = 0.025$, of the same order of magnitude of the ones found for nonproductive IgG sequences through standard annotation, see Fig 1A. The fraction of naive cells ($\mu = 0$), shown in inset, and the distribution of μ for experienced cells, shown in the main figure, are correctly inferred only by the probabilistic approach. An example of the error made by the deterministic approach involves naive sequences: sequences with no SHM are always labeled as “naive” by the best-scoring alignment, while the probabilistic method considers the possibility that they come from an experienced cell ($\mu > 0$) in which no SHM has had the chance to occur (which happens with probability $e^{-L\mu}$, where L is the length of the V segment). This leads the deterministic method to systematically overestimate the true number of naive sequences. The performance of the deterministic method degrades as the rescaled insertion rate β_{ins} is increased (and β_{del} with it, Fig 3C). In that regime of frequent indels, best-scoring alignments merge nearby events, underestimating their relative frequency and compensating resulting errors by an increased mutation rate μ . By contrast, the probabilistic method recovers the ground truth even at very high indel densities, corresponding to 3 deletion and 3 insertion events per sequence on average. The merging of nearby indels also causes the deterministic algorithm to find an excess of non-existing long insertion events above 30 base pairs, while the probabilistic one correctly assigns their frequency to zero (Fig 3D). The analogous of panels 3C and 3D for deletions are reported in S4 Fig.

Overall, this analysis on synthetic data shows that while the deterministic analysis correctly captures the main statistics of SHM, it may fail on other quantities such as the fraction of truly naive cells, or the frequency of rare, long indels.

Inference within the probabilistic framework

We next applied our probabilistic inference method to the analysis of nonproductive IgH sequences from the 9 donors of [8] to learn the parameters of the SHM process free of selection effects. We thus obtain individual-specific distributions of maturation ages (Fig 4A, green). For comparison, the naive deterministic estimate, based on the frequency of errors in best-scoring alignments, is also reported (red). We observe substantial variability across individuals, probably due to diverse histories of experience with antigens. The probabilistic method finds more sequences with a small, non-zero maturation age ($94.3\% \pm 4.1\%$) than the deterministic one ($83.1\% \pm 8.6\%$), see Fig 4B. Because such sequences often have no mutations, the deterministic method underestimates that number (as confirmed by our synthetic data results, see Fig

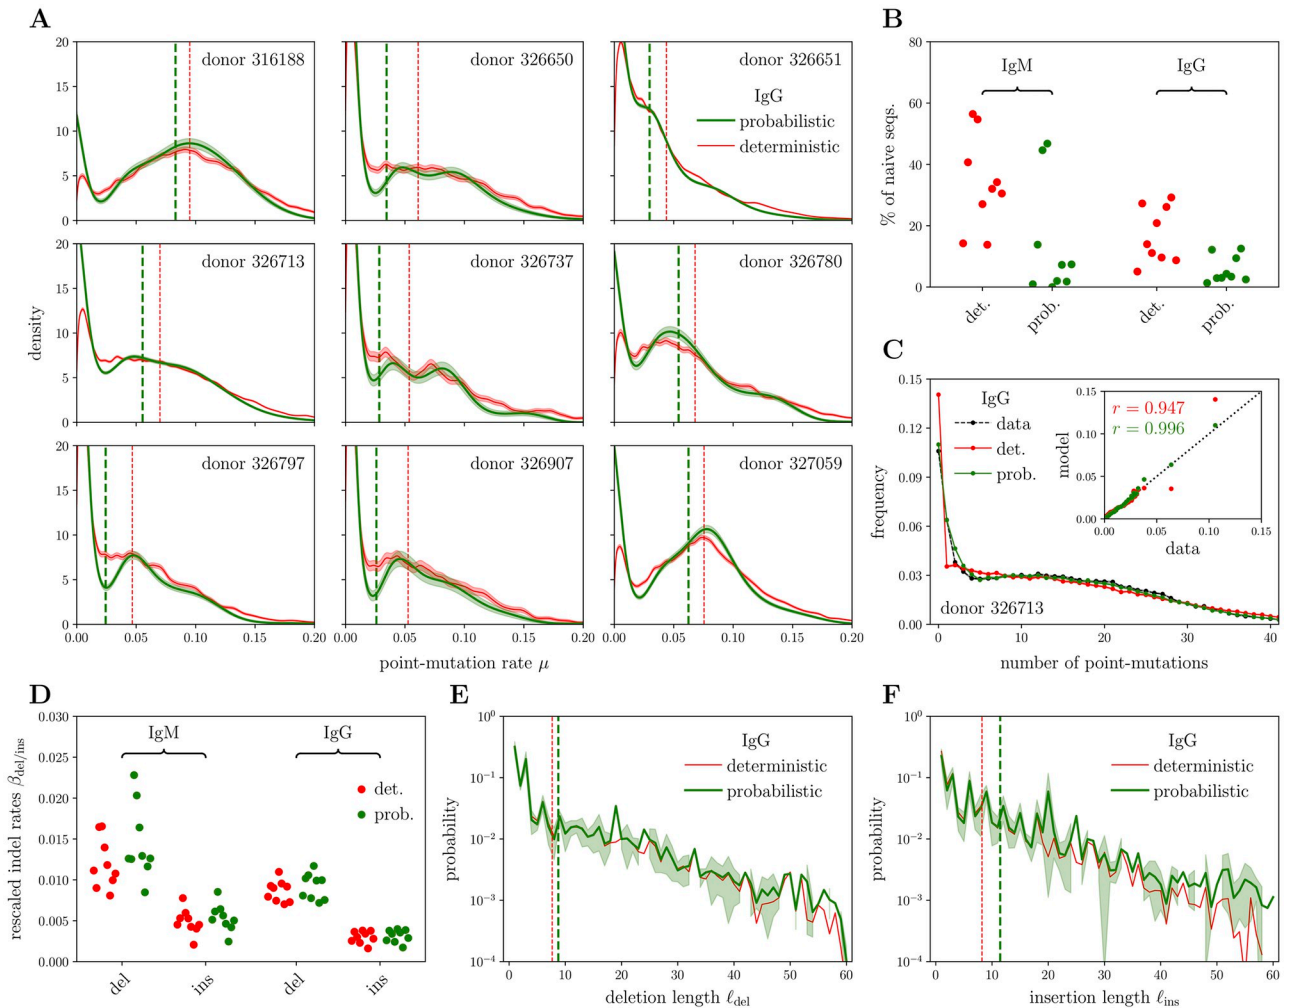


Fig 4. Results of probabilistic algorithm on IgG repertoire data. Probabilistic SHM-indel models trained on the 9 Briney healthy donors [8] (IgG sequences always shown, plus IgM only shown when relevant); for comparison, models based on deterministic assignments are also shown. (A) Point-mutation rate distributions; vertical lines locate average values for the distributions; deterministic distributions come from a smoothed version of SHM fraction histograms (more details in the main text); shaded areas show the confidence interval of the inference result (estimated from a Poisson noise assumption). (B) Corresponding estimates of the fraction of truly naive sequences, with a comparison also with the IgM case. (C) Validation of the inference results on one of the largest donors (326713); full histogram for the number of SHM in the main plot, corresponding scatter plot in the inset (correlation coefficients: $r_{\text{det.}} = 0.947$ vs $r_{\text{prob.}} = 0.996$). (D) Estimates for the rescaled (i.e. divided by the point-mutation rate) deletion and insertion rates per base pair; IgM is shown for comparison. (E) Inferred deletion length profile (solid line for the mean over the 9 donors, plus shaded area around for standard deviation), compared to the deterministic estimate (only mean shown). Vertical lines with the same color coding refer to the mean lengths. (F) Same as (E), but for insertions.

<https://doi.org/10.1371/journal.pcbi.1010167.g004>

3B), suggesting that the probabilistic approach gives a more accurate result. While we don't know the ground truth for the true distribution of maturation ages, we can verify this interpretation on the distribution of point mutations. In Fig 4C we compare predictions from the deterministic and probabilistic approaches for the IgG repertoire of one donor. Consistent with our intuition, the deterministic distribution overrepresents unmutated sequences, while the probabilistic distribution is in excellent agreement with the data. In addition, the model reproduces the super-Poissonian distribution of the number of indels (S5 Fig).

In Fig 4D we report estimates for rescaled indel rates, i.e. ratios of indel to point-mutation rates. They are consistent between the deterministic and probabilistic methods, as expected from the analysis on synthetic data, because of the relatively low indel rates in real data. They

are also well conserved across individuals. While the raw mutation rates depend on the history of affinity maturation, ratios between them are not expected to depend on it. They should instead be determined by universal biophysical mechanisms of DNA damage and repair, which explains their relative conservation. Similarly, the indel length profiles (Fig 4E and 4F) are consistent between the deterministic and probabilistic methods, as well as across individuals, pointing again to the universality of the biophysical mechanisms at play. However, the average deletion and insertion lengths are slightly larger in the probabilistic estimate, due to a more reliable evaluation of low-probability events corresponding to the tails of the distributions.

Discussion

Most studies of SHM insertions and deletions, and even point mutations, consider productive rearrangements. Instead, we focused our analysis on nonproductive sequences to study the statistics of mutation inherent to the biochemistry of the SHM process, free of subsequent selection effects. However, we found a strong correlation between the productive and nonproductive positional profiles of SHM, implying that most of the heterogeneity is of biochemical origin, rather than stemming from selection. Further studies of the differences in SHM statistics between productive and nonproductive sequences could shed light on the structural and functional constraints that selection imposes on antibodies during affinity maturation.

Because nonproductive sequences and indel SHM are both rare, the combination of both is extremely rare, necessitating a very large dataset such as the one provided by *Briney et al.* [8]. This study included 9 subjects, which allowed us to assess inter-individual variability. We found that features of the SHM process associated to the biochemistry of the process, such as the positional profiles, the ratio of the number of indels to point mutations, and the distribution of insertion and deletion lengths, were mostly conserved across subjects. By contrast, the distribution of the mutational age of antibodies, which reflects their unique immune history, was specific to each individual. While we studied individual sequences, we know that those are organized in distinct lineages originating from naive germline sequences. This could lead to multiple-counting mutations occurring in several sequences of the same lineage. However, since this effect is random and decoupled from the mutation process itself in nonproductive sequences, we only expect it to increase errors due to sampling. The low individual variability reported for most summary statistics provides an upper bound on that error.

We report a large positional heterogeneity, with a similar relative variability between point and indel hypermutations. Indels tend to concentrate around the CDR1 and CDR2 loops, as do point mutations [51], and consistently with previous reports on productive sequences [33]. However, our results show that the positional profiles of point, insertion, and deletion SHM are still distinct (albeit correlated), suggesting a complex picture of how the 3 processes result from AID selectivity and positional preference.

The precise mechanism causing SHM indels and point mutations is still elusive. Our results confirm the main prediction of the polymerase slippage model [45, 46] for indels, namely that inserted segments are homologous to their immediately flanking regions. This complements earlier observations from [33], but without the interference of selection, and on a much bigger dataset allowing for more detailed statistics. By contrast, we found that deleted segments are not significantly homologous to their neighborhood, suggesting the slippage is not driven by sequence homology, but reaches out to a random position along the DNA. An alternative hypothesis is that deletions are driven by double-strand breaks caused by AID, consistent with our and previous observations that they tend to co-localize with point mutations [34, 49]. In

the slippage model, the indel length corresponds to the distance by which the polymerase slips. Both have a characteristic exponential tail in their distribution, compatible with the processive motion of the slipping polymerase with random stopping.

Our probabilistic approach is not only useful for annotating sequences and inferring rates, but it can also be used as a generative model to produce synthetic datasets for computational controls. However, the model makes a few simplifying assumptions for the sake of computational tractability. First, its positional profiles are constant, in contradiction with the data. Second, inserted segments are assumed to be uniformly random, ignoring homology to flanking regions. While these approximations are unlikely to affect our results, future refinements of the generative model would be welcome to create more realistic datasets.

We evaluated the conditions under which the probabilistic approach we introduced added value to the classical deterministic annotation of mutated antibodies. For most of the bulk repertoire, where the rate of hypermutation is relatively low, both approaches yielded similar results, with a notable exception: the fraction of “naive” sequences that have not undergone any affinity maturation. This definition of naive is distinct from usual phenotypic characterization based on isotype classification, flow cytometry [52] or single-cell RNA sequencing data [53], although they are correlated. This proportion is always overestimated by the deterministic method. Since the SHM process is stochastic, some sequences having undergone affinity maturation may not have accumulated any mutation at all. These sequences are always considered naive by the deterministic approach, while the probabilistic method considers both possibilities, correcting that bias.

Beyond the bulk of the repertoire, the probabilistic approach is most useful for analyzing heavily mutated sequences, where it is difficult to distinguish point mutations from nearby insertion and deletion events. While a minority, such sequences are very important. They are found in the largest clonal families having undergone the longest affinity maturation, and so are the most experienced and selected ones. HIV1-related bnAbs, which often contain many indels, are examples of such sequences [54]. The relatively low frequency of indels that we report could explain why effective bnAbs are rare, and only emerge naturally in a small percentage of infected individuals. In that regard, evidence that indels frequently appear in response to some vaccination protocols [32] suggests promising strategies for HIV vaccination.

Methods

Data and preprocessing

We exploited publicly available IgM and IgG repertoire data from a recent ultra-deep repertoire study of immunoglobulin heavy chains (IgH), including approximately 3 billions sequences in total, coming from 9 different healthy donors [8]. The V region is well covered by the mRNA sequencing protocol, starting from nucleotide position ~ 66 along the templates for almost all the sequences; CDR3 and J gene are also covered.

We first preprocessed raw sequences as in [22], using pRESTO (release 0.5.13) from the Immcantation pipeline [55]. Then, we performed a standard annotation through the IgBLAST standalone tool [48], (release 1.13.0, with default values for penalties), referring to IMGT for germline templates [50] (only functional (F) genes are considered, totaling 265 IGHV different templates, alleles included). We also applied quality filters, keeping only sequences with properly annotated V gene, J gene and CDR3, and covering at least 200 base pairs of the V segment. We checked using synthetic data that this quality filter does not substantially alter the underlying indel distribution below length 60, although sequences with very long indels are slightly less likely to pass it (S6 Fig).

Unique molecular identifiers (UMI) were used in the original study to correct for PCR amplification biases. We grouped reads with the same UMI to build consensus sequences for downstream analysis in which read counts were ignored, allowing for a maximum error rate of 10% within each group. Sequences corresponding to low numbers of reads per UMI (e.g. 1 or 2) are expected to be more affected by sequencing errors, because of the inability to build a consensus sequence from multiple reads. We verified that our results are not affected by this potential source of error, by repeating the standard annotation analysis on UMIs represented by at least 3 reads (S7 Fig). Thus, to make the most of the available data, we did not discard sequences with low number of reads per UMI.

It is very hard to call mutations in the CDR3 because of its variability, and the J region is very short. Therefore we restricted ourselves to the V region, starting from the beginning of the read, up to the conserved cysteine before the beginning of the CDR3. We call that sequence s .

Modeling approach

We start from a repertoire $\{s\}$ of size N . Let $\{t\}$ be the set of (truncated) V templates used for the alignment (the same 265 templates exploited for IgBLAST annotation). For each s and t , a deterministic alignment algorithm, based on the standard Needleman-Wunsch (NW) algorithm for pairwise global alignments with affine score for gaps [56], provides *only* the alignment with the highest score $\mathcal{S}(s, t)$.

The value of the parameters used in the NW algorithm were chosen according to the actual frequency of related events in human IgG heavy chains, as estimated through IgBLAST annotation (Fig 1). Assuming an average point-mutation rate of $\sim 5\%$, an indel rate of 0.05%, and an exponentially decaying length distribution of indels with characteristic length $\ell = 10$, we have the following log-odd penalties (using natural logarithm): $\simeq -0.05$ for matching and $\simeq -3$ for mismatching aligned base pairs; $\simeq -7.7$ for opening a deletion gap; $\simeq -0.1$ for extending it; $\simeq -9.1$ for opening an insertion gap; and -1.5 for extending it (including the log(4) ≈ 1.4 penalty for choosing the base pair) [56]. Note that since the alignment is global, only relative values of the penalty matter.

The optimal alignment is built through a dynamic-programming strategy, relying recursively on the optimal alignments of shorter blocks. The key modification we introduce to the standard NW algorithm is to sum over all the possibilities during the recursive scheme, rather than choosing the optimal one at each step. The result is then no longer an optimal global alignment score, but rather the sum of the likelihoods of all the possible global alignments between s and t , summed over all templates t . Each possible alignment contributes to the final sum proportionally to its likelihood, yielding a *probabilistic* likelihood of the sequence, by contrast with the *deterministic* score of the standard NW algorithm.

The resulting likelihood of the sequence s , $\mathcal{L}(s; \theta)$, depends on the model parameters θ . The total likelihood of the whole repertoire is then given by:

$$\mathcal{L}(\theta) = \prod_s \mathcal{L}(s; \theta), \quad (1)$$

where we assumed independent sequences, ignoring the lineage structure of the repertoire (see Discussion). Maximizing the likelihood with respect to θ gives an estimate $\hat{\theta}$ for the alignment parameters that best describe the data.

Recursive algorithm for the alignment

We now give the details of the computation of $\mathcal{L}(s; \theta)$. Consider a sequence s and a template t . Indexing starts from zero, including an initial “placeholder” position for potential gaps,

running over actual base pairs from s and t (gaps are not taken into account directly in this indexing). So s_i is the i -th symbol along s , while $s_{0:i}$ is the portion of s running from the first actual symbol s_1 to i -th symbol s_i included, preceded by the “empty” placeholder position 0.

As in NW, the alignment procedure always starts from the beginning of the two sequences s and t , and goes through all the possibilities by extending the previous alignment for shorter pieces of s and t . In our case, however, we sum over these possibilities, rather than keep the best-scoring one. The indices on s and t are not constrained to increase at the same time; an asymmetric increase corresponds to a deletion or an insertion.

The alignment score for the portion s up to position i included ($s_{0:i}$) and the portion of t up to position j included ($t_{0:j}$) is computed by relying on the previous alignments of shorter portions, according to the following possibilities:

- s_{i-1} and t_{j-1} were previously aligned to each other, so the alignment advances by one position on both sequences, with a *match* ($s_i = t_j$) or a *mismatch* ($s_i \neq t_j$);
- t_j was previously aligned with s_k , $k < i$, so that a gap on s of length $i - k$ has to be taken into account (i. e., a *deletion* of templated bases);
- s_i was previously aligned with $t_{k'}$, $k' < j$, so that a gap on t of length $j - k'$ has to be taken into account (i. e., an *insertion* of non-templated bases).

In the standard NW algorithm, each of these possibilities is linked to a score, and only the largest among them is kept. Here, instead, we work directly with likelihoods, and then we sum over them according to the following recursion:

$$\begin{aligned} \mathcal{S}(s_{0:i}, t_{0:j}) = & \mathcal{S}(s_{0:i-1}, t_{0:j-1}) M(s_i, t_j) \\ & + \sum_{k < i} \mathcal{S}(s_{0:k}, t_{0:j}) \Gamma_{\text{del}}(i - k) \\ & + \sum_{k' < j} \mathcal{S}(s_{0:i}, t_{0:k'}) \Gamma_{\text{ins}}(j - k') \end{aligned} \tag{2}$$

where $M(s_i, t_j)$ is the match/mismatch probability between nucleotide s_i and t_j , $\Gamma_{\text{del}}(\ell)$ is the probability of a deletion of length ℓ , and $\Gamma_{\text{ins}}(\ell)$ is the probability of an insertion of length ℓ . The initial conditions of the recursion are:

$$\begin{aligned} \mathcal{S}(s_0, t_0) & = 1 \\ \mathcal{S}(s_{0:i}, t_0) & = \sum_{k < i} \mathcal{S}(s_{0:k}, t_0) \Gamma_{\text{del}}(i - k) \\ \mathcal{S}(s_0, t_{0:j}) & = \sum_{k' < j} \mathcal{S}(s_0, t_{0:k'}) \Gamma_{\text{ins}}(j - k'). \end{aligned} \tag{3}$$

The final score $\mathcal{S}(s, t) \equiv \mathcal{S}(s_{0:L_s}, t_{0:L_t})$ gives the likelihood of s aligning to t , obtained as the sum over all possible alignments.

Likelihoods M , P_{del} and P_{ins} can be further described as:

$$\begin{aligned} \Gamma_{\text{del}}(\ell) & \equiv \mu_{\text{del}} \times P_{\text{del}}(\ell) \\ \Gamma_{\text{ins}}(\ell) & \equiv \mu_{\text{ins}} \times P_{\text{ins}}(\ell) \times (1/4)^\ell, \end{aligned} \tag{4}$$

assuming uniformly random insertions, where $P_{\text{del/ins}}(\ell)$ is the distribution of deletion and

insertion lengths, and:

$$M(s_i, t_j) \equiv (1 - \mu_{\text{del}} - \mu_{\text{ins}}) \times [(1 - \mu) \mathbb{I}(s_i = t_j) + \mu \mathbb{I}(s_i \neq t_j)] \tag{5}$$

with $\mathbb{I}(\cdot)$ being the indicator function (equal to one if argument is true, otherwise equal to zero). The rates μ_{del} , μ_{ins} , and μ are interpreted as the deletion, insertion, and point mutation rates per base pair.

Based on evidence from data (Fig 1C), we assume that these rates scale together with the mutational age, since both types of mutations accumulate with time during affinity maturation cycles. We thus write:

$$\mu \equiv \mu_s, \quad \mu_{\text{del}} \equiv \mu_s \times \beta_{\text{del}}, \quad \mu_{\text{ins}} \equiv \mu_s \times \beta_{\text{ins}}, \tag{6}$$

where β_{del} and β_{ins} are constant (repertoire-specific) ratios, and μ_s is sequence specific, as suggested by the wide range of mutabilities in the data. The ratios $\beta_{\text{del/ins}}$ will also be referred to as *rescaled* deletion and insertion rates. The mutation rate μ_s is treated as a hidden variable, whose distribution $P(\mu)$ is a model parameter.

The likelihood of each sequence is obtained by summing over all mutation rates and templates:

$$\mathcal{L}(s; \theta) = \int d\mu_s P(\mu_s) \mathcal{L}(s|\mu_s; \phi), \tag{7}$$

$$\text{with } \mathcal{L}(s|\mu_s; \phi) \equiv \sum_t p_t \mathcal{S}(s, t|\mu_s; \phi), \tag{8}$$

with p_t uniform, and where the parameters θ have been separated into two components: $\theta = (P(\mu), \phi)$, with $\phi \equiv (\beta_{\text{del}}, \beta_{\text{ins}}, P_{\text{del}}(\ell), P_{\text{ins}}(\ell))$. In Eq (8) we made explicit the dependence of alignment score on the mutational age μ_s and indel parameters ϕ .

The repertoire is composed of a mixture of cells that have undergone some mutational process, presumably in germinal centers, and cells that have stayed naive. To properly model this, we further assume that the prior distribution of mutation rates has two parts:

$$P(\mu) = f \delta(\mu) + (1 - f) \tilde{P}(\mu), \tag{9}$$

where f is the fraction of naive cells, and $\tilde{P}(\mu)$ is a smooth probability density.

Likelihood maximization

To maximize the likelihood, Eqs (1) and (7), we employ a combination of Expectation-Maximization (EM) [57, 58] for updating $P(\mu)$, and gradient ascent for updating ϕ .

We denote by $\theta^t = (P^t(\mu), \phi^t)$ the value of the parameters at iteration t . Following EM, we define the pseudo log-likelihood to be maximized with respect to θ at each iteration:

$$\theta^{t+1} \leftarrow \arg \max_{\theta} Q(\theta|\theta^t), \tag{10}$$

with

$$Q(\theta|\theta^t) = \sum_s \langle \log[\mathcal{L}(s|\mu_s; \phi) P(\mu_s)] \rangle_{\mu_s|s;\theta^t}, \tag{11}$$

where the mean $\langle \cdot \rangle_{\mu_s | s; \theta^t}$ is with respect to the posterior distribution:

$$P(\mu_s | s; \theta^t) \equiv \frac{\mathcal{L}(s | \mu_s; \phi^t) P^t(\mu_s)}{\int d\mu_s \mathcal{L}(s | \mu_s; \phi^t) P^t(\mu_s)}. \tag{12}$$

Because of the separation between naive $\mu_s = 0$ and experienced cells, this posterior may also be decomposed as:

$$P(\mu_s | s; \theta^t) = f_s^t \delta(\mu_s) + (1 - f_s^t) \tilde{P}(\mu_s | s; \theta^t), \tag{13}$$

with

$$f_s^t = \frac{f^t \mathcal{L}(s | 0; \phi^t)}{f^t \mathcal{L}(s | 0; \phi^t) + (1 - f^t) \int d\mu_s \mathcal{L}(s | \mu_s; \phi^t) \tilde{P}^t(\mu_s)}, \tag{14}$$

and

$$\tilde{P}(\mu_s | s; \theta^t) = \frac{\mathcal{L}(s | \mu_s; \theta^t) \tilde{P}^t(\mu_s)}{\int d\mu_s \mathcal{L}(s | \mu_s; \theta^t) \tilde{P}^t(\mu_s)}. \tag{15}$$

The pseudo-likelihood Q can be decomposed into two terms corresponding to the two factors inside the logarithm of Eq (11), $\log \mathcal{L}(s | \mu_s; \phi)$ and $\log P(\mu_s)$. The first one only depends on θ through ϕ , and the second one only through $P(\mu)$, allowing for their independent maximization.

Maximization with respect to $P(\mu)$. To maximize Q under the normalization constraint, we define the functional with Lagrange multipliers:

$$\hat{Q}(\theta | \theta^t) \equiv Q(\theta | \theta^t) + \lambda(1 - \int d\mu P(\mu)). \tag{16}$$

The extremal condition with respect to $P(\mu')$ gives:

$$0 = \frac{\delta \hat{Q}}{\delta P(\mu')} = \sum_s \int d\mu_s P(\mu_s | s; \theta^t) \frac{\delta(\mu_s - \mu')}{P(\mu_s)} - \lambda \tag{17}$$

When plugging functional forms (13) and (9) in it, together with the extremal condition with respect to λ , we get:

$$\tilde{P}^{t+1}(\mu) \leftarrow \frac{\sum_s (1 - f_s^t) \tilde{P}(\mu | s; \theta^t)}{\sum_s (1 - f_s^t)} \tag{18}$$

$$f^{t+1} \leftarrow f = \frac{1}{N} \sum_s f_s^t. \tag{19}$$

Intuitively, the new estimate of the smooth part $\tilde{P}(\mu)$ of the repertoire-wide distribution is the weighted average of the sequence-specific posterior distributions over the repertoire. Similarly, the new estimate for the fraction f of naive cells is given by the average over the repertoire of the probability of being naive.

Maximization with respect ϕ . To maximize the first term of Q with respect to ϕ , we use Monte Carlo sampling to represent the integral over the posterior $P(\mu_s | s; \theta^t)$:

$$\sum_s \langle \log \mathcal{L}(s | \mu_s; \phi) \rangle_{\mu_s | s; \theta^t} \simeq \sum_s \sum_{n=1}^{N_{MC}} \log \mathcal{L}(s | \mu_{s,n}; \phi), \tag{20}$$

with $\mu_{s,n} \sim P(\mu_s | s; \theta^t)$.

Then, the extremal condition with respect to ϕ_i reads:

$$\nabla_i^t \equiv \frac{\partial}{\partial \phi_i} \sum_s \sum_n \log \mathcal{L}(s|\mu_{s,n}; \phi) = 0. \tag{21}$$

Since we could not find an exact update rule for ϕ_i , we relied on gradient ascent to maximize the likelihood. Each gradient component ∇_i^t can be estimated numerically by infinitesimally shifting ϕ_i :

$$\nabla_i^t \simeq \sum_s \sum_n \frac{\log \mathcal{L}_s(s|\mu_{s,n}; \tilde{\phi}_i^t) - \log \mathcal{L}_s(s|\mu_{s,n}; \phi^t)}{\varepsilon}, \tag{22}$$

where $\tilde{\phi}_i^t$ is equal to ϕ^t for all the components but i , which has been increased by ε .

We updated ϕ according to a momentum gradient-ascent update rule:

$$\phi^{t+1} \leftarrow \mathcal{P}(\phi^t + \alpha^t \nabla^t + \omega^t (\phi^t - \phi^{t-1})), \tag{23}$$

with learning rate α^t and inertial term ω^t , whose dependence on the iteration step t was chosen to optimize convergence and to avoid long-time oscillations [59]:

$$\alpha^t \propto e^{-2t/T}, \quad \omega^t \propto \frac{t}{t+3}, \tag{24}$$

with T being the maximum number of time iterations allowed. In Eq (23), \mathcal{P} denotes the projection onto the simplex defined by the constraints $\beta_{\text{del}} \geq 0, \beta_{\text{ins}} \geq 0, P_{\text{del}}(\ell) \geq 0, P_{\text{ins}}(\ell) \geq 0$, and $\sum_{\ell} P_{\text{del}}(\ell) = \sum_{\ell} P_{\text{ins}}(\ell) = 1$. Projection was done using the procedure described in Ref. [60].

Speeding up the computation

The algorithm described so far is computationally very costly. Just the basic step of computing the alignment likelihood $\mathcal{L}(s|\mu_s; \phi)$ for a single sequence at fixed μ_s is time-consuming: if we allow for a maximum size $\ell = \Theta$ for single-event deletions and insertions, then the requested number of operations roughly scales as $L_s \times \langle L_t \rangle \times N_t \times (2\Theta + 1)$, where N_t is the number of V templates considered and $\langle L_t \rangle$ is their average length. This has to be repeated for each sequence, for each mutation rate, for each template, and for each direction of the gradient, and all of this at each iteration. Below we describe approximations that considerably speed up the code.

Pruning the alignment matrix. Indels are quite rare in real Ig sequences, with most sequences having none or at most one, located around well defined regions. Allowing for possible indels (of size Θ at most) everywhere in every sequence contributes very little to the final cumulated alignment likelihood, wasting computational time. Dropping implausible terms would dramatically reduce the computational cost, pushing the factor $L_s \times \langle L_t \rangle \times (2\Theta + 1)$ almost down to L_s .

Plausible locations for gaps can be identified by computing, for each pair of positions (i, j) along s and t , the likelihood that the alignment goes through that pair. This is given by:

$$\mathcal{S}_{ij} \equiv \mathcal{S}(s_{1:i}, t_{1:j}) \times \mathcal{S}(s_{i+1:L_s}, t_{j+1:L_t}), \tag{25}$$

where the first term is computed as described above, and the second term similarly, but using a backward iteration on the reverted sequences. Normalizing by the total alignment likelihood, we obtain a relative probability that the correct alignment passes through (i, j) ,

$$\hat{\mathcal{S}}_{ij} = \mathcal{S}_{ij} / \mathcal{S}(s, t).$$

The entries of this matrix can now be used to *prune* the terms of the recursive Eq (2). When \hat{S}_{ij} is smaller than a fixed threshold $\vartheta = 10^{-5}$, we set $\mathcal{S}(s_{0:i}, t_{0:j})$ to zero and avoid computing it and including it in further sums.

To implement this strategy, we need to first run the forward and backward iterations without any pruning, for an initial guess of the parameters (μ_s, ϕ) . The \hat{S}_{ij} matrix computed in this way is then used to define the pairs of positions to be kept in future computations. After that, all subsequent evaluations of $\mathcal{S}(s, t)$ in the optimization algorithm are done with only those pairs of positions.

Reducing the number of templates. Another important performance improvement comes from pruning the list of plausible templates for each sequence s . In principle, the sum in Eq (8) runs over all the templates t , but in practice only one of them actually contributes in a non negligible way. Thus, we only keep the V template t_s^* closest to s , as given by the alignment score of the standard NW algorithm:

$$\mathcal{L}(s|\mu_s; \phi) = \sum_t p_t \mathcal{S}(s, t|\mu_s; \phi) \simeq \mathcal{S}(s, t_s^*|\mu_s; \phi). \tag{26}$$

The computational cost of $\mathcal{L}(s; \theta)$ is then reduced further by a factor N_t .

Averaging over the hidden age factor. Another demanding operation is the average over the hidden variable μ_s by Monte-Carlo sampling. In principle, the number of Monte-Carlo samples N_{MC} should be very large. In practice, we used a single $N_{MC} = 1$ sample at each iteration. This simplification works because parameters evolve slowly at each iteration, especially in the last convergence steps, which allows for time averaging as a substitute for large MC sampling.

Posterior approximation. To calculate and sample from the posterior $P(\mu_s|s; \theta)$, we evaluated it on a finite set of values of μ (called nodes), and interpolated the rest of the function using a smooth piecewise cubic interpolation [61]. The nodes are set for each sequences to be placed around the naive estimate of the mutation rate μ_s based on the NW alignment, in addition to 2 key nodes at $\mu = 0$ and 1. We used from 15 to 25 nodes depending on how mutated the sequence was.

As both alignment parameters ϕ and repertoire-wide distribution $P(\mu)$ get updated during inference, the positions of the nodes are updated for each sequence as well, by placing them around the previously estimated maximum of the posterior with a spread given by its inverse curvature.

The prior distribution $P(\mu)$ was numerically stored by binning uniformly the $[0, 1]$ interval with a bin size of 0.0005.

Generation of synthetic data

In order to mimic real distributions of the repertoire-wide point-mutation rate $P(\mu)$, we used a *shifted* Gamma distribution:

$$P(\mu) = \frac{1}{z(\alpha, \beta, \mu_0)} (\mu + \mu_0)^{\alpha-1} e^{-\beta\mu} \tag{27}$$

with α, β determined by the mode and standard deviation of the unshifted Gamma distribution: $(\alpha-1)/\beta = 0.07$ and $\sqrt{\alpha}/\beta = 0.04$; and where $z(\alpha, \beta, \mu_0)$ is a normalization constant. In practice, we drew values from the Gamma distribution with parameters α, β , and then subtracted 0.02 from it, discarding negative values. The resulting distribution has non-zero probability for infinitesimally small values and a mode of $\mu = 0.05$. We used this distribution for generating synthetic data when validating the software and reported as ground truth in Fig 3.

Supporting information

S1 Fig. Fraction of times the overlap between inserted base pairs and same-length flanking region on the 3' end is larger than the overlap with the 5' end along the sequence; averages and one standard deviation over the 9 Briney donors given by solid line and shaded area around, respectively. For comparison, overlap between random insertions and flanking regions is also reported (details in the main text). A weak preference for the 5' end is supported by a KS-test p-value of 0.04, when comparing true insertions with randomized ones. (TIFF)

S2 Fig. Overlap between deleted base pairs and same-length flanking regions on either 3' or 5' end along the sequence (larger is kept); averages and one standard deviation over the 9 Briney donors given by solid line and shaded area around, respectively. For comparison, overlap between random deletions and flanking regions is also reported (details in the main text). (TIFF)

S3 Fig. Distribution of the distance (in base pairs) separating (A) deletions or (B) insertions from the closest point mutation in nonproductive IgG sequences, normalized by the null expectation obtained by reshuffling indel and point mutations between sequences with the same V gene (to control for the positional biases of Fig 2). Average and error bars are over the 9 donors. (TIFF)

S4 Fig. (A) Probabilistic vs deterministic estimates of the rescaled deletion rate β_{del} at increasing values of indels density. Averages values plus one standard deviation error bars are obtained over the \mathcal{N} independent synthetic repertoires. **(B) Deletion length profiles for the largest value of β_{del} considered in panel (A);** mean and standard deviations over the \mathcal{N} independent realizations. (TIFF)

S5 Fig. Model prediction for frequency of (A) deletions, (B) insertions, and (C) both insertions and deletions, in IgG nonproductive sequences, compared to the data. Mean and variance are over the 9 donors. Poisson distributions with the same means as the data are shown for comparison. (TIFF)

S6 Fig. Distribution of deletion lengths in 5 synthetic repertoires of 500, 000 sequences each. Synthetic sequences were generated as in Fig 3D, but with a cut-off of 100 base pairs instead of 30 for indel lengths, and with indel-to-point-mutation rates $\beta_{\text{del}} = \beta_{\text{ins}} = 0.025$. The blue curve shows the true distribution of deletions lengths in the full dataset, while the orange curve shows the same distribution in sequences that have passed the quality filters described in Methods section. Mean and standard deviation over the 5 subsets shown. (TIFF)

S7 Fig. Indel statistics from Fig 1, compared to those obtained from sequences with at least 3 reads per UMI. No systematic error due to sequencing errors (expected to be larger in UMI with low counts) can be detected. **(A) Mutation rates per base pair, for both IgM and IgG. (B) Length profiles for deletions in IgG. (C) Length profiles for insertions in IgG.** Vertical lines show mean values. (TIFF)

Author Contributions

Conceptualization: Aleksandra M. Walczak, Thierry Mora.

Data curation: Natanael Spisak.

Formal analysis: Cosimo Lupo.

Funding acquisition: Aleksandra M. Walczak.

Investigation: Cosimo Lupo.

Methodology: Cosimo Lupo, Aleksandra M. Walczak, Thierry Mora.

Software: Cosimo Lupo.

Supervision: Aleksandra M. Walczak, Thierry Mora.

Validation: Cosimo Lupo.

Visualization: Cosimo Lupo.

Writing – original draft: Cosimo Lupo, Aleksandra M. Walczak, Thierry Mora.

Writing – review & editing: Aleksandra M. Walczak, Thierry Mora.

References

1. Hozumi N, Tonegawa S. Evidence for somatic rearrangement of immunoglobulin genes coding for variable and constant regions. *Proc Natl Acad Sci.* 1976; 73(10):3628. <https://doi.org/10.1073/pnas.73.10.3628> PMID: 824647
2. Boyd SD, Marshall EL, Merker JD, Maniar JM, Zhang LN, Sahaf B, et al. Measurement and Clinical Monitoring of Human Lymphocyte Clonality by Massively Parallel V-D-J Pyrosequencing. *Sci Transl Med.* 2009; 1(12):12ra23. <https://doi.org/10.1126/scitranslmed.3000540> PMID: 20161664
3. Glanville J, Zhai W, Berka J, Telman D, Huerta G, Mehta GR, et al. Precise determination of the diversity of a combinatorial antibody library gives insight into the human immunoglobulin repertoire. *Proc Natl Acad Sci.* 2009; 106(48):20216. <https://doi.org/10.1073/pnas.0909775106> PMID: 19875695
4. Larimore K, McCormick MW, Robins HS, Greenberg PD. Shaping of Human Germline IgH Repertoires Revealed by Deep Sequencing. *J Immunol.* 2012; 189(6):3221. <https://doi.org/10.4049/jimmunol.1201303> PMID: 22865917
5. Elhanati Y, Sethna Z, Marcou Q, Callan CG Jr, Mora T, Walczak AM. Inferring processes underlying B-cell repertoire diversity. *Philos Trans R Soc B Biol Sci.* 2015; 370(1676):20140243. <https://doi.org/10.1098/rstb.2014.0243> PMID: 26194757
6. DeWitt WS, Lindau P, Snyder TM, Sherwood AM, Vignali M, Carlson CS, et al. A Public Database of Memory and Naive B-Cell Receptor Sequences. *PLoS One.* 2016; 11(8):e0160853. <https://doi.org/10.1371/journal.pone.0160853> PMID: 27513338
7. Marcou Q, Mora T, Walczak AM. High-throughput immune repertoire analysis with IGoR. *Nat Commun.* 2018; 9(1):561. <https://doi.org/10.1038/s41467-018-02832-w> PMID: 29422654
8. Briney B, Inderbitzin A, Joyce C, Burton DR. Commonality despite exceptional diversity in the baseline human antibody repertoire. *Nature.* 2019; 566(7744):393. <https://doi.org/10.1038/s41586-019-0879-y> PMID: 30664748
9. Elsner RA, Shlomchik MJ. Germinal Center and Extrafollicular B Cell Responses in Vaccination, Immunity, and Autoimmunity. *Immunity.* 2020; 53(6):1136–1150. <https://doi.org/10.1016/j.immuni.2020.11.006> PMID: 33326765
10. Victora GD, Nussenzweig MC. Germinal Centers. *Annu Rev Immunol.* 2012; 30(1):429. <https://doi.org/10.1146/annurev-immunol-020711-075032> PMID: 22224772
11. Cobey S, Wilson PC, Matsen FA IV. The evolution within us. *Philos Trans R Soc B Biol Sci.* 2015; 370(1676):20140235. <https://doi.org/10.1098/rstb.2014.0235> PMID: 26194749
12. Mesin L, Ersching J, Victora GD. Germinal Center B Cell Dynamics. *Immunity.* 2016; 45(3):471. <https://doi.org/10.1016/j.immuni.2016.09.001> PMID: 27653600
13. Feng Y, Seija N, Di Noia JM, Martin A. AID in Antibody Diversification: There and Back Again. *Trends Immunol.* 2020; 41(7):P586. <https://doi.org/10.1016/j.it.2020.04.009> PMID: 32434680

14. Kleinstei SH, Louzoun Y, Shlomchik MJ. Estimating Hypermutation Rates from Clonal Tree Data. *J Immunol.* 2003; 171(9):4639. <https://doi.org/10.4049/jimmunol.171.9.4639> PMID: 14568938
15. Odegard VH, Schatz DG. Targeting of somatic hypermutation. *Nat Rev Immunol.* 2006; 6(8):573. <https://doi.org/10.1038/nri1896> PMID: 16868548
16. Yaari G, Vander Heiden JA, Uduman M, Gadala-Maria D, Gupta N, Stern JNH, et al. Models of Somatic Hypermutation Targeting and Substitution Based on Synonymous Mutations from High-Throughput Immunoglobulin Sequencing Data. *Front Immunol.* 2013; 4:358. <https://doi.org/10.3389/fimmu.2013.00358> PMID: 24298272
17. McCoy CO, Bedford T, Minin VN, Bradley P, Robins H, Matsen FA IV. Quantifying evolutionary constraints on B-cell affinity maturation. *Philos Trans R Soc B Biol Sci.* 2015; 370(1676):20140244. <https://doi.org/10.1098/rstb.2014.0244> PMID: 26194758
18. Cui A, Di Niro R, Vander Heiden JA, Briggs AW, Adams K, Gilbert T, et al. A Model of Somatic Hypermutation Targeting in Mice Based on High-Throughput Ig Sequencing Data. *J Immunol.* 2016; 197(9):3566. <https://doi.org/10.4049/jimmunol.1502263> PMID: 27707999
19. Sheng Z, Schramm CA, Kong R, Mullikin JC, Mascola JR, Kwong PD, et al. Gene-Specific Substitution Profiles Describe the Types and Frequencies of Amino Acid Changes during Antibody Somatic Hypermutation. *Front Immunol.* 2017; 8. <https://doi.org/10.3389/fimmu.2017.00537> PMID: 28539926
20. Hoehn KB, Lunter G, Pybus OG. A Phylogenetic Codon Substitution Model for Antibody Lineages. *Genetics.* 2017; 206(1):417. <https://doi.org/10.1534/genetics.116.196303> PMID: 28315836
21. Dhar A, Davidsen K, Matsen FA IV, Minin VN. Predicting B cell receptor substitution profiles using public repertoire data. *PLOS Comput Biol.* 2018; 14(10):e1006388. <https://doi.org/10.1371/journal.pcbi.1006388> PMID: 30332400
22. Spisak N, Walczak AM, Mora T. Learning the heterogeneous hypermutation landscape of immunoglobulins from high-throughput repertoire data. *Nucleic Acids Res.* 2020; 48(19):10702. <https://doi.org/10.1093/nar/gkaa825> PMID: 33035336
23. Wilson PC, de Bouteiller O, Liu YJ, Potter K, Banchereau J, Capra JD, et al. Somatic Hypermutation Introduces Insertions and Deletions into Immunoglobulin V Genes. *J Exp Med.* 1998; 187(1):59. <https://doi.org/10.1084/jem.187.1.59> PMID: 9419211
24. Wilson PC, Liu YJ, Banchereau J, Capra JD, Pascual V. Amino acid insertions and deletions contribute to diversify the human Ig repertoire. *Immunol Rev.* 1998; 162(1):143. <https://doi.org/10.1111/j.1600-065X.1998.tb01437.x> PMID: 9602360
25. Klein U, Goossens T, Fischer M, Kanzler H, Braeuninger A, Rajewsky K, et al. Somatic hypermutation in normal and transformed human B cells. *Immunol Rev.* 1998; 162(1):261. <https://doi.org/10.1111/j.1600-065X.1998.tb01447.x> PMID: 9602370
26. Fischer M, Küppers R. Human IgA- and IgM-secreting intestinal plasma cells carry heavily mutated VH region genes. *Eur J Immunol.* 1998; 28(9):2971. [https://doi.org/10.1002/\(SICI\)1521-4141\(199809\)28:09%3C2971::AID-IMMU2971%3E3.0.CO;2-3](https://doi.org/10.1002/(SICI)1521-4141(199809)28:09%3C2971::AID-IMMU2971%3E3.0.CO;2-3) PMID: 9754584
27. Goossens T, Klein U, Küppers R. Frequent occurrence of deletions and duplications during somatic hypermutation: Implications for oncogene translocations and heavy chain disease. *Proc Natl Acad Sci.* 1998; 95(5):2463. <https://doi.org/10.1073/pnas.95.5.2463> PMID: 9482908
28. Ohlin M, Borrebaeck CAK. Insertions and deletions in hypervariable loops of antibody heavy chains contribute to molecular diversity. *Mol Immunol.* 1998; 35(4):233. [https://doi.org/10.1016/S0161-5890\(98\)00030-3](https://doi.org/10.1016/S0161-5890(98)00030-3) PMID: 9736339
29. de Wildt RMT, van Venrooij WJ, Winter G, Hoet RMA, Tomlinson IM. Somatic insertions and deletions shape the human antibody repertoire. *J Mol Biol.* 1999; 294(3):701. <https://doi.org/10.1006/jmbi.1999.3289> PMID: 10610790
30. Küppers R, Goossens T, Klein U. The Role of Somatic Hypermutation in the Generation of Deletions and Duplications in Human Ig V Region Genes and Chromosomal Translocations. In: Melchers F, Potter M, editors. *Mech. B Cell Neoplasia 1998. Curr. Top. Microbiol. Immunol.* vol 246. Berlin: Springer; 1999. p. 193. Available from: http://link.springer.com/10.1007/978-3-642-60162-0_24.
31. Bemark M, Neuberger MS. By-products of immunoglobulin somatic hypermutation. *Genes, Chromosom Cancer.* 2003; 38(1):32. <https://doi.org/10.1002/gcc.10241> PMID: 12874784
32. Reason DC, Zhou J. Codon insertion and deletion functions as a somatic diversification mechanism in human antibody repertoires. *Biol Direct.* 2006; 1(1):24. <https://doi.org/10.1186/1745-6150-1-24> PMID: 16942619
33. Briney BS, Willis JR, Crowe JE. Location and length distribution of somatic hypermutation-associated DNA insertions and deletions reveals regions of antibody structural plasticity. *Genes & Immun.* 2012; 13(7):523. <https://doi.org/10.1038/gene.2012.28> PMID: 22717702

34. Yeap LS, Hwang JK, Du Z, Meyers RM, Meng FL, Jakubauskaitė A, et al. Sequence-Intrinsic Mechanisms that Target AID Mutational Outcomes on Antibody Genes. *Cell*. 2015; 163(5):1124. <https://doi.org/10.1016/j.cell.2015.10.042> PMID: 26582132
35. Zhou J, Lottenbach KR, Barenkamp SJ, Reason DC. Somatic Hypermutation and Diverse Immunoglobulin Gene Usage in the Human Antibody Response to the Capsular Polysaccharide of *S treptococcus pneumoniae* Type 6B. *Infect Immun*. 2004; 72(6):3505. <https://doi.org/10.1128/IAI.72.6.3505-3514.2004> PMID: 15155658
36. Wu X, Yang ZY, Li Y, Hogerkorp CM, Schief WR, Seaman MS, et al. Rational Design of Envelope Identifies Broadly Neutralizing Human Monoclonal Antibodies to HIV-1. *Science*. 2010; 329(5993):856. <https://doi.org/10.1126/science.1187659> PMID: 20616233
37. Walker LM, Phogat SK, Chan-Hui PY, Wagner D, Phung P, Goss JL, et al. Broad and Potent Neutralizing Antibodies from an African Donor Reveal a New HIV-1 Vaccine Target. *Science*. 2009; 326(5950):285. <https://doi.org/10.1126/science.1178746> PMID: 19729618
38. Walker LM, Huber M, Doores KJ, Falkowska E, Pejchal R, Julien JP, et al. Broad neutralization coverage of HIV by multiple highly potent antibodies. *Nature*. 2011; 477(7365):466. <https://doi.org/10.1038/nature10373> PMID: 21849977
39. Kepler TB, Liao HX, Alam SM, Bhaskarabhatla R, Zhang R, Yandava C, et al. Immunoglobulin Gene Insertions and Deletions in the Affinity Maturation of HIV-1 Broadly Reactive Neutralizing Antibodies. *Cell Host & Microbe*. 2014; 16(3):304. <https://doi.org/10.1016/j.chom.2014.08.006> PMID: 25211073
40. Krause JC, Ekiert DC, Tumphey TM, Smith PB, Wilson IA, Crowe JE. An Insertion Mutation That Distorts Antibody Binding Site Architecture Enhances Function of a Human Antibody. *MBio*. 2011; 2(1):e00345–10. <https://doi.org/10.1128/mBio.00345-10> PMID: 21304166
41. Pejchal R, Doores KJ, Walker LM, Khayat R, Huang PS, Wang SK, et al. A Potent and Broad Neutralizing Antibody Recognizes and Penetrates the HIV Glycan Shield. *Science*. 2011; 334(6059):1097. <https://doi.org/10.1126/science.1213256> PMID: 21998254
42. Wu X, Zhou T, Zhu J, Zhang B, Georgiev I, Wang C, et al. Focused Evolution of HIV-1 Neutralizing Antibodies Revealed by Structures and Deep Sequencing. *Science*. 2011; 333(6049):1593. <https://doi.org/10.1126/science.1207532> PMID: 21835983
43. Mascola JR, Haynes BF. HIV-1 neutralizing antibodies: understanding nature's pathways. *Immunol Rev*. 2013; 254(1):225. <https://doi.org/10.1111/imr.12075> PMID: 23772623
44. Steichen JM, Lin YC, Havenar-Daughton C, Pecetta S, Ozorowski G, Willis JR, et al. A generalized HIV vaccine design strategy for priming of broadly neutralizing antibody responses. *Science*. 2019; 366(6470):eaax4380. <https://doi.org/10.1126/science.aax4380> PMID: 31672916
45. Streisinger G, Okada Y, Emrich J, Newton J, Tsugita A, Terzaghi E, et al. Frameshift Mutations and the Genetic Code. *Cold Spring Harb Symp Quant Biol*. 1966; 31:77. <https://doi.org/10.1101/SQB.1966.031.01.014> PMID: 5237214
46. Golding GB, Gearhart PJ, Glickman BW. Patterns of Somatic Mutations in Immunoglobulin Variable Genes. *Genetics*. 1987; 115(1):169. <https://doi.org/10.1093/genetics/115.1.169> PMID: 3557109
47. Murugan A, Mora T, Walczak AM, Callan CG Jr. Statistical inference of the generation probability of T-cell receptors from sequence repertoires. *Proc Natl Acad Sci*. 2012; 109(40):16161. <https://doi.org/10.1073/pnas.1212755109> PMID: 22988065
48. Ye J, Ma N, Madden TL, Ostell JM. IgBLAST: an immunoglobulin variable domain sequence analysis tool. *Nucleic Acids Res*. 2013; 41(W1):W34–W40. <https://doi.org/10.1093/nar/gkt382> PMID: 23671333
49. Hwang JK, Wang C, Du Z, Meyers RM, Kepler TB, Neuberger D, et al. Sequence intrinsic somatic mutation mechanisms contribute to affinity maturation of VRC01-class HIV-1 broadly neutralizing antibodies. *Proc Natl Acad Sci*. 2017; 114(32):8614. <https://doi.org/10.1073/pnas.1709203114> PMID: 28747530
50. Giudicelli V, Duroux P, Ginestoux C, Folch G, Jabado-Michaloud J, Chaume D, et al. IMGT/LIGM-DB, the IMGT comprehensive database of immunoglobulin and T cell receptor nucleotide sequences. *Nucleic Acids Res*. 2006; 34(90001):D781–D784. <https://doi.org/10.1093/nar/gkj088> PMID: 16381979
51. Saini J, Hershberg U. B cell Variable genes have evolved their codon usage to focus the targeted patterns of somatic mutation on the complementarity determining regions. *Mol Immunol*. 2015; 65(1):157. <https://doi.org/10.1016/j.molimm.2015.01.001> PMID: 25660968
52. Glass DR, Tsai AG, Oliveria JP, Hartmann FJ, Kimmey SC, Calderon AA, et al. An Integrated Multi-omic Single-Cell Atlas of Human B Cell Identity. *Immunity*. 2020; 53(1):217–232.e5. <https://doi.org/10.1016/j.immuni.2020.06.013> PMID: 32668225
53. Horns F, Dekker CL, Quake SR. Memory B Cell Activation, Broad Anti-influenza Antibodies, and Bystander Activation Revealed by Single-Cell Transcriptomics. *Cell Rep*. 2020; 30(3):905–913.e6. <https://doi.org/10.1016/j.celrep.2019.12.063> PMID: 31968262

54. Sok D, Burton DR. Recent progress in broadly neutralizing antibodies to HIV. *Nat Immunol.* 2018; 19(11):1179. <https://doi.org/10.1038/s41590-018-0235-7> PMID: 30333615
55. Vander Heiden JA, Yaari G, Uduman M, Stern JNH, O'Connor KC, Hafler DA, et al. pRESTO: a toolkit for processing high-throughput sequencing raw reads of lymphocyte receptor repertoires. *Bioinformatics.* 2014; 30(13):1930. <https://doi.org/10.1093/bioinformatics/btu138> PMID: 24618469
56. Durbin R, Eddy SR, Krogh A, Mitchison G. *Biological Sequence Analysis.* Cambridge University Press; 1998.
57. Dempster AP, Laird NM, Rubin DB. Maximum Likelihood from Incomplete Data via the EM Algorithm. *J R Stat Soc Ser B.* 1977; 39(1):1.
58. McLachlan GJ, Krishnan T. *The EM Algorithm and Extensions.* Wiley; 2008.
59. Parikh N, Boyd S. Proximal Algorithms. *Found Trends Optim.* 2014; 1(3):127. <https://doi.org/10.1561/2400000003>
60. Duchi J, Shalev-Shwartz S, Singer Y, Chandra T. Efficient projections onto the ℓ_1 -ball for learning in high dimensions. In: *ICML'08 Proc. 25th Int. Conf. Mach. Learn.* New York, New York, USA: ACM Press; 2008. p. 272.
61. Kluge T. C++ cubic spline interpolation; 2015. Available from: <https://kluge.in-chemnitz.de/opensource/spline/> <https://github.com/ttk592/spline>.