



Peer Community Journal

Section: Mathematical & Computational Biology

Research article

Published
2025-04-30

Cite as
Cyprien Borée, Mathieu Giraud and Mikaël Salson
(2025) *Alignment-free detection and seed-based identification of multi-loci V(D)J recombinations in Vidjil-algo*, Peer Community Journal, 5: e44.

Correspondence
contact@vidjil.org

Peer-review
Peer reviewed and recommended by
PCI Mathematical & Computational Biology,
<https://doi.org/10.24072/pci.mcb.100268>



This article is licensed under the Creative Commons Attribution 4.0 License.

Alignment-free detection and seed-based identification of multi-loci V(D)J recombinations in Vidjil-algo

Cyprien Borée¹, Mathieu Giraud ¹, and Mikaël Salson ¹

Volume 5 (2025), article e44

<https://doi.org/10.24072/pcjournal.547>

Abstract

The diversity of the immune repertoire is grounded on V(D)J recombinations in several loci. Many algorithms and software detect and designate these recombinations in high-throughput sequencing data. To improve their efficiency, we propose a multi-loci seed identification through an Aho-Corasick like automaton as well as a seed-based gene filtration. These algorithms were implemented into Vidjil-algo, used routinely by several labs for the analysis of hematologic malignancies. We benchmark the results of Vidjil-algo and of MiXCR on five datasets, evaluating the specificity and sensitivity of the detection, as well as the adequation of the designation to manually curated sequences. Compared to the previous algorithms, the new algorithms implemented in Vidjil-algo bring speedups between 3× and 30×, with a smaller memory footprint and without quality loss in results. They enable to precisely annotate in a few minutes millions of sequences coming from V(D)J recombinations, including incomplete V(D)J-like recombinations, improving our knowledge on immune repertoires.

¹Univ. Lille, CNRS, Centrale Lille, UMR 9189 CRISTAL, F-59000 Lille, France



1. Introduction

V(D)J recombinations are genetic events occurring in immature immunologic cells, the lymphoblasts. These recombinations are an important factor of the diversity of the receptors on B- and T-cells (Tonegawa, 1983). A V(D)J recombination is the result of a random process combining a V gene, possibly a D gene, and a J gene on the genome, building either VDJ or a VJ recombination. The recombination involves enzymes called recombination-activating genes (RAGs). At the junction of these V, D, and J gene segments, nucleotides can be removed and other random ones can be added (see Figure 1), improving again the diversity. Depending on the cell type, these recombinations occur at several loci. Human B-cell receptors have an heavy chain (IGH) and a light chain λ or κ (denoted here IGL and IGK), whereas T-cell receptors either have γ and δ chains (TRG and TRD) or α and β chains (TRA and TRB).

T- and B-cell receptors are crucial for recognizing antigens, which is a key part of the adaptive immune response. Identifying and possibly quantifying V(D)J recombinations helps thus describe the immune response more efficiently. Moreover, a V(D)J recombination can be seen as an identifier – that may be unique – of a *clonal population*, or shortly *clone*, that is a cell population coming from a same lymphoblast. This identifying sequence is called a *clonotype*, but there may be several clonotypes of the same clone (at different loci or alleles) or several different clones with the same clonotype (that may differ elsewhere). In hemato-oncology, V(D)J recombinations are thus used to identify and track clonotypes along time in blood cancers (Cavé et al., 1998). Library preparation for these studies was recently standardized by the EuroClonality-NGS working group (Brüggemann et al., 2019; Langlois de Septenville et al., 2022; Villarese et al., 2022).

On the software side, since the 1980s, after the pioneering work by the Universität zu Köln with DNAPLOT, many tools for the in-depth analysis of V(D)J recombinations were developed by IMGT (Giudicelli et al., 1998; Lefranc, 2011). In the 2010s, new methods and software were proposed to analyze these V(D)J recombinations in high-throughput sequencing data with millions of sequences, such as (Arnaout et al., 2011), IgBlast (Ye et al., 2013), Decombinator (Thomas et al., 2013), miTCR (Bolotin et al., 2013), TCRKlass (Yang et al., 2014), Vidjil (Giraud et al., 2014), MiXCR (Bolotin et al., 2015), IMSEQ (Kuchenbecker et al., 2015), Partis (Ralph and Iv, 2016), IgReC (Shlemov et al., 2017), and IGoR (Marcou et al., 2018). Afzal et al. (2019) did a comparison of several of those software tools. These Adaptive Immune Receptor Repertoire (AIRR-Seq) methods and software, also called Repertoire Sequencing (RepSeq) (Benichou et al., 2012), developed and invented text algorithmic methods adapted to the specificity of V(D)J recombinations.

We call here:

- *detection*, the process of identifying the presence of a V(D)J recombination within a given DNA sequence and determining the related chain (the locus);
- *designation*, the process of determining the specific germline V, (D), and J genes that have undergone recombination, as well as identifying nucleotide deletions, or insertions that may have occurred at the junction between those genes, and possibly detecting the mutations in the whole sequence;
- *clusterization*, the process of gathering equal or similar recombinations, according to some criteria, into clonotypes.



Figure 1 – A VDJ recombination. After removal of a few nucleotides at the junction of V, D, and J, genes, AATA was inserted between gene segments V and D and GCT was inserted between gene segments D and J.

Table 1 – Regular and incomplete/irregular human V(D)J and V(D)J-like recombinations analyzed by `vidjil-algo`. There are seven loci with at least 16 recombination possibilities (called *recombination systems*). For example, on the 14q11.2 TRD locus, apart from the regular recombination, there are at least three known systems of irregular recombinations (TRD+) as well as two systems involving the very close TRA locus (TRA+D). These recombination systems are formally described in the JSON file [homo-sapiens.g](#), and the corresponding 5' and 3' genes (V and J genes for the regular recombinations) are indexed in a Aho-Corasick automaton. Note also that, for short genes (J genes, and D genes involved in incomplete recombination systems), both the genes and 60 bp of their downstream/upstream region are indexed (Duez et al., 2016).

locus	regular recombinations	incomplete/irregular recombinations
14q1.12	TRA Va-Ja	
7q34	TRB Vb-(Db)-Jb	TRB+ Db-Jb
14q11.2	TRD Vd-(Dd)-Jd	TRD+ Vd-Dd3, Dd2-(Dd)-Jd, Dd2-Dd3 TRA+D Vd-(Dd)-Ja, Dd-Ja
7p14	TRG Vg-Jg	
14q32.33	IGH Vh-(Dh)-Jh	IGH+ Dh-Jh
22q11.2	IGL VI-JI	
2p11.2	IGK Vk-Jk	IGK+ Vk-KDE, INTRON-KDE

Most V(D)J analysis software detect and designate V(D)J recombinations at the same time, for each of the input sequences, and some of them cluster sequences thereafter. Vidjil-algo rather cluster sequences *before* the full designation (Giraud et al., 2014). Indeed, for many applications, the designation for each sequence is not useful and more efficient alignment-free approaches, including *k*-mers indexing, can cluster the sequences after detection.

A mature B- or T- cell needs only two V(D)J recombinations, on only one allele. However, V(D)J recombinations may occur at seven different loci (IGH, IGL, IGK, TRG, TRD, TRA, TRB, see Table 1). One may find unproductive recombinations on the other allele, or even in the TR loci in B-cells. Some of these non-productive recombinations may be *incomplete* or *irregular*, such as D-J recombinations on the IGH locus or V-KDE on the IGK, the RAGs enzymes handling the KDE sequence as a J gene. These bi-allelic, incomplete, or irregular recombination are even more frequent in pathological samples, and are also tracked in hemato-oncology studies (Brüggemann et al., 2019). Altogether, one frequently studies datasets with V(D)J or V(D)J-like recombinations on different loci. Table 1 lists $\ell = 16$ different known recombination systems. Some AIRR-Seq/RepSeq software allow to analyze several systems. The previous Vidjil-algo algorithm was able to detect recombinations in time $O(\ell n)$, where n is the length of the input sequence and ℓ the number of recombination systems.

We propose here a new algorithm able to process at once all recombinations systems, detect any V(D)J or V(D)J-like recombinations in time $O(\ell' n)$, where $\ell' \leq \ell$ is an average number of gene labels per position, with usually $\ell' \ll \ell$ (Section 2). We also propose a filtering algorithm to improve the V(D)J designation, in time $O(M' n)$ instead of the previous $O(M n)$, where M is the total size of genes and usually $M' \ll M$ (Section 3). These algorithms were implemented in `vidjil-algo`: We finally report benchmarks, both on simulated and real datasets, on the quality and the speed of detection and designation (Section 4)¹. The new algorithms bring speedups between 3× and 30×, with a smaller memory footprint and without quality loss in results.

¹*Transparency on Benchmarks.* The preprint version 2 on which the recommendation is based (<https://doi.org/10.24072/pci.mcb.100268>) includes benchmarks in Sections 4.3 and 4.4, originally conducted using the development version (commit 7623483f3) of `vidjil-algo`. The final data reported here corresponds to the stable release 2025.02 of `vidjil-algo`. The merge of the development branch into the release branch involved minor updates, particularly related to seed adjustments. To ensure long-term reproducibility, all benchmarks have been re-run using the stable release. As a result, small variations were observed, and Sections 4.3 and 4.4 have been updated accordingly.




Figure 2 – Aho-Corasick automaton for $P(ACAC, \#-\#) \cup P(CCAC, \####)$, that is for the projected seeds A-A, C-C, and CCAC. The gray arrows correspond to the failure transitions in the Aho-Corasick automaton. They can be seen as ϵ -transitions and are removed in a pre-processing step, without changing the number of states. The accepting states have been labeled with V+ for the first seed and with J+ for the second one. Thus the state CCAC is labeled with J+ but it is also labeled with V+ as its failure function points to an accepting state labeled with V+.

Input: a reduced label sequence $a = a_1 a_2 \dots a_n$
 $\delta \leftarrow 0$; $\delta_{\max} \leftarrow 0$
 $i \leftarrow 0$; $j \leftarrow 0$

For each q from 1 to n
 Invariant: $\delta = |a[1, q - 1]|_V - |a[1, q - 1]|_J$
 if $a_t = V$, then $\delta \leftarrow \delta + 1$
 if $a_t = J$, then $\delta \leftarrow \delta - 1$
 if $\delta > \delta_{\max}$, then $\delta_{\max} \leftarrow \delta$ and $i \leftarrow q$
 if $\delta = \delta_{\max}$, then $j \leftarrow q$

End for
 Return i and j




Figure 3 – (Left) $O(n)$ time search of the (i, j) plateau reached around the V-J recombination zone. The actual implementation (affectanalyser.hpp) uses bitsets to check the values of a_t , and also computes, in the same linear time, the values $|a[1, i]|_V$, $|a[1, i]|_J$, $|a[j, n]|_V$ et $|a[j, n]|_J$ for filters and the statistical evaluation. (Right) On this reduced label sequence, the maximal values of δ are $\delta(11) = \dots = \delta(15) = 4$. Thus here $i = 11$ and $j = 15$.

The *accepting states*, that is here the end of seed occurrences, are labeled with the list of the labels $T(g)$ of the genes g occurring at that point – there can be one or several such genes. Note that, contrarily to a lookup table index, the Aho-Corasick automaton may store words with different lengths, representing different spaced seeds (possibly with different weights) according to the recombination system, the gene V/J type, or even an individual gene.

2.3. Querying

Querying a sequence s simply means traversing it while following the transitions in the automaton, in $O(|s|)$ time. All accepting states encountered indicate the end of at least one spaced seed that occurs in the sequence. The output is a label sequence, that is a sequence of list of gene labels such as $Aff(s, \mathcal{G}) = -J_G^+ - V_H^+ V_H^+ V_H^+ J_H^+ - - V_H^+ V_H^+ - V_B^- - - J_H^+ J_H^+ J_H^+ -$. As there may be several gene labels per state, the label sequence can have a maximum size of $O(nL)$, where $L \leq \ell$ is the maximum number of gene labels per state. L is considered constant (on the usual germlines, $L \leq 6$). Calling ℓ' the average number of gene labels per state, the label sequence has an average size of $O(n\ell')$. On the usual germlines, $\ell' = 1.14$ gene labels per accepting state.

One can then focus on a reduced label sequence $Aff(s, \mathcal{G})_{LOC}^\pm$ focusing on V and J of one locus LOC and one strand. For example, focusing on the labels V_H^+ and J_H^+ , we consider a reduced label sequence $Aff(s, \mathcal{G})_{IGH}^+ = ---V V V J --V V-----J J J-$.

Given a gene label t , we denote by $|s|_t = |\{P(g, u) \cap Fact(s) \text{ such that } T(g) = t\}|$, the number of t in the label sequence $Aff(s, \mathcal{G})$, with the seed u being used for t . On the same example, $|s|_{V_H^+} = 5$.

2.4. Locus estimation

The label sequence is analyzed according to the two most probable gene labels² (see *p-value estimation* section below). These two gene labels may represent complete or incomplete V(D)J recombinations (Table 1), but unexpected recombinations can also be detected and are tagged as such.

2.5. Recombination detection

Given these two most probable gene labels, the algorithm detects a 5' zone with seed occurrences from a given label followed by a 3' zone with seed occurrences of another label, but allowing other labels, such as a few (random) J_H^+ or even V_B^- in a significant V_H^+ zone.

Let $a = Aff(s, \mathcal{G})_{LOC}^\pm$ a reduced label sequence. We look for positions t such as $\delta(t) = |a[1, t]|_V - |a[1, t]|_J$ is maximal (many V and few J at the left). This is equivalent to maximize $\delta'(t) = |a[t, n]|_J - |a[t, n]|_V$ (many J and few V at the right), because, for every t , $\delta(t) - \delta'(t) = |a|_V - |a|_J$ is constant. The algorithm described on Figure 3 computes, in linear time, both positions $i \leq j$ that are the first and the last to maximize δ , allowing to detect in the label sequence the 5' and 3' zones.

2.6. p-value estimation

To estimate the significance of the (i, j) zone split and exclude chimeric sequences such as VVVV--JJJ--VVV--JJ, the first check is that the 5' zone has significantly more 5' seeds than the 3' zone, that is $|a[1, i]|_V \geq \tau \cdot |a[j, n]|_V$ with $\tau = 2$, as well as the symmetrical check for the 3' zone. The p-value of a recombination is then estimated as follows. We call p'_V the probability to observe as many labels V in a random label sequence, and estimate $p'_V = B(p_V, |a[1, i]|_V, i)$, where $B(p, k, n) = \sum_{k \leq t \leq n} \binom{n}{t} p^t (1-p)^{(n-t)}$ is the cumulated probability to have an event of probability p at least k times out of n in a Bernoulli schema. This is a very simple model, as the occurrences of seeds are actually not independent.

The probability p to have a label V on fg random seed of weight k (its number of match symbols) is estimated as $p = N_V/4^k$, where N_V is the number of seeds V stored in the index.

²The actual implementation, in $O(\ell^* n)$ time, builds a bitset for each of the $\ell^* \leq 4\ell$ gene labels for at least one seed in the read. The bitsets of the two most probable gene labels are later reused for the recombination detection. The step could however be improved in time $O(\ell' n)$, that is the size of the label sequence, by counting at once the gene labels, then by building the bitsets only for the two most probable ones.




Figure 4 – V(D)J designation by dynamic programming, declared in [segment .h](#). See Jones and Pevzner (2004, chapter 6) for an introduction on dynamic programming methods to compare sequences. Grayed out triangles show parts excluded from the computation. a) Search of the best alignments between the read with a V gene and a J gene, in $O(Mn)$ time, where M is the total size of indexed genes and n the size of the read. Candidate V and J segments are independently predicted. b) When the best alignments make the candidate V and J segments overlap on $o \leq n$ positions, the best split point is found by another search in time $O(mo)$, where $m \ll M$ is the total size of considered V and J genes. c) In the case of a VDJ-like recombination, the central segment is predicted by a local alignment in time $O(M'(d + 2\delta))$, where M' is the total size of indexed D genes and $d + 2\delta \leq n$ is the size of the zone where the D segment is searched. Overlaps between V and D or between D and J candidate segments are handled as previously. The algorithm finally runs in time $O((M + M')n)$.

We similarly define $p'_j = B(p_j, |a[j, n]|_J, n - j + 1)$, and roughly estimate the p-value of a V-J recombination as $p'_V + p'_J$. As a multiple testing correction, this p-value is multiplied by the number of processed sequences, giving an E-value.

Altogether, when this E-value is below a given threshold, a V(D)J recombination has been *detected*. The middle of the recombination zone can then be estimated at around $(j + i + k - 1)/2$. The next section details how we precisely *designate* such a recombination.

3. Fast V(D)J designation through seed-based heuristics

Designating a V(D)J recombination requires to compare the sequence against all V(D)J germline genes from the detected locus. Precisely aligning a gene against V, (D), and J germline genes can be done with dynamic programming techniques (Figure 4). This alignment is done in time $O(Mn)$, where M is the total length of the considered V and J genes. As V genes are about 300 bp in length, this is time consuming. Banded alignments (Chao et al., 1992) bring some improvements, however due to the deletions that occur at the end of V genes or the start of J genes, the constraint on the alignment is imposed on a single part on the gene to prevent restrictions on the number of deletions (Figure 4a and c). However, aligning *all* the genes of a given locus to a given sequence is still very long, in particular for some locus such as the B-cell heavy chain (IGH) with about 350 genes and alleles.

3.1. Selecting candidate genes with seed-based heuristics

Several V(D)J designation methods use seed-based heuristics (eg. (Bolotin et al., 2015; Thomas et al., 2013; Ye et al., 2013)). We propose here to use the previous filtering phase to speed-up the designation phase. The V(D)J detection heuristic presented in the previous section, giving an information on the gene label in sequences, is extended to identify genes which have seeds

occurring in any sequence, in order to determine against which genes the sequence s will be aligned.

For each gene g , the accepting states in the Aho-Corasick automaton are now marked with the gene identifier $I(g)$ along with the gene label $T(g)$. Computing, still in linear time, the labels on the gene identifiers, we consider $C_G(s) = \{g \in \mathcal{G} \mid |s|_{I(g)} > 0\}$ the set of genes, having at least one seed in the sequence s , and $n_{max} = \max_{g \in \mathcal{G}} |s|_{I(g)}$ the maximal number of seeds from s that matched on a gene. We want to align s against all the genes whose number of matched seeds is *close* to n_{max} . This proximity is determined by assuming that the number of matched seeds follows a binomial distribution. Thus we compute a confidence interval for n_{max} (with a confidence interval of 99.9% by default) which gives us a range $[n_a, n_b]$. We finally consider the set of candidate genes $C_G^*(s) = \{g \in \mathcal{G} \mid |s|_{I(g)} \geq n_a\}$: All the genes from \mathcal{G} which have at least n_a seeds matching on s are aligned against s using dynamic programming as on the Figure 4. The designation algorithm runs in $O(M'n)$, where $M' \leq M$ is the total size of these candidate genes. For example, on the IGH germlines, a typical value of $n_{max} = 200$ on sequences of length about 300 gives $[n_a, n_b] = [174, 216]$, and there are usually less than 10 genes/alleles (out of 350+) matching at least 174 seeds, thus $M' \approx 10 \times 200 \ll M \approx 350 \times 200$.

4. Evaluation and results

4.1. Datasets

Five datasets were used to benchmark the detection and the designation of V(D)J recombinations.

Evaluation of the specificity of the detection.

- A. 10^6 random DNA sequences, generated with %GC ratios and sequence length similar to the V(D)J IGH germline genes, in which no V(D)J recombination should be found.

Evaluation of the sensitivity of the detection.

- B. $2.3 \cdot 10^6$ synthetic sequences on all loci, both for regular and “incomplete” recombinations. For each possible combination of V, D, and J genes, 10 sequences were generated, by taking the full gene lengths. Insertions and deletions at the junctions were added according to a normal distribution of 5 ± 5 , and substitutions to the whole sequence on 2% of the nucleotides to take into account sequencing artifacts but also individual variations. We also generated datasets with 5% and 10% substitutions (resp. $B_{5\%}$ and $B_{10\%}$).
- C. TRB simulated sequences from the benchmark of Afzal et al. (2019). We focus on their datasets with more than a single clone and with some errors (low, .1%, and medium, 1%). Each dataset is made of 1M sequences of 250bp. For each error rate, we average the results obtained for the datasets with varying level of clonal populations, as they were very similar.

Evaluation of the correctness of the V(D)J designation.

- D. 1,351 sequences from LIGM-DB (Giudicelli et al., 2006), focusing on the two most represented loci in LIGM-DB, IGH and IGK.
- E. 301 sequences from patient data with curated VDJ designations (Salson et al., 2016).

4.2. Software

Methods described here were implemented in C++ in Vidjil-algo version 2025.02 and compared to Vidjil-algo 2018.02. We benchmarked against MiXCR (versions 3.0.13 and 4.4.1) (Bolotin et al., 2015). MiXCR is widely used, and, although it is not open-source, its code is available. After systematic comparison between several V(D)J analysis tools, MiXCR was assessed by Afzal et al. (2019) as the most balanced generic tools in terms of flexibility and accuracy. Note that




Figure 5 – Detection results of MiXCR and Vidjil-algo on random sequences (dataset A)

MiXCR and Vidjil-algo were launched on the same germline sequences coming from the same IMGT GENE-DB version. Details on germline sequences and used seeds are shown in Supplementary File 1, Table 1. The four programs were launched on one thread on a server with 2 Intel Xeon Gold 6130 processors (2.10 GHz, 32 MB cache) and 128 GB RAM. The benchmark is fully reproducible (from the retrieval of the data to the production of the paper’s figures) using Snakemake (Köster and Rahmann, 2012).

4.3. V(D)J detection

4.3.1. Specificity and sensitivity. We restrict MiXCR to launch the smallest analysis it can do – performing a V(D)J designation on each sequence. Thus the results they produce are much more detailed than what Vidjil-algo provides. On the opposite, Vidjil-algo first tries to identify a V(D)J recombination, then determines an identifier for this V(D)J recombination and clusters all sequences sharing the same identifier into clonotypes. Then, only the 100 most abundant clonotypes are designated (which is not discussed in this section).

On random sequences (dataset A, see Figure 5), MiXCR3 detects hundreds of spurious V(D)J recombinations among the million random sequences. However we note a large improvement in specificity, with a two-fold improvement between MiXCR 3 and MiXCR 4. Vidjil-2025.02 reports 7 spurious recombinations, while Vidjil-2018.02 was the only one not to report any.

Sensitivity is assessed with the dataset B (Figure 6), consisting of randomly generated V(D)J recombinations with 2% substitutions. Our new algorithm is very sensitive to detect those recombinations. It is the most sensitive among all the tested tools. It is the only one to reach 100% sensitivity on all loci. This is a noticeable improvement compared to previous versions of Vidjil-algo. In Supplementary File 1, Figure 1, we show that even with 10% errors the new heuristic shows very good results. On average, on the complete recombinations, it reaches 99.1% detection which is better than Vidjil-2018.02 (66.7%) and MiXCR 3 (94.8%) but less than MiXCR 4 (100.0%). In Supplementary File 1, Figure 3, we show that with 5% errors (20% of them being indels), the detection ratios of Vidjil-2025.02 and MiXCR 4 are mainly unchanged. Vidjil-2018.02 is much more affected by indels. This is due to the spaced seeds used in Vidjil-algo. In Vidjil-2018.02, we couldn’t use seeds of different lengths for V and J, thus the seeds are too long in the J gene to still have a significant number of hits matching in spite of the indels. This shows that our new approach is much more robust to noisy data.

On the TRB benchmark from Afzal et al. (2019) (dataset C, Figure 7) our new algorithm is the most sensitive in the low error condition (.1%) as it is the only one to detect all the recombinations. In the medium error condition (1%), it detects 99.983%. In this second condition, this is slightly less than MiXCR3 (99.996%) but more than MiXCR 4 (99.94%). On IGH and IGK recombinations from LIGM-DB (dataset D, Figure 8), MiXCR 3 gives again the best results for the detection. However Vidjil-algo’s new heuristic improves the former one, and for IGH, is now closer to MiXCR 4 results. The improvement is similar on IGK. However in this second case, our new heuristics gives better detection results than MiXCR 4. The reason why IGK results are worse is that IGK recombined sequences in LIGM-DB are very short on the J side, with 60% of the J gene being at most 20bp long and even 20% sequences at most 10bp long, whereas germinal IGKJ genes are 38-39bp long. There are 62 sequences that Vidjil-algo did not detect

Table 2 – Detection on MiXCR, and Vidjil-algo on curated V(D)J designations (dataset E)

locus	nb. seq.	mixcr3	mixcr4	vidjil-2018.02	vidjil-2025.02
IGH	94	95	94	92	90
IGK	2	2	2	2	2
IGL	2	2	2	2	2
TRA	1	1	1	1	1
TRB	16	15	15	16	16
TRD	18	18	18	18	18
TRG	31	30	31	30	31
IGH+	23			23	23
IGK+	29			29	29
TRA+D	28			28	26
TRB+	20			19	19
TRD+	31			25	25

and MiXCR 3 did: All of them had an IGKJ gene 12bp long or shorter. With so short sequences, a single mutation may prevent any spaced seed to match the J sequence.

On the dataset E, results are shown in Table 2. Results are very similar between the four programs for complete recombinations (IGH, IGK, IGL, TRA, TRB, TRD, TRG). As expected, only Vidjil-algo detects some incomplete recombinations, as MiXCR does not deal with them.

4.3.2. Speed and memory. On top of that, the new algorithm is much quicker than the former version of Vidjil-algo (at least three times quicker on large enough datasets). For datasets A, B, and C, it is between 15 and 40 times quicker than MiXCR. The difference is less striking on small datasets due to the construction of the Aho-Corasick automaton at each startup. In spite of this data structure, the memory consumption is lower than before because it used lookup tables (that were sparse). MiXCR almost systematically has the highest memory consumption, apart from dataset with a large number of distinct recombinations (IGH on Figure 6).

4.4. V(D)J designation

V(D)J designations predicted by the software, taking the same reference genes, are compared to the reference datasets D and E (Table 3 and Figure 8). Only the names of the V and of the J genes are checked (or D and J genes for incomplete recombinations). As some genes are very similar, designating one gene is very dependant to fine tunings in the scoring of the sequence comparisons. The output of both software could thus be considered correct in some cases even when departing from those reference datasets – and note that the dataset E already contains some alternative acceptable designations.

Both Vidjil-algo and MiXCR designate the same V(D)J recombinations than in the reference datasets in most of the sequences. Designations on the complete recombinations (IGH, TRB, TRD, TRG) are particularly close to the curated dataset E, with more than 90% correct designations, showing that the software do not have any specific difficulty to identify the V and J genes involved (see also Supplementary File 1, Figure 2). Surprisingly, MiXCR 4 with the default parameters does not designate some sequences in the way that MiXCR 3 did. This may come from stricter parameters to improve specificity, as it was also shown on random sequences in Figure 5. Note also that in some cases (see IGK in Figure 8, dataset D), the detection step of Vidjil-algo can assign borderline sequences to an incorrect locus. During the designation step such an error would be fixed.

Anyway, while staying very specific, the new Vidjil-algo designation is now similar to the results obtained with MiXCR 4. Moreover, as expected, only Vidjil is capable of handling incomplete recombinations (Table 3), with 82% correct designations on incomplete recombinations.

Table 3 – Correct designation on V(D)J recombinations on manually curated V(D)J sequences (dataset E) with MiXCR and Vidjil-algo. *nb. seq.* is the number of sequences in the dataset with the given locus.

locus	nb. seq.	mixcr3	mixcr4	vidjil-2018.02	vidjil-2025.02
IGH	94	89	87	86	87
IGK	2	2	2	2	2
IGL	2	2	2	2	2
TRA	1	1	1	1	1
TRB	16	14	15	14	15
TRD	18	15	14	15	18
TRG	31	28	28	28	28
IGH+	23	1		21	20
IGK+	29			19	18
TRA+D	28	2		21	23
TRB+	20	2		19	19
TRD+	31			26	27

On the E dataset, our new heuristic to avoid the alignment against many genes improves time efficiency by over 10 times. Thus, Vidjil time consumption of designation with our new heuristic becomes comparable to MiXCR, while Vidjil-algo didn't optimize the alignment by itself (by using SIMD for instance). Note that in the classic usage, and due to its approach, Vidjil-algo can limit the designation to the 100 most abundant clonotypes.

5. Discussion and perspectives

Studying immune repertoire by high-throughput sequencing for immunological or onco-hematological applications requires adapted methods. We introduced a seed-based alignment-free algorithm, based on an Aho-Corasick automaton, to detect in a single pass, in almost linear time ($O(\ell'n)$), V(D)J recombinations coming from different loci, as well as a filtering algorithm improving the designation of V(D)J gene segments from a recombination. Both algorithms are fast and sensitive, and come with a statistical evaluation of their results, including on irregular or incomplete recombinations.

Our solution is another example where alignment-free approaches, here seed-based heuristics, provide pertinent results to analyze huge datasets, using a fraction of the resources required by alignment-based approaches. With this contribution, we almost have a linear processing of sequences. Other improvements on time or memory consumption are still possible but, as shown in our benchmarks, improvements on the quality of the results can only be marginal.

Our new version of Vidjil-algo is hence one of the fastest available programs for analyzing large datasets with billions of immune recombinations, and is moreover released under an open-source licence. The two algorithms provided an up to $5\times$ speed-up compared to the previous Vidjil-algo version, still keeping excellent sensibility and specificity and a low memory footprint. Other software provide more information – notably the V(D)J designation of each sequence – but they are not necessarily needed in several applications. We also show that Vidjil-algo is highly effective to identify and filter sequences that do not exhibit V(D)J recombinations. This is a feature of interest to analyze large sequencing datasets, such as RNA-seq data that contain very few V(D)J recombinations.




Figure 6 – Detection by MiXCR and Vidjil-algo on synthetic V(D)J recombinations on all human loci (dataset B). The X-axis on the time diagrams is logarithmic.




Figure 7 – Detection on Afzal et al. (2019) TRB benchmark (dataset C) by MiXCR and Vidjil-algo.




Figure 8 – Correct detection and designation of V(D)J recombinations on LIGM-DB version LIGMDB_V12 (dataset D) with MiXCR and Vidjil-algo.

Vidjil-algo is already used in reference protocols dealing with sequencing and analyzing immunogenetical data (Langlois de Septenville et al., 2022; Villarese et al., 2022). With the rise of large-scale analysis of public datasets (Edgar et al., 2022), Vidjil-algo could become one of the preferred methods for V(D)J detection on such large-scale analyses. In a second time, MiXCR could be used to obtain detailed informations on the V(D)J recombinations detected by Vidjil-algo.

As the Aho-Corasick automaton stores words projected from different seeds, *seed optimization* could be further studied. The smaller the seed, the more sensitive, but the less specific. Gene repertoires have very different sizes according to the locus, with for example more than 200 kB of sequences on IGH V-J and just a few nucleotides for the TRD+ Dd2/Dd3 (see Supplementary File 1, Figure 1). In the assessed version, shorter seed sizes were selected for J genes, enabling a better recognition. Further work could optimize the seed lengths and weights depending on each recombination system.

More generally, research could include efficient detection and designation of recombined sequences with three or more segments, as well as improving again the statistical evaluation of recombinations.

Acknowledgments

We thank the Mésocentre de Lille for their computing resources, and Inria for the support as well as computing resources. We also thank Nika Abdollahi (Abdollahi, 2021) and the Bonsai team for discussions on alignment-free methods. We finally thank the VidjilNet consortium, users of Vidjil as well as the EuroClonality-NGS consortium. Their feedback helped us improve the algorithm. Preprint version 2 of this article has been peer-reviewed and recommended by Peer Community In Mathematical and Computational Biology (<https://doi.org/10.24072/pci.mcb.100268>); (Pibiri, 2024).

Conflict of interest disclosure

Mathieu Giraud and Mikaël Salson are members of the Scientific and Technical Committee of the not-for-profit VidjilNet consortium. They do not receive any financial compensation from that consortium. The authors of this preprint declare that they have no financial conflict of interest with the content of this article.

Data, scripts, and code

Vidjil 2025.02 is released under the GNU GPLv3, or any other later version, and is available from <https://www.vidjil.org/releases/> and also from the long-term Software Heritage archive: <https://archive.softwareheritage.org/swh:1:snp:ac4da737d1ebb2930d25ca4064fc23ad95170f4c>. The code to get the data and to run the benchmarks reported in this article is available at <https://gitlab.inria.fr/vidjil/data/2025-algo> which is archived on Software Heritage: <https://archive.softwareheritage.org/swh:1:dir:0111ff89845695647674f8688a81f8be5f74c10e>. MiXCR code is available at <https://github.com/milaboratory/mixcr>. However, it is not released under an open-source licence.

Supplementary File 1 is available from <https://hal.science/hal-04361907>.

References

Abdollahi N (2021). *B cell receptor repertoire analysis in clinical context : new approaches for clonal grouping, intra-clonal diversity studies, and repertoire visualization*. <https://hal.science/tel-03482030>. PhD thesis. Sorbonne Université.

- Afzal S, Gil-Farina I, Gabriel R, Ahmad S, Kalle C, Schmidt M, Fronza R (2019). *Systematic comparative study of computational methods for T-cell receptor sequencing data analysis*. *Briefings in Bioinformatics* **20**, 222–234. <https://doi.org/10.1093/bib/bbx111>.
- Aho AV, Corasick MJ (1975). *Efficient string matching: An aid to bibliographic search*. *Communications of the ACM* **18**, 333–340. <https://doi.org/10.1145/360825.360855>.
- Arnaut R, Lee W, Cahill P, Honan T, Sparrow T, Weiland M, Nusbaum C, Rajewsky K, Koralov SB (2011). *High-Resolution Description of Antibody Heavy-Chain Repertoires in Humans*. *PLoS ONE* **6**, e22365. <https://doi.org/10.1371/journal.pone.0022365>.
- Benichou J, Ben-Hamo R, Louzoun Y, Efroni S (2012). *Rep-Seq: uncovering the immunological repertoire through next-generation sequencing*. *Immunology* **135**, 183–91. <https://doi.org/10.1111/j.1365-2567.2011.03527.x>.
- Bolotin DA, Shugay M, Mamedov IZ, Ekaterina V Putintseva MAT, Zvyagin IV, Britanova OV, Chudakov DM (2013). *MiTCR: software for T-cell receptor sequencing data analysis*. *Nature Methods* **10**, 813–814. <https://doi.org/10.1038/nmeth.2555>.
- Bolotin DA, Poslavsky S, Mitrophanov I, Shugay M, Mamedov IZ, Putintseva EV, Chudakov DM (2015). *MiXCR: software for comprehensive adaptive immunity profiling*. *Nature Methods* **12**, 380–381. <https://doi.org/10.1038/nmeth.3364>.
- Brown DG (2008). *Bioinformatics Algorithms: Techniques and Applications*. In: chap. A survey of seeding for sequence alignment, pp. 126–152. <https://doi.org/10.1002/9780470253441.ch6>.
- Brüggemann M, Kotrová M, Knecht H, Bartram J, Boudjogrha M, Bystry V, Fazio G, Froňková E, Giraud M, Grioni A, et al. (2019). *Standardized next-generation sequencing of immunoglobulin and T-cell receptor gene recombinations for MRD marker identification in acute lymphoblastic leukaemia; a EuroClonality-NGS validation study*. *Leukemia* **33**, 2241–2253. <https://doi.org/10.1038/s41375-019-0496-7>.
- Cavé H, Werff Ten Bosch J, Suciú S, Guidal C, Waterkeyn C, Otten J, Bakkus M, Thielemans K, Grandchamp B, Vilmer E, Nelken B, Fournier M, Boutard P, Lebrun E, Méchinaud F, Garand R, Robert A, Dastugue N, Plouvier E, Racadot E, et al. (1998). *Clinical significance of minimal residual disease in childhood acute lymphoblastic leukemia*. *New England Journal of Medicine* **339**, 591–598. <https://doi.org/10.1056/NEJM199808273390904>.
- Chao KM, Pearson WR, Miller W (1992). *Aligning two sequences within a specified diagonal band*. *Bioinformatics* **8**, 481–487. <https://doi.org/10.1093/bioinformatics/8.5.481>.
- Duez M, Giraud M, Herbert R, Rocher T, Salson M, Thonier F (2016). *Vidjil: A web platform for analysis of high-throughput repertoire sequencing*. *PLOS One* **11**, e0166126. <https://doi.org/10.1371/journal.pone.0166126>.
- Edgar RC, Taylor J, Lin V, Altman T, Barbera P, Meleshko D, Lohr D, Novakovsky G, Buchfink B, Al-Shayeb B, et al. (2022). *Petabase-scale sequence alignment catalyses viral discovery*. *Nature* **602**, 142–147. <https://doi.org/10.1038/s41586-021-04332-2>.
- Giraud M, Salson M, Duez M, Villenet C, Quief S, Caillault A, Gardel N, Roumier C, Preudhomme C, Figeac M (2014). *Fast multiclonal clusterization of V(D)J recombinations from high-throughput sequencing*. *BMC Genomics* **15**, 409. <https://doi.org/10.1186/1471-2164-15-409>.
- Giudicelli V, Chaume D, Mennessier G, Althaus HH, Müller W, Bodmer J, Malik A, Lefranc MP (1998). *IMGT, the international ImMunoGeneTics database: a new design for immunogenetics data access*. In: *MEDINFO'98*. IOS Press, pp. 351–355. <https://doi.org/10.3233/978-1-60750-896-0-351>.
- Giudicelli V, Duroux P, Ginestoux C, Folch G, Jabado-Michaloud J, Chaume D, Lefranc MP (2006). *IMGT/LIGM-DB, the IMGT® comprehensive database of immunoglobulin and T cell receptor nucleotide sequences*. *Nucleic acids research* **34**, D781–D784. <https://doi.org/10.1093/nar/gkj088>.

- Jones NC, Pevzner PA (2004). *An introduction to bioinformatics algorithms*. MIT Press.
- Köster J, Rahmann S (2012). *Snakemake—a scalable bioinformatics workflow engine*. *Bioinformatics* **28**, 2520–2522. <https://doi.org/10.1093/bioinformatics/bts480>.
- Kuchenbecker L, Nienen M, Hecht J, Neumann AU, Babel N, Reinert K, Robinson PN (2015). *IMSEQ – a fast and error aware approach to immunogenetic sequence analysis*. *Bioinformatics* **31**, btv309. <https://doi.org/10.1093/bioinformatics/btv309>.
- Langlois de Septenville A, Boudjoghra M, Bravetti C, Armand M, Salson M, Giraud M, Davi F (2022). *Immunoglobulin Gene Mutational Status Assessment by Next Generation Sequencing in Chronic Lymphocytic Leukemia*. In: *Immunogenetics*. Ed. by Anton W. Langerak. Vol. 2453. *Methods in Molecular Biology*. Springer, pp. 153–167. https://doi.org/10.1007/978-1-0716-2115-8_10.
- Lefranc MP (2011). *IMGT, the International ImMunoGeneTics Information System*. *Cold Spring Harbor Protocols* **2011**, pdb.top115. <https://doi.org/10.1101/pdb.top115>.
- Marcou Q, Mora T, Walczak AM (2018). *High-throughput immune repertoire analysis with IGoR*. *Nature communications* **9**, 561. <https://doi.org/10.1038/s41467-018-02832-w>.
- Pibiri G (2024). *An accelerated Vidjil algorithm: up to 30X faster identification of V(D)J recombinations via spaced seeds and Aho-Corasick pattern matching*. *Peer Community in Mathematical and Computational Biology* **100268**. <https://doi.org/10.24072/pci.mcb.100268>.
- Ralph DK, Iv FAM (2016). *Consistency of VDJ Rearrangement and Substitution Parameters Enables Accurate B Cell Receptor Sequence Annotation*. *PLOS Comput Biol* **12**, e1004409. <https://doi.org/10.1371/journal.pcbi.1004409>.
- Salson M, Caillault A, Duez M, Ferret Y, Fievet A, Kotrova M, Thonier F, Villarese P, Wakeman S, Wright G, Giraud M (2016). *A Dataset of Sequences with Manually Curated V(D)J Designations*. <https://hal.science/hal-01331556v1>. Workshop on Immune Repertoire Sequencing : Bioinformatics and Applications in Hematology and Immunology (RepSeq 2016).
- Shlemov A, Bankevich S, Bzikadze A, Turchaninova MA, Safonova Y, Pevzner PA (2017). *Reconstructing antibody repertoires from error-prone immunosequencing reads*. *The Journal of Immunology* **199**, 3369–3380. <https://doi.org/10.4049/jimmunol.1700485>.
- Thomas N, Heather J, Ndifon W, Shawe-Taylor J, Chain B (2013). *Decombinator: a tool for fast, efficient gene assignment in T-cell receptor sequences using a finite state machine*. *Bioinformatics* **29**, 542–550. <https://doi.org/10.1093/bioinformatics/btt004>.
- Tonegawa S (1983). *Somatic generation of antibody diversity*. *Nature* **302**, 575–581. <https://doi.org/10.1038/302575a0>.
- Villarese P, Abdo C, Bertrand M, Thonier F, Giraud M, Salson M, Macintyre E (2022). *One-Step Next-Generation Sequencing of Immunoglobulin and T-Cell Receptor Gene Recombinations for MRD Marker Identification in Acute Lymphoblastic Leukemia*. In: *Immunogenetics. Methods and Protocols*. Ed. by Anton W. Langerak. Vol. 2453. *Methods in Molecular Biology*. Springer, pp. 43–59. https://doi.org/10.1007/978-1-0716-2115-8_3.
- Yang X, Liu D, Lv N, Zhao F, Liu F, Zou J, Chen Y, Xiao X, Wu J, Liu P, Gao J, Hu Y, Shi Y, Liu J, Zhang R, Chen C, Ma J, Gao GF, Zhu B (2014). *TCRklass: A New K-String-Based Algorithm for Human and Mouse TCR Repertoire Characterization*. *Journal of Immunology* **194**, 446–454. <https://doi.org/10.4049/jimmunol.1400711>.
- Ye J, Ma N, Madden TL, Ostell JM (2013). *IgBLAST: an immunoglobulin variable domain sequence analysis tool*. *Nucleic Acids Research* **41**, W34–W40. <https://doi.org/10.1093/nar/gkt382>.