Reconstructing the evolutionary history of a BCR lineage with minimum spanning tree and clonotype abundances

Nika Abdollahi¹, Anne de Septenville², Frédéric Davi^{2,*} and Juliana S. Bernardes^{1,*}

¹Sorbonne Université, CNRS, UMR 7238, Laboratoire de Biologie Computationnelle et Quantitative, Paris, France

²Sorbonne Université, AP-HP, Hôpital Pitié-Salpêtrière, Department of Biological Hematology, Paris, France

*To whom correspondence should be addressed; E-mail: juliana.silva.bernardes@upmc.fr.

Motivation: B cell receptor (BCR) genes exposed to an antigen undergo somatic hypermutations and Darwinian antigen selection, generating a large BCR-antibody diversity. This process, known as B cell affinity maturation, increases antibody affinity, forming a specific B cell lineage that includes the unmutated ancestor and mutated variants. In a B cell lineage, cells with a higher antigen affinity will undergo clonal expansion, while those with a lower affinity will not proliferate and probably be eliminated. Therefore, the abundance of different genotypes provides a valuable perspective on the ongoing evolutionary process. Phylogenetic tree inference is often used to reconstruct B cell lineage trees and represents the evolutionary dynamic of BCR affinity maturation. However, such methods should process B cell population data derived from experimental sampling that might contain different cellular abundances. There are a few phylogenetic methods for reconstructing the evolutionary history of B cell lineages; best-performing solutions are time-demanding and restricted to analyze a reduced number of BCR sequences, while time-efficient methods do not consider cellular abundances. We propose ClonalTree, a low-complexity and accurate approach to reconstruct BCR lineage trees that incorporates genotype abundances into minimum spanning tree (MST) algorithms. Using both simulated and experimental data, we demonstrated that ClonalTree outperforms MST-based algorithms and achieves a similar performance compared to a method that explores tree generating space exhaustively. However, ClonalTree has a lower running time, being more convenient for reconstructing phylogenetic lineage trees from high-throughput BCR sequencing data, mainly in biomedical applications, where a lower computational time is appreciable. It is hundreds to thousands of times faster than exhaustive approaches, enabling the analysis of a large set of sequences within minutes or seconds and without loss of accuracy. The source code is freely available at github.com/julibinho/ClonalTree.

key words: B cell immune repertoire, Lineage tree, phylogenetics

Bcells are essential components of the adaptive immune system. They express a cell surface receptor, the B cell receptor (BCR), recognizing a vast array of antigens. The main components of the BCR are immunoglobulins (IG), which can be secreted in a soluble form as antibodies. IGs are heterodimers composed of two identical heavy chains (IGH) and two identical light chains (IGL); each chain possesses a variable and a constant region. The variable regions are responsible for antigen-binding specificities, whereas the constant regions are associated with cell-signaling components of the BCR and also facilitate interactions with other immune-system molecules. The variable regions are encoded by Variable (V), Diversity (D) in the case of heavy chains, and Joining (J) genes, which are rearranged during lymphopoiesis by a complex genetic mechanism known as V(D)J recombination. During this process, the combinatorial diversity generated by the ligation of these genes into an exon is further enhanced by random deletions and insertion at their joining sites. These mechanisms allow newly-formed naive B cells to express a vast repertoire of distinct BCRs ($> 10^{13}$) (1).

When exposed to an antigen, naive B cells' genes encoding the IG undergo multiple rounds of mutations, e.g., somatic hypermutations (SHM) and Darwinian antigen selection. This stage of the B cell development is known as affinity maturation since it leads to a progressive increase of the IG's affinity for their cognate antigens and occurs in specialized structures of secondary lymphoid organs, the germinal centers. The naive B cell and its variants (generated during affinity maturation) form a B-cell lineage. Natural selection, occurring during affinity maturation, selects for B-cells with higher IG-antigen affinities, which proliferate and undergo clonal expansion, while those with lower affinity will be eliminated (2). Thus, one might find cells with different IG-antigen affinities within a B-cell lineage. The number of unique variant sequences and their respective abundances provides an important perspective on the ongoing evolutionary process and helps elucidate clonal selection. While this model predicts a reduced clonal diversity along with clonal selection, recent studies with complex antigens have shown an increased diversity in IGH variable genes repertoire within germinal centers. Thus affinity maturation is rather characterized by a complex interplay of intraclonal and interclonal diversification (3, 4).

Understanding BCR repertoire evolution is necessary to answer fundamental biological questions

such as clonal selection during antigen challenges, immune senescence, development of efficient vaccines, therapeutic monoclonal antibodies, or further understanding B cell tumor development. Recently, an evolutionary approach has been able to quantify dissimilarity when comparing BCR repertoires of young and aged individuals after influenza vaccination (5). Antibody evolutionary studies have also been used to guide the clonal reconstruction of BCR repertoires (6, 7).

Several methods have been used to analyze BCR phylogenies where clonally related B cells derive from a common ancestor (8). These lineage trees are largely shaped by antigen-driven selection occurring during an immune response. Genetic evolution, such as observed during affinity maturation, is often studied through phylogenetic inference, a well-known methodology that describes the evolution of related DNA or protein sequences in various species. Theoretically, phylogenetic inference methods could reconstruct BCR lineage trees by replacing species with BCR sequences (with different mutations). However, in a phylogenetic tree, the root is usually unknown, the observed sequences are represented usually only in the leaves, and the inner nodes represent the relationships amongst sequences. Conversely, in a BCR lineage tree, the root or the BCR sequence of the naive B cell giving rise to the lineage can be deduced by aligning the IG variable region sequences with sequences stored in a reference database for identifying the germline V, (D) and J genes from which they derive (9). Another distinct point is B cells with different BCR mutations can coexist; therefore, the observed BCR sequences can be leaves or internal nodes in the tree (3, 4). Due to simultaneous divergence, multi-furcations are also common (10). IG sequences are under intense selective pressure, and the neutral evolution assumption is invalid. Moreover, the context dependence of SHM violates the assumption that sites evolve independently and identically. Under these circumstances, conventional phylogenetic tree algorithms seem not appropriate for reconstructing BCR lineage trees. The performance of such methods varies substantially in terms of the tree topology and the ancestral sequence (10, 11).

Some computational tools were explicitly designed to reconstruct BCR lineage trees. Igtree (12) employs the maximum parsimony criterion to find the minimal set of events that could justify the observed sequences. It first constructs a preliminary tree, which only contains observed sequences,

then uses a combined score based mainly on sequence mutations to add internal nodes (unobserved sequences) gradually. IgPhyML combines the maximum likelihood approach (13) with a codon substitution model that uses a Markov process to describe substitutions between codons (14). IgPhyML has modified the codon substitution model to incorporate hot- and cold-spot biases observed in BCR sequences (15).

GCtree (16) employs the maximum parsimony principle and additionally incorporates the cellular abundance of a given genotype in phylogenetic inference. This information is used for ranking parsimonious trees, obtained by dnajpars (17) with the assumption that more abundant parents are more likely to generate mutant descendants. GCtree uses a likelihood function based on the Galton-Watson Branching process (18). It is an accurate method, but its computational complexity is high, especially for a high number of sequences. GLaMST (19) is another method for reconstructing BCR lineage trees. It is a minimum spanning tree (MST)-based algorithm and iteratively grows the lineage tree from the root to leaves by adding minimal edge costs. GLaMST is more time-efficient than GCtree, but it ignores genotype abundance information.

Here we propose ClonalTree, a method to reconstruct BCR lineage trees, combining MST and genotype abundance to infer maximum parsimony trees. ClonalTree starts from the root (the ancestral sequence) and iteratively adds nodes to the tree presenting minimal edge cost and maximum genotype abundance; therefore, it optimizes a multi-objective function rather than a single function based only on edge costs as implemented in GLaMST. Using simulated and experimental data, we demonstrate that ClonalTree outperforms GLaMST and achieves a comparable performance to GCtree. ClonalTree has a lower time complexity and great potential for many applications, particularly in clinical settings where time constraint is important.

0.1 Approach

ClonalTree reconstructs BCR lineage trees based on the minimum spanning tree (MST) and genotype abundances. We start from a formal description of the BCR lineage tree reconstruction problem, BCR sequence distance, and minimum spanning tree algorithms. Next, we describe how we modified

Prim's algorithm (20) to incorporate genotype abundance information. Ultimately, we show how trees can be improved by creating intermediate nodes that describe non-observed sequences or performing editing operations.

0.1.1 Problem statement

Given a set of observed BCR sequences with the same rearrangement event, we look for a minimum-sized directed tree structure (B-cell lineage tree) that might represent the affinity maturation process. Vertices (nodes) represent BCR sequences, and the weight of edges connecting vertices represents the distance between sequences in terms of mutation, insertion, and deletion operations. All observed sequences are reachable from the root, e.g., the naive/unmutated BCR sequence.

0.1.2 Computing BCR sequence distances

The construction of BCR lineage trees requires the calculation distance between clonally-related sequences. Pairwise edit distances such as Levenshtein (21) usually require dynamic programming algorithms that become computationally expensive in time and memory when the number of sequences per lineage tree is large. Another weak point of pairwise distances is that they provide a local landscape without considering mutated regions across all the sequences. ClonalTree performs a multiple sequence alignment to consider evolutionary events such as mutations, insertions, and deletions over all clonally-related sequences. Next, we compute a normalized hamming distance (?) between each pair of sequences. This pairwise distance is then used as the weight edge connecting two sequences.

0.1.3 Minimum spanning Tree

Given a connected, undirected graph (V, E) , where V is a set of vertices, and E is the weight edges, its minimum spanning tree (MST) is a subset of vertices and edges that form a tree (a graph without cycles/loops) so that the sum of the weights of all the edges, the cost, is at minimum. For a given connected graph, several MSTs can exist. All trees have the same cost, but their topologies are different. Most MST construction algorithms are greedy approaches, where edges are sorted according

to their weights and selected according to some criteria. In each step, greedy algorithms make a locally-optimal choice hoping that this choice will lead to a globally optimal solution.

0.1.4 A modified Prim's algorithm

Prim's (20) and Kruskal's (22) are algorithms for finding the minimum spanning tree of a graph. Both are greedy approaches and present low time complexity. However, Prim's algorithm runs faster than Kruskal in dense graphs (23). Therefore, we modified Prim's algorithm to reconstruct BCR lineage trees. We start at the root and add all its neighbors with minimum edge weight to a priority queue. We then iteratively extract from the priority queue the node with the lowest edge weight and highest genotype abundance. If no cycle is formed, the node and the edge will be added to the tree. For each added node, all its neighbors with minimum edge weights are included in the priority queue. We keep on adding nodes and edges until we cover all nodes. In order to decrease the time complexity of the algorithm, we add each node only once at the priority queue. The original Prim's algorithm has only one objective function, which minimizes the sum of edge weights (cost). Here we include a second objective function to maximize genotype abundance. If a set of edges have the same weight, we will choose the one that connects nodes with high abundance. Prim's algorithm has a time complexity of $O(|V|^2)$ in the worst case, but can be improved up to $O(|E| + \log |V|)$ when using data structures based on Fibonacci heaps (24). Figure 1 shows a simple example of the tree construction process.

0.1.5 Editing the reconstructed lineage tree

A greedy algorithm makes the optimal choice at each step, attempting to find the optimal way to solve the optimization problem. It never reconsiders its choices, while optimal algorithms always find the best solution. A way to review decisions and improve the ClonalTree algorithm is to edit the obtained lineage tree. Thus, we implemented two strategies: add unobserved intermediate nodes to the tree and detach/reattach sub-trees.

Unobserved internal nodes might represent unobserved sequences that were not sampled or disappeared during the affinity maturation. In those cases, the evolutionary relationships were also lost. One way to recover them is to analyze the reconstructed tree to identify common ancestors not yet

represented. This process is similar to building a phylogenetic tree among species, where unobserved internal nodes represent common ancestors of descendants. However, only leaves' nodes are observed in a classical phylogenetic tree, and all internal nodes are unobserved, while in a BCR lineage tree, internal nodes can be observed or unobserved. We add an unobserved internal node when we detect a missing common ancestral in the tree. It generally happens when we observe a distance between sister nodes that is smaller or equal to the distance for their parent, see nodes {d, e} in [Figure S1-A](#). To add unobserved internal nodes, we traverse the tree in a pre-order manner, and for each pair of sister nodes, we verify if a common ancestor is missing. If it is the case, we add an unobserved internal node in the tree, connect it to the observed nodes by direct edges [Figure S1-B](#), and update edge weights [Figure S1-C](#). For this later, let d_{pm} and d_{pn} be the edge weights connecting a pair of sister nodes m and n to their parent p , respectively. Let d_{mn} be the distance between the sister nodes and i an unobserved internal added to the tree. The edge weights are updated as follows:

$$\begin{aligned} d_{pi} &= \max(d_{pm} - d_{mn}, d_{pn} - d_{mn}, 1) \\ d_{im} &= d_{pm} - d_{pi} \\ d_{in} &= d_{pn} - d_{pi} \end{aligned} \tag{1}$$

We can detach a sub-tree from an internal node by removing its edge and reattaching it to another internal node or leaf. We perform this editing operation to reduce the depth of the lineage tree by keeping the overall cost. The motivation here is to try to find the most parsimonious tree. We consider all branching nodes (i.e., nodes with more than one descendent) for this edition operation. Then, we try to detach and reattach each node (under edition condition) to any other node in the lineage tree. If this operation reduces the tree depth, we accept it and examine the resulting lineage tree again for additional edition operations that may further reduce the tree depth. We repeat this process until no editing operation can reduce the depth of the tree (see an illustration in [Figure S2](#)).

Figure 1: **ClonalTree construction example.** We start with a complete connected weighted graph (a) where nodes represent BCR sequences, edge weights their distances and node weights the genotype abundances. Then, we first place the ancestral sequence or root (b), and iteratively add nodes to the tree with lowest edge weights and highest genotype abundances (c,d), when edges have the same weight (e) we choose those connected to the node with higher abundance (f), we repeat until all nodes were added to the tree (g), the final tree is shown in (h).

0.2 Methods

In the following, we describe how data sets were obtained or produced, how the performance of BCR lineage tree algorithms was evaluated, and the tools considered in inter-comparisons.

0.2.1 Data sets

We used two types of data sets to measure the ClonalTree performance and compare it with existing algorithms: simulated and experimental. Simulated data were produced by GCtree simulator (16), while one of two experimental data sets was created during chronic lymphocytic leukemia (CLL) routine diagnostic procedures at the Pitié Salpêtrière hospital (Paris-France), and the second one is a public data set (25).

Simulated lineage trees In order to create simulated lineage trees, we used the B cell lineage simulator provided by GCtree. The simulator produces a B cell lineage by randomly selecting IGHV, IGHD, and IGHJ germline genes from the IMGT database. Then, nucleotide(s) can be added to or removed from the junction region: IGHV-IGHD and IGHD-IGHJ. Next, it performs a branching process, and point mutations can be included in the descendants. Somatic hypermutations are simulated by a sequence-dependent process, where mutations are preferentially introduced within specific hot- and cold-spot motifs (26). We kept simulator default parameters and generated 92 artificial lineage trees. The sizes of simulated trees ranged from 6 to 99 nodes, the number of observed sequences between 20 and 200, the degree of roots vary from 1 to 42, and depth trees from 2 to 7, see Table 1.

Table 1: **Characteristics of simulated lineage trees**

	min	mean	max	std
Tree depth	2	4	7	1
Root degree	1	11	42	9
number of nodes	6	34	99	21
number of leaves	5	25	82	16
number of observed sequences	30	81	200	48

Experimental data We used a public data set from a previously reported experiment (25), where the authors combined multiphoton microscopy and single-cell mRNA sequencing to obtain IGH sequences extracted from germinal B cells of a lineage sorted from the mouse germinal center. The data set, labeled as TAS-42, contains 65 IGHV sequences and 42 genotype abundances; it is available at <https://github.com/matsengrp/GCtree/tree/master/example>. The ancestor IGHV gene was inferred with Partis (27), and the final data set contains 66 sequences: 65 from the original data set plus one representing the reconstructed germline sequence. The second experimental data set was generated by sampling sequences from the most abundant clone of a CLL monoclonal repertoire. It contains 20 different genotypes originating from a clonal IGHV-IGHD-IGHJ rearrangement, totaling 3406 sequences. These sequences were annotated with IGHV4-34*01/IGHD3-3*01/IGHJ4*02 genes using IMGT/HighV-QUEST (28). We label this data set as CLL-20. To predict the hypothetical naive sequence or the root of this lineage tree, we considered the germline sequences of the corresponding IGHV and IGHJ genes provided by IMGT/HighV-QUEST. For the junction, we took it from the sequence with the lowest number of mutations on the IGHV gene when compared to the germline determined by IMGT/HighV-Quest. Eventually, we concatenate germline IGHV, the junction sequence, and the IGHJ sequences to obtain the hypothetical naive sequence.

0.2.2 Tree comparison and evaluation

To measure the performance of B cell lineage reconstruction algorithms, we used two metrics to compare trees: graph editing distances (29), and ancestral sequence inferences (10, 16).

Graph Editing Distance Let G_1 and G_2 be two graphs; the Graph Editing Distance (GED) finds the minimum set of graph transformations able to transform G_1 into G_2 through edit operations on G_1 . A graph transformation is any operation that modifies the graph: insertion, deletion, and substitutions of vertices or edges. GED is similar to string edit distances such as Levenshtein distance (21) when we replace strings by connected directed acyclic graphs of maximum degree one. We used two versions of GED, one applied to the whole tree (GED tree-based), and another applied to each branch separately (GED path-based). The latter version is more stringent than the first one since any difference in the

path from each leaf to the root is considered a transformation.

The problem of computing graph edit distance is NP-complete (30), and there is no optimal solution in a reasonable time. This problem is difficult to approximate, and most approximation algorithms have cubic computational time (31, 32). Here we could use an optimal algorithm implementation since the number of nodes of evaluated trees was small. Nevertheless, we favored a grid cluster to compute GED for trees with more than 50 nodes. The code is available at <https://github.com/julibinho/ClonalTree>.

Ancestral reconstruction distances We also compared trees by measuring their ancestral sequence reconstruction agreement. For that, we used two measures: The Most Recent Common Ancestor (MRCA) (10), and The Correctness Of Ancestral Reconstruction (COAR) (16). Both measures emphasize the importance of a correct ancestral reconstruction and do not penalize minor differences in the tree topologies whether the ancestral reconstruction is accurate. The classical MRCA distance is calculated by iterating through all pairs of leaves. It compares the most recent ancestral sequence of each pair of leaves presenting into two trees simultaneously — for instance, the ground truth, the inferred tree, or any pair of a compared tree. Since internal nodes in the BCR lineage tree can also represent observed sequences, we modified MRCA to consider all node pairs instead of leaves only. For a given pair of nodes i and j belonging to two trees A and B , the $MRCA_{ij}$ is the normalized hamming distance between the most recent ancestral sequences of A and B trees; the MRCA metric is the average of $MRCA_{ij}$ for all node pairs. COAR (10) is another measure allowing the reconstruction of ancestral sequences. It compares paths in the trees going from the root (the naive sequence) to any node. COAR is the average number of path mismatches across compared trees. The numeric values of both measures range between 0 and 1, where 0 represents a perfect ancestral sequence reconstruction, and 1 is the worst case. To see more details of these distance calculations, please refer to (10) and (16).

0.2.3 Computational tools used for comparisons

There exists a few BCR lineage reconstruction software; some of them are not available for download (12), and others are specific for Windows operational system (33). Among the available methods, we selected two state-of-art tools to compare with ClonalTree: GCtree (16) and GLaMST (19). GCtree is an optimal solution that uses dnapsars (17) to find a parsimony forest and then ranks parsimony trees according to genotype abundance information. On the other hand, GLaMST (19) is a time-efficient algorithm that uses MST, but it does not consider sequence affluence. The authors have shown a higher performance of these tools in the respective publications, GCtree outperformed IgPhyML (10), and GLaMST surpassed IgTree. Thus, we decided to keep only GCtree and GLaMST for the evaluation and comparison.

0.3 Results and Discussion

We first validated our method with several artificial lineage trees that simulate germinal center BCR diversification. Then, we used two data sets containing experimental BCR lineages for biological validations.

0.3.1 Reconstructing BCR lineage trees from simulated data

To evaluate ClonalTree performance and compare it to GCtree and GLaMST, we generated simulated data sets using 92 different simulation settings, varying the root sequence genes and the relative probabilities of mutation, insertion, and deletion (see section 0.2.1 and Table 1). The artificial lineage trees served as ground truth that we would like to recover using BCR lineage tree algorithms. Thus, the performance measures how close reconstructed trees are from simulated ground truths. To quantify this resemblance, we used graph editing distances (GED) that measure the dissimilarity between two graphs/trees and two distances related to the correctness of common ancestral inferences. We computed two types of GED distances, based on the entire tree (GED tree-based) and separated paths (GED path-based), see Section 0.2.2.

Figure 2 shows boxplots of GED distances for each compared method on the 92 simulated lin-

Figure 2: **Performance comparison between GTree, ClonalTree, and GLaMST using GED distances.** Box-plots in (a) present GED tree-based distances, while (b) display GED path-based ones

eages trees. We observed that GTree and ClonalTree had comparable performances, but ClonalTree outperformed GLaMST. Reconstructed BCR lineage trees of GTree and ClonalTree displayed similar topologies, while trees produced by GLaMST are different. For GED tree-based distances (Figure 2-A), median values were 0, 2, and 12 and the highest distance 37, 38, and 120 for GTree, ClonalTree and GLaMST, respectively. GLaMST presented the highest median value and the highest distance. ClonalTree produced 39 correct trees (GED tree-based distance equal to zero), while GLaMST produced only two trees.

GED path-based (Figure 2-B) compares each tree path, from leaves to the root, between inferred and ground truth trees; therefore, it is more sensitive to topology changes, resulting in higher distances

Figure 3: Most Recent Common Ancestor (MRCA) distance distributions. (a) compares ClonalTree, and GLaMST with ground truth trees, while (b) with trees generated by GTree.

than the GED tree-based ones. Figure 2-B confirms that GTree and ClonalTree reconstruct BCR lineage trees with similar paths. Median values were 0, 4, and 18, with the highest distance 56, 58, and 180 for GTree, ClonalTree, and GLaMST. As observed for GED tree-based, GLaMST presented the highest median value and distance. ClonalTree produced 39 correct paths (GED path-based distance equal to zero), while GLaMST produced only one path. In order to better evaluate the performance, we split the trees into three categories according to their number of sequences: small (between 30 and 50), medium (between 60 and 80), and large (having more than 90 sequences), see Figure S3. We observed a slight difference between GTree and ClonalTree, mainly in the small and large groups.

Figure 4: Correctness Of Ancestral Reconstruction (COAR) distance distributions. (a) compares ClonalTree, and GLaMST with ground truth trees, while (b) with trees generated by GTree.

On the other hand, we observed that GLaMST had difficulties in all groups; GED distances increased as the number of sequences grew.

GED measures depend on the tree topologies, mainly GED-based-path that penalizes each path change. Therefore, it is also essential to estimate the accuracy of ancestral reconstructions without penalizing minor differences in the tree topologies. For that, we used two measures: the MRCA and the COAR, see Section 0.2.2. MRCA distance focuses on the most recent common ancestor and does not consider the entire evolutionary lineage. On the other hand, COAR measures the correctness of ancestral reconstruction from the root to any leaf. We first compared ClonalTree and GLaMST to

ground truth trees and then with GTree. The latter comparison is important since it gives us a basis for evaluating these methods on experimental data sets, where the true lineage evolution is unknown. Figure 3 shows MRCA distance distributions for ClonalTree and GLaMST when compared to ground truth trees (A) or GTree (B). For both plots, we observed better performance for ClonalTree, which could reconstruct recent ancestral relationships more accurately. Figure 4 shows COAR distance distributions when comparing ClonalTree and GLaMST with ground truth trees (A) or with GTree ones (B). Similarly, we observed that the trees produced by our approach are closer to both ground truth and GTree trees. We noted that ClonalTree reconstructed trees that preserved the correctness of ancestral reconstruction, mainly when compared to ground truth trees.

0.3.2 Biological validation using BCR sequencing data

We performed a biological validation on two experimental data sets: TAS-42 and CLL-20 (Section 0.2.1). Since ground truth trees are unavailable for these samples, we compared the inferred trees of ClonalTree and GLaMST with the trees inferred by GTree. We considered trees generated with GTree as references for the experimental validation since such a tool achieved the best performance on simulated data sets. The TAS-42 is a public data set and was generated by lineage tracing and single-cell germinal center BCR sequencing (25). TAS-42 contains 66 sequences, used as input for GTree, GLaMST, and ClonalTree, see details in Section 0.2.1. The 2nd and 3rd columns of Table 2 show tree distances for ClonalTree and GLaMST when compared to GTree, respectively. The GED tree-based distance was smaller for ClonalTree, producing a tree with ten differences than GTree. GLaMST produced a more disparate tree than ClonalTree; Its GED tree-based distance was 47. GED path-based distances also showed that ClonalTree produced more similar evolutionary paths to GTree than GLaMST. Only 22 evolutionary paths differ from GTree against 230 for GLaMST. Most GLaMST paths contain various unobserved nodes, producing many mismatches compared to GTree. We temporarily removed them for evaluating the tool's performance. In the absence of unobserved nodes, ClonalTree presented just one mismatched path, while GLaMST produced seven. We also evaluated the agreement of ancestral sequence reconstructions, see Section 0.2.2. MRCA val-

ues were low for ClonalTree and GLaMST; both tools detected recent ancestral accurately for most node pairs. We did not observe a significant difference between MRCA values presented in Table 2. However, COAR distance values indicated that ClonalTree had better reconstructed entire evolutionary paths; we observed significantly lower COAR values for ClonalTree than GLaMST. Since ClonalTree achieved a lower GED path-based distance, it was expected to present a lower COAR distance. COAR also considers sequence dissimilarities rather than only differences in the paths as GED path-based distance. Both measures confirmed that ClonalTree produces trees that are closer to GCTree ones.

Table 2: Performance evaluation of ClonalTree, and GLaMST on experimental BCR repertoire data sets.*GED path-based when removing unobserved nodes

Metric	TAS-42		CLL-20	
	ClonalTree	GLaMST	ClonalTree	GLaMST
GED tree-based	10	47	3	36
GED path-based	22	230	18	1616
GED path-based *	7	2	13	
MRCA * 10^{-4}	5.16	5.7	6.4	7.5
COAR	0.02	0.54	0.11	0.78

The CLL-20 data set was obtained from blood samples of a patient with CLL. We only analyzed the most abundant clone, which contains 3406 sequences and 20 different genotypes. For these sequences, IMGT/HighV-QUEST (28) inferred naive ancestor V(D)J genes using germline genetic references from IMGT (34) database, see Section 0.2.1. The 4th and 5th columns of Table 2 show tree distances for ClonalTree and GLaMST, respectively. ClonalTree GED tree-based distance was the smallest among experimental data sets, showing that ClonalTree generated a tree with very similar topology to the one generated with GCTree. Likewise, GED path-based distance was smaller for ClonalTree. GLaMST presented an extremely high GED path-based distance; it added many unobserved nodes, causing several mismatches when comparing tree paths. We also evaluated tree path agreement without considering unobserved nodes. ClonalTree presented only two different paths against 13 of GLaMST. As observed in the data set TAS-42, we noted a slight difference between ClonalTree and GLaMST MRCA values, indicating that both tools reconstructed the most recent ancestral properly. However, a notable difference was again observed in COAR, ClonalTree distance smaller than

GLaMST one. This result is coherent with a better GED path-based distance since COAR accounts for the entire path from a leaf to the root.

The experimental results showed that the trees generated by ClonalTree were closer to GCtree ones; when compared to GLaMST, all distance metrics were smaller for ClonalTree, especially GED path-based and COAR distances. In summary, ClonalTree reconstructed more accurately the entire evolutionary lineage than GLaMST.

0.3.3 Time complexity and running time

The time complexity of a phylogenetic reconstruction algorithm is a function that represents the computing time required to analyze an instance of the problem. Typically, it depends on the number of treated sequences n and possibly other parameters. The computational cost of algorithms that explore the tree generating space exhaustively, producing eventually optimal solutions, increases considerably with input size. On the other hand, time-efficient approaches solve problems faster and more efficiently but can sacrifice optimality, accuracy, precision, or completeness. GCtree is exhaustive and finds the optimal solution, the most parsimonious tree. It uses dnaps, which have $O(n^3)$ complexity, to produce a forest of equally parsimonious trees. Then GCtree ranks equally parsimonious trees through the abundance of genotypes. Although ranking parsimony trees requires just a polynomial increase in the dnaps runtime, finding the parsimony forest is computationally demanding and restricted to analyzing small-sized problems. MST-based algorithms are time-efficient; they are faster than GCtree because they construct a single tree instead of a forest. Their time complexity is basically determined by the complexity of the MST algorithm and the pairwise distance calculation. GLaMST and ClonalTree use Prim's algorithm to grow the tree from the root node and interactively add edges and nodes (possibly unobserved). The time complexity of Prim's algorithm is $O(|E| + \log |V|)$, where $|V| = n$ and $|E|$ is the number of edges or connections among sequences. However, GLaMST computes pairwise edit distances using a dynamic programming algorithm with time complexity around $O(n^2)$. Contrarily to GCtree and GLaMST, ClonalTree uses hamming distance since sequences are previously aligned with Clustal Omega (35). Both alternatives decrease the ClonalTree time com-

plexity substantially. Hamming distance is computed in $O(n)$, while multiple sequence alignment in $O(n * (\log(n))^2)$. In order to measure the running time of different tools, we used the data set provided by GLaMST, containing 684 observed sequences. The three algorithms were compared on a desktop workstation with i7 processor at 3.40GHz and 32G memory. GCtree took 4.5 days, GLaMST spent 55 mins, while ClonalTree took less than 15 min to reconstruct a lineage tree for this data set. Note that we used a pre-compiled version of GLaMST that is slower than the source code. GLaMST was coded in Matlab, a proprietary programming language, and only an executable for Windows is available for users that do not possess a Matlab license. Since ClonalTree is coded in python and uses freely available libraries and software, it is ready for the scientific community.

1 Conclusion

We addressed the computational problem of reconstructing lineage trees from high-throughput BCR sequencing data. It is challenging since this data often contains a subset of BCR sequences, partially representing the dynamic process of BCR affinity maturation. Therefore, efficient algorithms need to reconstruct the entire evolutionary history from partially observed sequences. We proposed ClonalTree, a fast method that combines minimum spanning trees with genotype abundance information to reconstruct accurate evolutionary trees. ClonalTree outperformed GLaMST, a method based only on the minimum spanning tree approach, on simulated and experimental data. ClonalTree has systematically produced more maximum parsimonious trees than GLaMST.

Compared to GCtree, an exhaustive sequence-based phylogenetic method, ClonalTree presented a comparable performance, mainly on experimental data, where we observed few differences in the produced trees. However, ClonalTree is computationally more efficient, achieving equivalent results with a lower runtime. ClonalTree was hundreds to thousands of times faster than GCtree, allowing the analysis of large data sets within minutes or seconds and with minor loss of accuracy. ClonalTree's high and fast performance allows the users to consider all the available sequences when reconstructing lineage trees. It can help researchers understand B cell receptor affinity maturation lineages, mainly when data from a dense quantitative sampling of diversifying loci are available. Integrating Clonal-

Tree into existing BCR sequencing analysis frameworks could speed up lineage tree reconstructions without compromising the quality of evolutionary trees.

Acknowledgements

Authors are grateful to Lucile Jausset and Thibaud Verny for their insightful comments and fruitful discussion. We also acknowledge the anonymous reviewers that contributed to improve the manuscript and the method.

Funding

This work has been supported by "2016 Programme Doctoral de Cancérologie" and SIRIC CURA-MUS.

References

1. Susumu Tonegawa. Somatic generation of antibody diversity. *Nature*, 302(5909):575–581, 1983.
2. Gabriel D Victora and Michel C Nussenzweig. Germinal centers. *Annual review of immunology*, 30:429–457, 2012.
3. Masayuki Kuraoka, Aaron G Schmidt, Takuya Nojima, Feng Feng, Akiko Watanabe, Daisuke Kitamura, Stephen C Harrison, Thomas B Kepler, and Garnett Kelsoe. Complex antigens drive permissive clonal selection in germinal centers. *Immunity*, 44(3):542–552, 2016.
4. Jeroen MJ Tas, Luka Mesin, Giulia Pasqual, Sasha Targ, Johanne T Jacobsen, Yasuko M Mano, Casie S Chen, Jean-Claude Weill, Claude-Agnès Reynaud, Edward P Browne, et al. Visualizing antibody affinity maturation in germinal centers. *Science*, 351(6277):1048–1054, 2016.
5. Charles FA de Bourcy, Cesar J Lopez Angel, Christopher Vollmers, Cornelia L Dekker, Mark M Davis, and Stephen R Quake. Phylogenetic analysis of the human antibody repertoire reveals

- quantitative signatures of immune senescence and aging. *Proceedings of the National Academy of Sciences*, 114(5):1105–1110, 2017.
6. Yana Safonova and Pavel A Pevzner. Igevolution: clonal analysis of antibody repertoires. *bioRxiv*, page 725424, 2019.
7. Donald W Lee, Ilja V Khavrutskii, Anders Wallqvist, Sina Bavari, Christopher L Cooper, and Sidhartha Chaudhury. Brilia: integrated tool for high-throughput annotation and lineage tree assembly of b-cell repertoires. *Frontiers in immunology*, 7:681, 2017.
8. Kenneth B Hoehn, Anna Fowler, Gerton Lunter, and Oliver G Pybus. The diversity and molecular evolution of b-cell receptors during infection. *Molecular biology and evolution*, 33(5):1147–1157, 2016.
9. Eltaf Alamyar, Patrice Duroux, Marie-Paule Lefranc, and Véronique Giudicelli. Imgt® tools for the nucleotide analysis of immunoglobulin (ig) and t cell receptor (tr) v-(d)-j repertoires, polymorphisms, and ig mutations: Imgt/v-quest and imgt/highv-quest for ngs. In *Immunogenetics*, pages 569–604. Springer, 2012.
10. Kristian Davidsen and Frederick A Matsen IV. Benchmarking tree and ancestral sequence inference for b cell receptor sequences. *Frontiers in immunology*, 9:2451, 2018.
11. Alexander Yermanos, Victor Greiff, Nike Julia Krautler, Ulrike Menzel, Andreas Dounas, Enkelejda Miho, Annette Oxenius, Tanja Stadler, and Sai T Reddy. Comparison of methods for phylogenetic b-cell lineage inference using time-resolved antibody repertoire simulations (ab-sim). *Bioinformatics*, 33(24):3938–3946, 2017.
12. Michal Barak, Neta S Zuckerman, Hanna Edelman, Ron Unger, and Ramit Mehr. Igtree©: creating immunoglobulin variable region gene lineage trees. *Journal of immunological methods*, 338(1-2):67–74, 2008.

13. Joseph Felsenstein. Evolutionary trees from dna sequences: a maximum likelihood approach. *Journal of molecular evolution*, 17(6):368–376, 1981.
14. Pietro Liò and Nick Goldman. Models of molecular evolution and phylogeny. *Genome research*, 8(12):1233–1244, 1998.
15. Kenneth B Hoehn, Gerton Lunter, and Oliver G Pybus. A phylogenetic codon substitution model for antibody lineages. *Genetics*, 206(1):417–427, 2017.
16. William S DeWitt III, Luka Mesin, Gabriel D Victora, Vladimir N Minin, and Frederick A Matsen IV. Using genotype abundance to improve phylogenetic inference. *Molecular biology and evolution*, 35(5):1253–1265, 2018.
17. Joseph Felsenstein. *PHYLIP (phylogeny inference package), version 3.5 c*. Joseph Felsenstein., 1993.
18. Theodore Edward Harris. The theory of branching process. 1964.
19. Xingyu Yang, Christopher M Tipton, Matthew C Woodruff, Enlu Zhou, F Eun-Hyung Lee, Ināki Sanz, and Peng Qiu. Glamst: Grow lineages along minimum spanning tree for b cell receptor sequencing data. *BMC genomics*, 21(9):1–11, 2020.
20. RC Prim. Shortest connection networks and some generalizations. *Bell System Technical Journal*, 36(6):1389–1401, 1957.
21. Vladimir I Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, pages 707–710, 1966.
22. Joseph B Kruskal. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical society*, 7(1):48–50, 1956.

23. Feixue Huang, Pengfei Gao, and Yu Wang. Comparison of prim and kruskal on shanghai and shenzhen 300 index hierarchical structure tree. In *2009 International Conference on Web Information Systems and Mining*, pages 237–241, 2009.
24. Michael L Fredman. Fibonacci Heaps and Their Uses in Improved Network Optimization Algorithms. Technical Report 3.
25. Visualizing antibody affinity maturation in germinal centers. *Science*, 351(6277):1048–1054, 2016.
26. Gur Yaari, Jason Vander Heiden, Mohamed Uduman, Daniel Gadala-Maria, Namita Gupta, Joel NH Stern, Kevin O’Connor, David Hafler, Uri Laserson, Francois Vigneault, et al. Models of somatic hypermutation targeting and substitution based on synonymous mutations from high-throughput immunoglobulin sequencing data. *Frontiers in immunology*, 4:358, 2013.
27. Duncan K Ralph and Frederick A Matsen IV. Likelihood-based inference of b cell clonal families. *PLoS computational biology*, 12(10):e1005086, 2016.
28. Marie-Paule Lefranc, Patrice Duroux, Shuo Li, Véronique Giudicelli, and Eltaf Alamyar. IMGT/highv-quest: the IMGT web portal for immunoglobulin (ig) or antibody and t cell receptor (tr) analysis from ngs high throughput and deep sequencing. *Immunome Research*, 08(01), 2012.
29. Alberto Sanfeliu, Alberto Sanfeliu, and King Sun Fu. A Distance Measure Between Attributed Relational Graphs for Pattern Recognition. *IEEE Transactions on Systems, Man and Cybernetics*, SMC-13(3):353–362, 1983.
30. M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. Mathematical Sciences Series. W. H. Freeman, 1979.
31. M. Neuhaus and H. Bunke. *Bridging the Gap Between Graph Edit Distance and Kernel Machines*. Series in machine perception and artificial intelligence. World Scientific, 2007.

32. K Riesen. *Structural Pattern Recognition with Graph Edit Distance: Approximation Algorithms and Applications*. Advances in Computer Vision and Pattern Recognition. Springer International Publishing, 2016.
33. Thomas B Kepler. Reconstructing a b-cell clonal lineage. i. statistical inference of unobserved ancestors. *F1000Research*, 2, 2013.
34. Véronique Giudicelli, Denys Chaume, and Marie-Paule Lefranc. IMGT/GENE-DB: a comprehensive database for human and mouse immunoglobulin and T cell receptor genes. *Nucleic Acids Research*, 33(Database issue):D256–D261, 2004.
35. Fabian Sievers and Desmond G Higgins. Clustal omega, accurate alignment of very large numbers of sequences. In *Multiple sequence alignment methods*, pages 105–116. Springer, 2014.