






Sequence analysis

Detecting high-scoring local alignments in pangenome graphs

Tizian Schulz ^{1,2,3}, Roland Wittler ^{1,2}, Sven Rahmann ⁴, Faraz Hach ^{5,6} and Jens Stoye ^{1,2,*}

¹Faculty of Technology and Center for Biotechnology (CeBiTec), Bielefeld University, Bielefeld 33615, Germany, ²Bielefeld Institute for Bioinformatics Infrastructure (BIBI), Bielefeld University, Bielefeld 33615, Germany, ³Graduate School 'Digital Infrastructure for the Life Sciences' (DILS), Bielefeld University, Bielefeld 33615, Germany, ⁴Genome Informatics, Institute of Human Genetics, University Hospital Essen, University of Duisburg-Essen, Essen 45122, Germany, ⁵Vancouver Prostate Centre, Vancouver, BC V6H 3Z6, Canada and ⁶Department of Urologic Sciences, University of British Columbia, Vancouver, BC V6T 1Z4, Canada

*To whom correspondence should be addressed.

Associate Editor: Janet Kelso

Received on August 24, 2020; revised on December 2, 2020; editorial decision on January 25, 2021; accepted on January 29, 2021

Abstract

Motivation: Increasing amounts of individual genomes sequenced per species motivate the usage of pangenomic approaches. Pangenomes may be represented as graphical structures, e.g. compacted colored de Bruijn graphs, which offer a low memory usage and facilitate reference-free sequence comparisons. While sequence-to-graph mapping to graphical pangenomes has been studied for some time, no local alignment search tool in the vein of BLAST has been proposed yet.

Results: We present a new heuristic method to find maximum scoring local alignments of a DNA query sequence to a pangenome represented as a compacted colored de Bruijn graph. Our approach additionally allows a comparison of similarity among sequences within the pangenome. We show that local alignment scores follow an exponential-tail distribution similar to BLAST scores, and we discuss how to estimate its parameters to separate local alignments representing sequence homology from spurious findings. An implementation of our method is presented, and its performance and usability are shown. Our approach scales sublinearly in running time and memory usage with respect to the number of genomes under consideration. This is an advantage over classical methods that do not make use of sequence similarity within the pangenome.

Availability and implementation: Source code and test data are available from <https://gitlab.uni-bielefeld.de/gi/plast>.

Contact: jens.stoye@uni-bielefeld.de

Supplementary information: [Supplementary data](#) are available at *Bioinformatics* online.

1 Introduction

1.1 Motivation

Substantial technological advances in DNA sequencing made genomic data become one of the largest types of information kept by humankind (Stephens *et al.*, 2015). Thus, finding efficient ways of storing and analyzing these data is of high importance. The discipline of *computational pangenomics* tries to cope with this challenge (Marschall *et al.*, 2016). A *pangenome* is defined as a set of genomic sequences that may be stored and analyzed collectively while being represented as a single entity. Sequences within these sets are usually strongly related and thus highly similar. Commonly, a pangenome comprises all genomic sequences of a species, but this definition may be widened or tightened to any other taxonomic unit or kind of

distinction. The pangenomic approach allows a high memory saving potential as sequence parts shared by multiple genomes have to be stored only once. In addition, it enables the simultaneous comparison of a large number of individual genomes while avoiding classical reference-based analyses that turned out to have shortcomings in various cases (Brandt *et al.*, 2015; Degner *et al.*, 2009).

Pangenomes can be represented in many forms, ranging from pure collections of raw sequences (Vernikos *et al.*, 2015) over alignment-based approaches using multiple sequence alignments (Edgar and Batzoglou, 2006; Notredame, 2007) to graphical structures (e.g. Dilthey *et al.*, 2015; Garrison *et al.*, 2018; Iqbal *et al.*, 2012). In this work, we focus on graphical structures. In particular, we represent pangenomes as compacted colored de Bruijn graphs (Marcus *et al.*, 2014; Minkin *et al.*, 2017). Advantages of this representation

over others include low memory and storage footprints, versatility in accepting different input types (raw sequences or assemblies) and fast and alignment-free construction.

1.2 Background

Basic kinds of comparisons on a pangenome are different variants of sequence searches. Detecting exact or approximate matches can serve to answer questions like presence or absence queries. More complex analyses often involve alignment-based methods.

Algorithms for sequence-to-graph alignment have already been studied for some time. Early works date back to 1989 where a graph was used for approximate regular expression mapping (Myers and Miller, 1989). In 2000, an algorithm to align sequences to arbitrary graphs was proposed in the context of hypertext search (Navarro, 2000). Lee et al. (2002) introduced *partial order alignment* (POA) on directed acyclic graphs in 2002. Works by Diltney et al. (2015, 2016) are restricted to acyclic graphs as well. The tool ‘vg’ uses POA on general graphs (Garrison et al., 2018) by *unrolling* cyclic parts. Read mapping on de Bruijn graphs can be done by a heuristic proposed by Limasset et al. (2016). Recently, a method was published allowing exact read mapping on general graphs (Rautiainen et al., 2019). However, covering a distance based scenario only, a generalization to local alignment is non-trivial and makes this approach limited to semi-global alignment. Other solutions have been presented by Antipov et al. (2016) and Kavya et al. (2019).

These solutions address a sequence mapping scenario where query sequences are aligned either globally or locally to the graph region that best matches the query. This approach is usually designed for rapidly mapping a huge number of queries gaining speed by the underlying assumption that a query either maps to exactly one position in the graph or to none.

In this work, we study the problem of finding high-scoring local alignments between a query sequence and a graph that are likely to represent sequence homology. The exact notion of ‘high-scoring’ is based on statistical considerations. Hence, we are within the regime of the popular Basic Local Alignment Search Tool (BLAST; Altschul et al., 1990). Even though BLAST is still widely used, many other solutions have been presented for heuristic sequence alignment searches (e.g. Kent, 2002). Some exploit new algorithmic techniques or data structures to improve sensitivity, run time or both (e.g. Edgar, 2010; Frith and Shrestha, 2018; Steinegger and Söding, 2017). Others exclusively focus on protein alignment (Buchfink et al., 2015; Suzuki et al., 2015; Vaser et al., 2016; Zhao et al., 2012). Recently, the tool BlastFrost (Luhmann et al., 2021) appeared, which enables sequence queries on a pangenome graph. However, it does not calculate alignments.

1.3 Contribution

In comparison to the abovementioned alignment tools that work on collections of individual sequences, our approach stores genomic sequences in a graph to analyze them jointly. Moreover, by querying the graph, we are able to compare these sequences not only to the query but also among each other. This has an advantage over other approaches where an all-against-all post-processing of results would add a quadratic number of comparisons to their running time. To this end, we introduce the notion of a *quorum* and a *search color set* to allow for customized searches in specific parts of the pangenome. A quorum here is meant as the number of individual sequences of the pangenome that have to share a graph sequence to consider it during search. A search color set enables the search to be focused on graph sequences that appear in individual sequences from the set. For instance, a user might be interested only in alignments shared by the majority of genomes showing a certain phenotype. Finally, we present first results of alignment statistics considering that all sequences in a pangenome are related and highly similar. This is in contrast to the general initial, underlying independence assumptions. To our knowledge, our approach is unique and has never been studied before.

The remainder of this manuscript is organized as follows. In Section 2, we define our model, formally state the problem and

describe the algorithmic procedure. Section 3 contains use cases of our algorithm and comparisons of its performance. In Section 4, we discuss our results. The source code of our implementation, instructions on how to generate samples for our statistical parameter estimation and test data used in this article are available from <https://gitlab.ub.uni-bielefeld.de/gi/plast>.

2 Materials and methods

2.1 Basic definitions

A *string* is a sequence of characters drawn from a finite, non-empty set, called *alphabet*. For a given string s , we denote its length by $|s|$, the character at position i by $s[i]$ and the substring starting at position i and ending at position j by $s[i..j]$. A string of length k is called a *k-mer*. For any decomposition $s = xy$, the (potentially empty) substrings x and y are called *prefix* and *suffix* of s , respectively.

In this work, we assume that all strings are over the DNA nucleotide alphabet $\Sigma_{\text{DNA}} = \{\text{A, C, G, T}\}$. Then, a *query* is a finite string q over Σ_{DNA} . A *genome* is a set of strings over Σ_{DNA} that can be many millions of short sequences representing raw data produced by a sequencing machine, or a few long sequences representing chromosomes or contigs of a complete or draft assembly. Each genome is identified with a unique color of the *universal color set* U , and the color is assigned to all strings of the genome to distinguish between sequences from different genomes.

2.2 Compacted colored de Bruijn graphs

Let $k \geq 2$. A *compacted colored de Bruijn graph* of dimension k over an alphabet Σ and a color set U is a vertex-labeled directed graph $G = (V, E, \lambda, C)$, where each vertex $v \in V$ is labeled with a sequence $\lambda(v)$ of length at least k , each k -mer r appears in the label of at most one vertex of the graph, there exists an edge $(v, w) \in E$ from vertex v to vertex w if and only if the $(k-1)$ -length suffix of $\lambda(v)$ equals the $(k-1)$ -length prefix of $\lambda(w)$, and C assigns a color set $C(r) \subseteq U$ to each k -mer r that appears in any vertex label. For convenience, we write $|v|$ instead of $|\lambda(v)|$ for the length of the label of v . In addition, $\lambda'(v)$ denotes the label of v except its last $k-1$ characters, i.e. $\lambda'(v) := \lambda(v)[1..(|v| - k + 1)]$.

2.3 Locations and truncated paths

We represent a *location* in a compacted colored de Bruijn graph $G = (V, E, \lambda, C)$ by a pair (v, i) with $v \in V$ and $1 \leq i \leq |v|$. We say that a k -mer r *overlaps* a location $l = (v, i)$ if and only if r is a substring of $\lambda(v)$ starting at position o such that $\max(1, i - k + 1) \leq o \leq \min(i, |v| - k + 1)$. For an interval $[b..e]$ of positions within v , $1 \leq b \leq e \leq |v|$, we define the *interval location* as $L(v, b, e) = \{(v, b), (v, b + 1), \dots, (v, e)\}$. The interval location $L(v, 1, |v|)$ of a complete vertex v is denoted by $L(v)$.

We say that $p = (v_0, v_1, \dots, v_z)$ is a *path* in G if $v_0, v_1, \dots, v_z \in V$ and $(v_i, v_{i+1}) \in E$ for all i with $0 \leq i < z$. A triple $t = (p, b, e)$ is called a *truncated path* in G from location (v_0, b) to location (v_z, e) passing vertices v_1, \dots, v_{z-1} if and only if (i) $p = (v_0, v_1, \dots, v_z)$ is a path in G ; (ii) $b \in \{1, \dots, |v_0|\}$; (iii) $e \in \{1, \dots, |v_z|\}$; and (iv) if $z = 0$ then $b \leq e$ (see Fig. 1 for an example). A truncated path $t = (p, b, e)$ has a sequence

$$\lambda(t) := \begin{cases} \lambda(v_0)[b..e], & \text{if } z = 0, \\ \lambda(v_0)[b..|v_0| - k + 1]\lambda'(v_1) \dots \\ \lambda'(v_{z-1})\lambda(v_z)[1..e], & \text{otherwise.} \end{cases}$$

The *path location* of a truncated path $t = (p, b, e)$ is defined as

$$L(t) := \begin{cases} L(v_0, b, e), & \text{if } z = 0, \\ L(v_0, b, |v_0|) \cup \bigcup_{v \in \{v_1, \dots, v_{z-1}\}} L(v) \\ \cup L(v_z, 1, e), & \text{otherwise.} \end{cases}$$

The color set of a path location L is then defined as $C(L) := \bigcup_{r \text{ overlaps any } l \in L} C(r)$. Given a quorum value

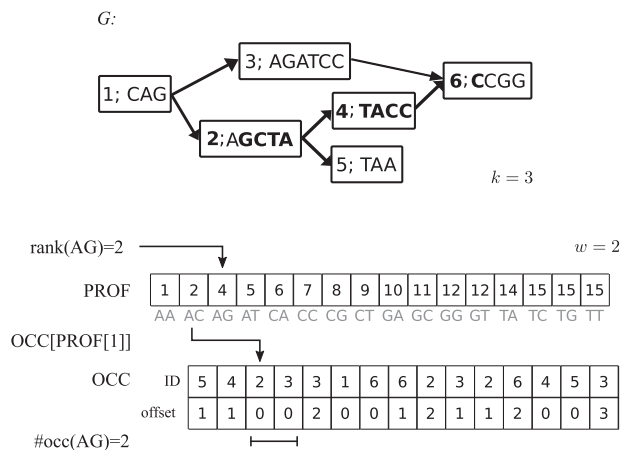


Fig. 1. Used data structures. **Top:** Example of a compacted de Bruijn graph G . The truncated path $t = (1, (2, 4, 6), 0)$ is marked in bold. **Bottom:** Visualization of an index search using arrays PROF and OCC. Occurrences of the 2-mer AG in G are accessed by looking up AG's position in PROF that provides the offset in OCC where #occ(AG) locations corresponding to occurrences of AG in G are stored consecutively. Array and sequence indices start with 0

$m \in \{1, \dots, |C|\}$, we say a truncated path t is *quorum fulfilling* if and only if $|C(l)| \geq m$ for every location $l \in L(t)$.

2.4 Formal problem statement

Our algorithm finds high-scoring local alignments between a given query sequence q and a pangenome represented as a compacted colored de Bruijn graph $G = (V, E, \lambda, C)$ over the DNA alphabet and a color set U . Apart from q and G , it takes as input a non-empty color set $C_{\text{search}} \subseteq U$ called *search color set* and a quorum $m \in \{1, \dots, |C_{\text{search}}|\}$. It outputs a set \mathcal{A} , where each $A \in \mathcal{A}$ is a pair $(\alpha, [c_1, c_2, \dots])$ consisting of an alignment α between a substring of q and the sequence $\lambda(t)$ of a quorum fulfilling truncated path t , and a list of checkpoints c_1, c_2, \dots partitioning t into maximal subpaths with the same color set, each containing a subset of C_{search} of cardinality at least m . Our algorithm allows to find local alignments for both q and its reverse complement. For the sake of simplicity, we forgo to mention the reverse complement in the following descriptions. Procedures may be assumed to work similarly for the reverse complementary sequence.

2.5 Algorithm

The algorithm consists of three basic algorithmic steps. In the *seed detection step*, maximal exact matches of a minimal length $w \leq k$ are searched between q and all sequences of G . Found seeds are extended without gaps in the *seed extension step*. Statistical considerations are used to extract biologically meaningful alignments. These are then recalculated in the *gapped recalculation step* to obtain the finally reported results.

2.5.1 Preprocessing: Graph construction and index generation

While the query sequence q has to be specified as an input parameter, the graph G may be given or initially constructed from a set of genomes in a preprocessing step. Once G has been created for a specific value of k , an additional index is built according to a minimal seed length $w \leq k$.

The index consists of two arrays that allow a fast location of exact matches between q and sequences stored in G . Array PROF is a w -mer profile that contains, in form of a cumulative sum, for any w -mer ω the number of occurrences of ω in G , denoted #occ(ω). In other words, PROF[0] = #occ(ω_0) and

$$\text{PROF}[i] := \#occ(\omega_i) + \text{PROF}[i-1] \quad \text{for all } i, 1 \leq i < 4^w,$$

where ω_i is the i -th w -mer over Σ_{DNA} in lexicographic order. The second array, OCC, harbors the occurrences of all w -mers in the

graph. An occurrence of a w -mer ω in G is defined as a pair $\text{occ}(\omega) := (v, o)$ where $v \in V$ and o represents the offset in $\lambda(v)$ at which ω starts. In OCC, all occurrences of a w -mer are stored consecutively in some arbitrary order such that OCC[PROF[$i-1$]] contains the first occurrence of ω_i and OCC[PROF[i]] the last. It may be accessed by using a ranking function to get position i in PROF. In case a w -mer falls into the $(k-1)$ -length label overlap of two consecutive vertices $u \rightarrow v$, it is counted only once and only its occurrence in v is stored. An example of the index and its use can be found in Figure 1.

2.5.2 Seed detection step

The index is used to look up all w -mers occurring in q . All matches are checked whether they fulfill the quorum requirement. Consecutive matches may be merged into a single interval location. A match is stored in a list directly linked to the vertex whose sequence it was found in. This allows a quick access to the match if needed during seed extension. Inside the list, matches are ordered by increasing starting offset in the vertex label. Match order is used to terminate a list iteration as soon as it is clear that a demanded match cannot be part of a list anymore, which leads to an additional speed gain.

2.5.3 Seed extension step

Each match represents an interval location used as seed for an ungapped extension. The extension is performed according to an X-drop algorithm as used in BLAST along q and all truncated paths of G starting at an initial interval location (v, b, e) . The paths are generated by a depth-first search traversal through G outgoing from vertex v . The exploration of a path stops as soon as either the current extension's score drops below X or the quorum requirement is no longer fulfilled. Our experience shows that the exhaustive processing of all existing truncated paths starting from v through G is possible without a notable slowdown of the algorithm most of the time even though the number of paths may be very large. However, an iteration over all truncated paths can become infeasible if G has some dense regions consisting of vertices with short labels and high degrees, producing truncated paths with highly similar sequences. If these sequences are similar to the query q , the X-drop criterion does not suffice to terminate path explorations leading to suboptimal alignments within an acceptable time frame. Therefore, we introduced a threshold that limits the number of vertices that can be visited during the extension of a seed.

2.5.4 Gapped recalculation step

The result of the ungapped extension is a set of path locations consisting of a few biologically meaningful and many more spurious hits. Statistical significance criteria, described in Section 2.6, are used to rank the hits and separate both kinds. The most significant alignments are recalculated using standard gapped alignment with banded dynamic programming. The bandwidth is chosen according to the quality of the alignment. The final alignments are reported, together with statistics for gapped alignments.

2.6 Alignment statistics on a pangenome graph

In classical pairwise (linear) sequence comparison, an extensive statistical theory exists, sometimes referred to as Karlin-Altschul theory (Karlin and Altschul, 1990), but subsequently extended, refined and made practical by many others (e.g. Pearson, 1998; Waterman and Vingron, 1994). No statistical theory currently exists for sequence queries against (graphical) pangenomes, so we provide a baseline here.

A key question is as follows: Given a score value s , how many hits H_s of a random query sequence of length n against the given reference (genome or pangenome) of size m are observed whose score reach or exceed s ? The expected value $E_s = \mathbb{E}[H_s]$ is called the *E-value* of an observed score s . Let S be the score of the best hit. Then we may observe the event $\{S \geq s\}$ that is equivalent to $\{H_s \geq 1\}$ (as one implies the other), and the probability $p_s := \mathbb{P}[S \geq s] = \mathbb{P}[H_s \geq$

1] of the event that there exists at least one hit with score reaching s is called the *p-value* of an observed score s . If E_s is small (say ≤ 0.05) by Poisson approximation and first-order approximation of the exponential function, then we have

$$p_s = 1 - \mathbb{P}[H_s = 0] \approx 1 - \exp(-E_s) \approx E_s,$$

and we need not distinguish in practice between E-value and p-value. We summarize an approximate version of the classical Karlin–Altschul statistics for pairwise sequence alignment and then discuss our proposal about how to generalize the theory for query-to-pangenome alignment.

2.6.1 Summary of approximate Karlin–Altschul statistics

Comparing the expected numbers of hits $E_s, E_{s+1}, E_{s+2}, \dots$ with increasing scores $s, (s+1), (s+2), \dots$, when already E_s is small, one observes that $E_{s+k} \approx \mu^k \cdot E_s$ for some factor $0 < \mu < 1$, mostly independent of s , as long as s is large enough. Factor μ is typically written as $e^{-\lambda}$ with some $\lambda > 0$. This holds if the average score when comparing two single nucleotides is negative; otherwise, there exist arbitrarily long high-scoring matches. If we increase the length n of the query or length m of the reference, we provide more possibilities for a hit, and the expected number of such hits increases linearly with both n and m . This leads to

$$p_s \approx E_s \approx Kmn \cdot e^{-\lambda s}, \tag{1}$$

where $K > 0$ and $\lambda > 0$ are constants, n is the query length and m is the reference genome length. Note that this approximation is only valid for the extreme tail of the distribution (large s , small $E_s \approx p_s \leq 0.05$). The functional form of (1) has been empirically shown to be very robust, valid for ungapped and gapped alignments and even when considering compositional bias (difference in GC content) between query and reference, or when considering a fixed query against a random reference (Wolfsheimer *et al.*, 2011). Constants K and λ depend on the scoring scheme, including the gap costs for gapped alignment. In practice, λ must be determined by sampling and simulation (Altschul *et al.*, 2001; Wolfsheimer *et al.*, 2011).

2.6.2 Statistics for pangenome alignment

We hypothesize that a relation as in Equation (1) can be observed when considering the top score of a random query aligned against a pangenome, if s is sufficiently large such that hits reaching score s are rare. We further hypothesize that such a relation holds for both ungapped high-scoring pairs after seed extension and for final gapped alignments, albeit with different values of $\lambda > 0$ and $K > 0$. However, the dependency of λ and K on sequence relatedness and diversity within a pangenome may be complex, and it is out of scope of this work to investigate the details. Instead, we investigate whether indeed there holds an affine–linear relationship $\log p_s \approx C - \lambda s$ with $C := \log(Kmn) \in \mathbb{R}$ for fixed query length n and a pangenome graph of size m .

2.6.3 Parameter estimation by importance sampling

To obtain a good estimate of $\lambda > 0$ and $C \in \mathbb{R}$, we need good estimates of small probabilities p_s for large s . Using random sequences, large s with small $p_s < 10^{-6}$ are by definition rare, so too many samples would be needed for accurate estimates. Hence, we only use this ‘naive’ random sampling strategy to obtain an initial estimate of C and λ and then resort to importance sampling, using a Metropolis–Hastings Markov Chain Monte Carlo strategy similar to the one described by Wolfsheimer *et al.* (2011).

In brief, let $\pi_s := \mathbb{P}[S = s]$ be the unknown score distribution on integers s . We construct a Markov chain in such a way that the probability to sample a random sequence with score s is exponentially biased toward higher scores, $r_s := \pi_s \cdot \exp(\lambda_0 \cdot s)/Z$, where $\lambda_0 < \lambda$ should slightly underestimate the true λ and can be derived from the initial naive sampling step, and Z is the appropriate (unknown) normalization constant such that $\sum_s r_s = 1$. To avoid

computing Z explicitly, we use the Metropolis–Hastings method: Given a current DNA sequence x , a new candidate sequence y is proposed from a neighborhood of x (see Wolfsheimer *et al.*, 2011 for the precise definition of the neighborhood). Roughly, a single nucleotide can be inserted, deleted or substituted at any position, deleting or inserting a nucleotide at the left or right end to keep the sequence length n constant. Thus, the edit distance between x and the new proposal y is at most 2. This creates a Markov chain where, in equilibrium, each sequence is equally probable, similarly to the naive simulation, where each nucleotide is drawn independently from a uniform distribution. Now, to bias the samples toward higher scores, the scores s_x and s_y of x and proposal y , respectively, are compared. We accept y with probability $\min\{1, (r_{s_y}/\pi_{s_y})/(r_{s_x}/\pi_{s_x})\} = \min\{1, \exp(\lambda_0 \cdot (s_y - s_x))\}$, i.e. a score increase is always accepted, and a decrease only with a certain probability. When a proposal is rejected, x stays the current sequence and another proposal is generated. A score sample is drawn after a large number of accepts that allows the query sequence to change considerably in comparison to the previous sample. Typically, $2n/3$ accepts suffice for a query of length n . The first few samples are discarded to allow the Markov chain to reach equilibrium. Since the defined chain is rapidly mixing (there exist short paths from every sequence to every other sequence), we found it sufficient to discard the first five samples.

The procedure yields uncorrelated score samples, which stem from the biased distribution $r = (r_s)$. Let R_s be the absolute number of times that score s was sampled, and let $T_s := \sum_{s' \geq s} R_{s'} \cdot \exp(-\lambda_0 \cdot s')$. Then λ can be estimated by fitting a line to points $(s, \log T_s)$ in an interval of s where the counts R_s are consistently high, say $R_s \geq 50$. Then, C is estimated from the 10% of highest scores in the initial naive sampling step.

3 Results

We implemented the algorithm described in Section 2.5 in C++ using Bifrost (Holley and Melsted, 2020) as the underlying realization of a compacted colored de Bruijn graph. Among several other implementations (e.g. Almodaresi *et al.*, 2017; Chikhi *et al.*, 2016; Holley *et al.*, 2016; Iqbal *et al.*, 2012; Muggli *et al.*, 2017), we chose Bifrost since it is an efficient, easy-to-use implementation, allows the usage of assembled and raw sequencing data and provides the possibility to assign any kind of data to vertices of the graph. Our implementation called PLAST¹ (‘Pangenome Local Alignment Search Tool’) is available from <https://gitlab.uni-bielefeld.de/gi/plast>. We present results of our statistical parameter estimation, followed by a performance analysis of our method. In particular, we show the advantage in runtime, memory usage and result aggregation when searching local alignments inside a pangenome with our method compared to a conventional search and analysis using other BLAST-like software tools. Afterwards, we present a more advanced use case and show that our method scales even to human data. Unless stated differently, all graphs have been built for a k -mer length of 31 (Bifrost default), and searches were performed using default parameters. To obtain statistical parameters, we used a combination of initial naive sampling followed by importance sampling, as described in Section 2.6.

3.1 Statistical parameter estimation

We tested the hypothesis from Section 2.6 that $\log p_s \approx C - \lambda s$ for constants $C \in \mathbb{R}$ and $\lambda > 0$ for both ungapped and gapped alignments. After initial confirming results on simulated pangenomes (not shown), we considered a real pangenome of 220 *Salmonella enterica* genomes of the same lineage (Para C) taken from Zhou *et al.* (2018). Naive simulation was performed with one million random DNA sequences of length $n=200$, yielding the empirical complementary cumulative distribution function (ccdf) of the best hit’s

¹ Not to be confused with a software of the same name introduced in (Van Nguyen and Lavenier, 2009) parallelizing the conventional BLAST algorithm.

score for each query for both the gapped and ungapped case. A least-squares fit of affine linear functions in the distribution's tail, considering logarithmic cumulative relative frequencies between 10^{-2} and 10^{-4} , yielded initial estimates of $C=17.45$ and $\lambda=1.085$ for ungapped alignments and $C=14.81$ and $\lambda=0.852$ for the final gapped alignments (see Fig. 2). As expected, $\lambda_{\text{gapped}} < \lambda_{\text{ungapped}}$, as gaps provide more freedom to achieve a higher score with the same query length. The affine relationship cannot hold for much higher p-values because our approximation assumes small $p_s \leq 10^{-2}$, and we also cannot make a statement for much lower p-values with 'only' 10^6 simulations.

To gain access to the rare-event tail, we performed importance sampling as described above. Per sampled sequence and score, we need to evaluate $(2n/3)/\alpha$ sequences, where α is the average acceptance rate, which we typically find to be around 0.46 to 0.75, which amounts to approximately $(2n/3)/(1/2) = 4n/3$ evaluated sequences per drawn sample. Thus, importance sampling introduces a 250-fold overhead over naive sampling for $n=200$. However, it allows us to sample from high scores that are unobtainable by naive sampling, even if billions of samples were used, yielding much higher efficiency by several orders of magnitude (cf. importance sampling

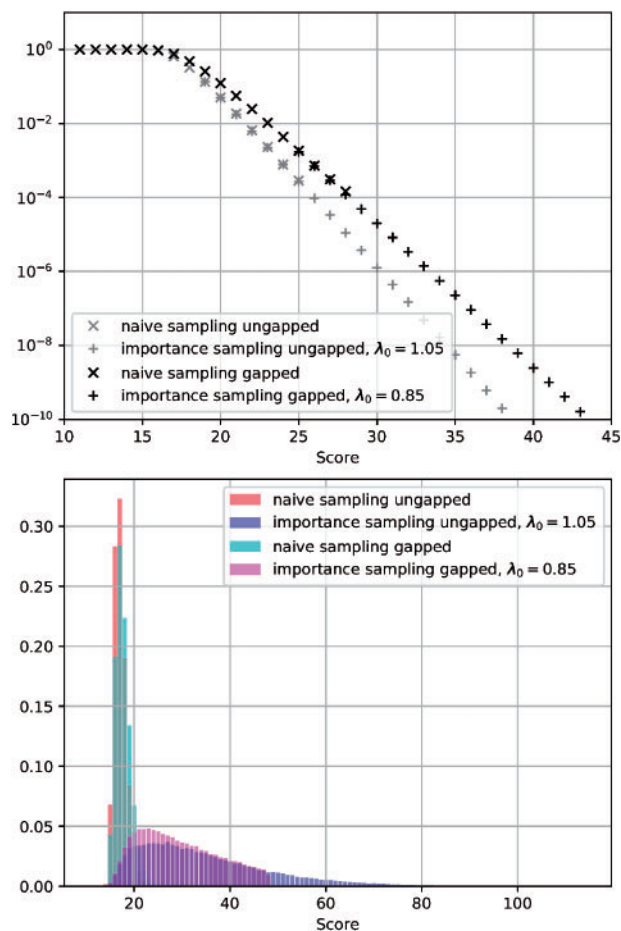


Fig. 2. Results of naive and importance sampling. **Top:** Logarithmic plot of p-values (complementary cumulative distribution function) of highest ungapped and gapped alignment scores for random queries (without color or quorum constraints) against a pangenome of 220 *Salmonella enterica* genomes. Naive simulation with 10^6 samples yields accurate estimates for p-values in the range between 10^{-2} and 10^{-4} . Importance sampling enables a better view of the rare-event tail. For p-values $\leq 10^{-2}$, the hypothesis of an affine dependency holds with values $C \approx 20.78$ and $\lambda \approx 1.136$ for ungapped alignments and $C \approx 16.11$ and $\lambda \approx 0.898$ for gapped alignments, estimated from importance sampling. **Bottom:** Histograms of (normalized) sample counts per score value for ungapped and gapped alignments, using naive sampling and importance sampling. Importance sampling gives access to a broader interval of scores in the rare-event tails

tails in Fig. 2). We obtain further refined estimates of λ and C for both ungapped and gapped alignments. As expected, using weight factors of $\lambda_0 = 1.05$ for ungapped and $\lambda_0 = 0.852$ for gapped that are slightly smaller than the 'true' $\lambda = 1.136$ for ungapped and $\lambda = 0.898$ for gapped, we observe an almost flat and slowly decreasing histogram of score counts (Fig. 2) and thus sample from a broad interval of scores.

3.2 Comparison to other tools

We evaluated our approach by comparing our implementation against MMseqs2 (as of January 8, 2020), blastn (2.6.0+), BLAT (36x4) and UBLAST (11.0.667). The tools DIAMOND, SWORD and GHOSTZ do not support DNA to DNA alignment. RAPSearch2 was not possible to install due to a reported, but unresolved issue. LAST is known to run very slow on highly redundant datasets—index building was terminated after 10 days. BLAT was designed to search genomes (represented as target databases) for query sequences. It was only possible to run the tool for 750 genomes at once. Runs for larger pangenome sizes had to be split into separate program calls. Results and runtimes were aggregated. Similarly, UBLAST's freely available 32-bit version has an upper database size limit such that pangenomes consisting of more than 100 genomes had to be distributed on multiple databases.

For the comparison, we downloaded 5000 randomly selected *Salmonella Typhimurium* assemblies from a total of 19237 that were annotated as serovar 'Typhimurium' from EnteroBase (Alikhan et al., 2018). As queries we chose 100 random substrings of length 1000 from the *Salmonella* reference genome assembly (RefSeq assembly accession GCF_000195995.1). We obtained queries with an average percentual identity of about 86.9% (SD 31.1) per query and genome where 55 953 query genome combinations without any alignment were considered with an identity of zero.

If possible, tools were run with the same scores for match, mismatch, gaps and the same X-drop value. Additionally, we set the maximum number of reported results high enough to get all existing results. This was necessary, because all other tools do not compare their results internally and would otherwise report only the best result for each genome separately, hiding all further findings. For all remaining parameters, default values were used. Concrete program calls are documented at <https://gitlab.uni-bielefeld.de/gi/plast>. Calculations were performed single threaded on a virtual machine with 28 cores and 256 GB of RAM.

To compare the results, we scanned the output of each tool and identified corresponding alignments of PLAST and BLAST. In case of PLAST, this identification was possible using the color sets that are part of the program's output for each alignment as explained in Section 2.4. Two results are considered *matching* if they overlap by at least 90% of the shorter alignment with respect to the query sequence. We call an alignment *unique* if it does not match any alignment of another tool. The comparison is shown in Table 1. We see that PLAST reports the lowest total number of results. This is due to the fact that one alignment of PLAST may correspond to several genomes of the pangenome. All other tools report such alignments separately, and an additional post-processing step would be required to merge them. UBLAST reports by far the highest number of alignments. Thus, it seems to be most sensitive in our experiment. Apart from UBLAST, the number of results that could be found by some tool but not PLAST is below 0.03% (columns 'Tool\PLAST' and 'BLAST\Tool', first row). The percentage of results unique to PLAST varies between $\sim 4\%$ and $\sim 6\%$ among the tools. However, only 178 PLAST alignments (2.4%) were unique with respect to all other tools. Their score was low (mean 18.5 and maximum 25) and they were short (mean length 28 bp and maximum length 112 bp). They are a result of the different statistical parameters used by PLAST in comparison to other tools. As explained in Section 1, these tools do not consider a pangenomic use case where all sequences in the database are highly similar. Thus, they overestimate the chance for a random hit. Using the same statistical parameters as BLAST, all 178 results are filtered out by PLAST due to the E-value threshold. We observed 306 PLAST alignments (4.0%) that were unique

Table 1. Comparison of PLAST to other alignment tools

Tool	Results	Tool\ PLAST	PLAST\ Tool	Tool\ BLAST	BLAST\ Tool
PLAST	7 565	–	–	357 4.72 %	290 0.02 %
BLAST	1 246 221	290 0.02 %	357 4.72 %	–	–
BLAT	457 089	1 0.00 %	456 6.03 %	508 0.11 %	49 798 4.00 %
MMseqs2	695 792	6 0.00 %	322 4.26 %	800 0.12 %	21 022 1.69 %
UBLAST	4 881 509	111 386 2.28 %	272 3.59 %	220 577 4.52 %	5 459 0.44 %

Note: 100 random substrings of length 1 000 from the *Salmonella* reference genome assembly have been aligned to the ‘Typhimurium’ dataset. The columns list the tool names, absolute number of alignments and a comparison to PLAST and BLAST, where X\Y denotes the number of results of X that did not match any result of Y.

with respect to at least one other tool but at the same time found by at least one other. 56 of them had a score higher than 25 (maximum 678) or were longer than 112 bp (maximum 999 bp). Generally, we found that PLAST was able to find alignments with even very low sequence identity values of up to 58%, which is comparable to BLAST (63%) in this setting.

Next, we randomly selected subsamples from the 5 000 *Salmonella* genome assemblies to generate pangenomes of different sizes and compared the tools’ performance. Runtime and memory usage are shown in Figure 3. Generally, we can see that most tools quickly lose speed with a growing pangenome size. Only MMseqs2 is able to keep a speed comparable to PLAST even for a pangenome of 5 000 individual genomes. However, its speed comes along with a memory requirement between ~9 GB and 176 GB, which is far beyond that of all remaining tools.

All tools spend a large amount of memory for loading the highly redundant datasets and additional index data structures for fast alignment searches. Our graphical representation, in turn, allows a maximal exploitation of sequence similarity. Most parts of the graph represent sequences shared by many individual genomes that have to be stored only once. An alignment calculated for these parts is valid for all sharing genomes and does not have to be calculated multiple times. Thus, we observe a strongly reduced growth in run




Fig. 3. Run time and memory usage comparison of all tools. The areas marked by a dashed rectangle in the plots on the right are shown in separate plots on the left. Values on less than 1 000 genomes have been averaged over five random subsamples each. The same legend applies to all plots

time and memory consumption for PLAST. Considering a pangenomic use case with constantly growing numbers of individual genomes, our approach will always be superior compared to conventional methods not making use of sequence similarity.

When comparing PLAST's run time for different quorum values and pangenome sizes, we can generally observe two driving forces (data not shown). On the one hand, the usage of a quorum has a beneficial influence on PLAST's run time, because it allows to disregard parts of the graph harboring rarely appearing variations that may prune it considerably, especially if the overall diversity within the pangenome is high. At the same time, a quorum check is comparatively cheap if the graph has only a moderate total number of colors in it. On the other hand, for larger pangenomes, the usage of a high quorum becomes an increasing burden as color coverages have to be checked on every vertex using Bifrost's API. The high degree of color compression generates a noticeable loss in speed in this case, which outstrips any gain by pruning from a certain pangenome size. In our experiments, it doubled computation time compared to a run without quorum for the largest pangenome size.

To prevent this loss in speed, a first filter was implemented that, instead of only checking color presence, also considers the number of missing colors on a vertex and rejects it as soon as this number becomes too high during iteration. Other heuristics to speed up a quorum check are currently under development and discussed in Section 4.

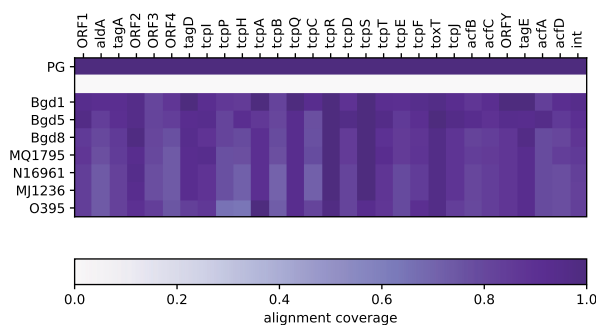
3.3 Pathogenicity islands in *Vibrio cholerae*

Pandemic strains in *V. cholerae* are known to contain the *Vibrio Pathogenicity Island-1* (VPI-1) that consists of a whole set of genes whose sequence and order inside the island can be diverse among different strains [Karaolis et al., 2001].

We examined a recent collection of 21 *V. cholerae* genomes, 7 of which have been obtained from clinical samples and are labelled 'pandemic genomes' (PG), and the remaining 14 have been obtained from non-clinical samples and are labelled 'environmental genomes' (EG) [Shapiro et al., 2017, primary dataset]. Some of the samples were available as assemblies, whereas others could be found only as raw read datasets. Nevertheless, one graph was built from all 21 samples, using the functionality of the underlying Bifrost library to construct a de Bruijn graph from both types of data: read data, which are automatically filtered for low coverage k -mers, as well as assembly data that are not filtered.

As a proof of concept, we used PLAST to search VPI-1 inside the *V. cholerae* pangenome—once restricting the search to alignments with PG genomes to verify the existence of VPI-1 in PG, and once restricting the search to alignments with EG genomes to verify the absence of VPI-1 in EG. For the search, we set an E-value cutoff of 0.01, the maximum number of alignments to maximum, and used standard or automatically determined parameters otherwise.

The order of VPI genes within PG genomes may be rearranged, and EG genomes might contain only small fragments of the island. Even so, we used the complete island sequence (accession no. AF325734 [Karaolis et al., 2001]) of length 41272 bp as query.



PLAST's ability to compute *local* alignments and to report many suboptimal findings, nevertheless, allowed an easy detection of each VPI gene separately. Furthermore, conserved gene orders could be detected by local alignments spanning larger fragments of the island. We observed 64 alignments in the PG search that were longer than 2814 bp, i.e. three times the median length of coding sequences in VPI-1.

For each coding sequence in VPI-1, we determined the maximum overlap by any local alignment. As can be seen in Figure 4, when restricting the alignment to PG, all coding sequences are covered by alignments almost completely (row 'PG'). In contrast, when restricting the alignment to EG, only very few (three of twenty-nine) coding sequences are covered by any alignment by 50% or more (row 'EG'). We want to highlight here that the detection of such outstanding sequences that are contained in any of a whole group of genomes is possible by a single PLAST search. For further investigations, it would be possible to analyze the PLAST alignments for individual genomes, as exemplified in the remaining rows in Figure 4, where the maximum overlap is determined among those parts of the alignments that are supported by the corresponding individual genomes.

3.4 Beyond bacterial pangenomes

To test PLAST's applicability beyond bacterial pangenomes, we built a human pangenome using data from the 1000 Genomes (1KG) Project phase 3 (1000 Genomes Project Consortium et al., 2015). We used bcftools consensus (https://github.com/samtools/bcftools) to generate chromosome-wise genomic sequences for chromosomes 2 and 15 of all 2504 human individuals by inserting all reported variations into the GRCh37 reference sequence and built one graph ($k=63$) for each of the two chromosomes. Instead of building a single graph from all human chromosomes at once, we had to perform this chromosome-wise approach as our system could not store all 2504 human genome sequences including indexing structures at once for a comparative call of MMsseqs2 (see below). Our method alone also scales to a whole human pangenome of this size.

We exemplarily investigated a known polymorphism within the human pangenome. The single nucleotide polymorphism *rs1426654* is reported to influence skin pigmentation. Its reference allele indicates a light skin color that is common to a West Eurasian ancestry [Lamason et al., 2005; Soejima and Koda, 2006]. It is located on exon 3 of gene *SLC24A5* on chromosome 15. Searching the exon sequence from reference genome GRCh37 within the pangenome of chromosome 15 resulted in two full size alignments: a perfect match representing the reference allele, and one having a single mismatch representing the variant allele. Restricting our search to the European core genome (search set of all European samples and a quorum of 99%), we exclusively found the reference allele, confirming the results of the earlier studies.

To evaluate PLAST's performance on the human pangenome, we used the graph of chromosome 2 ($\approx 8\%$ of the complete human

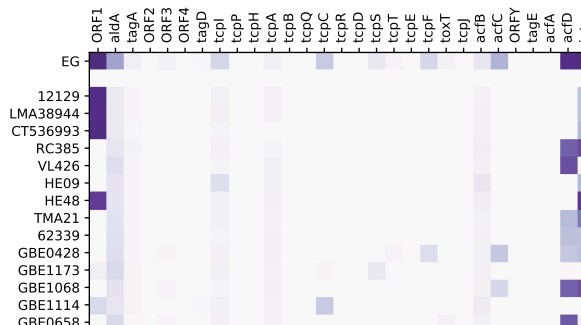


Fig. 4. PLAST search results for VPI-1 [Karaolis et al., 2001] inside the *Vibrio cholerae* dataset [Shapiro et al., 2017] with search color sets PG (left) and EG (right). Columns correspond to coding sequences in VPI-1 ordered as appearing in the query sequence. Rows 'PG' and 'EG' show the maximum overlap of any alignment with a coding sequence inside the pangenome using PG or EG as search color set, respectively. Rows 'Bgd1' to 'O395' and '12129' to 'GBE0658' show the maximum coverage of a sample in any alignment overlapping a coding sequence using PG or EG as search color set, respectively

genome) and searched 100 queries of length 1000 randomly drawn from the human reference genome. For a comparison, we also ran MMseqs2, next to PLAST the fastest tool according to Section 3.2, on the same queries. PLAST took 317 s per query on average for the complete dataset using a maximum of 24 GB of memory and running on a single thread. Using only a subset of 1000 chromosomes, MMseqs2 took 339 s per query on average using a maximum of 259 GB of memory and running on all available 28 cores. Furthermore, its input files (sequences and index structures) required 2.2 TB of disk space. PLAST's input files occupied only 9.2 GB on disk in total. Both tools were run with default parameters.

4 Discussion

We presented a new BLAST-like method to find highest scoring local alignments between a query sequence and a pangenome represented as a colored de Bruijn graph. Unlike read mapping tools developed to find the best or possibly a few suboptimal mapping positions for potentially many query sequences in a short time, our aim is to find all such alignments with statistically significant score. Using a minimal seed length much smaller than k increases our detection sensitivity that goes far beyond the scope of a k -mer-based seeding approach, while alignment statistics enable us to filter the results for biologically meaningful ones.

By working on a graph, our method is able to exploit the high degrees of sequence similarity within a pangenome. On the one hand, this allows to draw conclusions not only about the similarity toward a query sequence but also to compare the diversity of genomic sequences in the graph with regard to the query. On the other hand, it avoids the storage and processing of highly redundant information and lets the run time and memory usage of our algorithm scale sublinearly with respect to the number of genomes inside the pangenome. Both advantages make our approach superior in comparison to conventional local alignment search tools working on a database of multiple individual genomes, scaling linearly in the number of sequences and enabling us to handle even large eukaryotic pangenomes. We showed this in a comparison to other state-of-the-art BLAST-like alignment search tools and by using human data from the 1KG Project.

Moreover, we introduced the usage of a quorum and a search color set that allow to limit searches on customized regions of the pangenome. This is extremely useful for answering specific research questions. Additionally, it avoids repeated construction of the graph for different database sequence sets, which is especially important if pangenomes are large and graph construction becomes prohibitively expensive. As a practical application of our algorithm, we demonstrated its usability on a classical pangenomic use case in Section 3.3.

Although our implementation PLAST performs well in practice, it has to be seen as a proof of concept implementation so far. For example, no affine gap cost model has been incorporated yet and is subject to ongoing work. Also, many heuristics are strongly oriented on common techniques for plain sequences. We are convinced that further efforts on the development of heuristics exploiting the special conditions prevailing in a graphical pangenome may lead to even more efficient algorithms.

Yet, quorum checks may become a time-determining factor if large pangenomes are considered and quorums are high. The reason is that Bifrost compresses color information within binary matrices for each vertex, and accessing this information requires a time consuming iteration over the matrix. Currently, we are working on ways to aggregate this information as a preprocessing step next to index building to avoid matrix iterations during the search. Aggregated quorum information could be stored for each vertex, e.g. by using five integers encoding the minimum quorum fulfilled on the vertex, the potentially higher quorum being fulfilled at its beginning and at the end and the sequence offsets at which this higher quorum breaks. A different idea would be to store the number of colors covering a k -mer separately for each k -mer using a single byte. The memory footprint for both ideas would be small especially in pangenomes of very closely related individual genomes.

We established that, similarly to the statistical behavior of pairwise alignment, sequence-to-graph alignment p-values exhibit exponential tails, $\log p_s \approx C - \lambda \cdot s$ for constants $C \in \mathbb{R}, \lambda > 0$. Still, many interesting open questions remain about these statistics. The calculation of precise statistical parameters for each pangenome graph is based on compute-intensive simulations so far. To benefit from existing simulations on similar graphs, instead of starting a new simulation for every pangenome graph, we would like to better understand how the statistical parameters change with graph properties and scores. While the dependency on gap scores has been explored in the past, nothing is known about how (and which) graph properties influence the parameters, and preliminary experiments show complex patterns, which we intend to investigate in future work.

Furthermore, we would like to further explore the computational limits using our implementation. The here presented datasets were the largest we considered so far in terms of individual samples (Section 3.2) and the number of k -mers (Section 3.4).

We also plan to further extend the concept of quorum and search color set. In particular, we would like to give more freedom to the user by allowing not only to focus on certain parts of the graph but also to explicitly exclude regions from analysis.

Funding

This work is supported by the BMBF-funded de.NBI Cloud within the German Network for Bioinformatics Infrastructure (de.NBI) [031A537B, 031A533A, 031A538A, 031A533B, 031A535A, 031A537C, 031A534A and 031A532B]. It was funded in part by the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie agreement [872539] and the International DFG Research Training Group GRK 1906 to T.S.; and National Science and Engineering Council of Canada (NSERC) Discovery Grants (RGPIN-05952) and Michael Smith Foundation for Health Research (MSFHR) Scholar Award (SCH-2020-0370) to F.H.

Conflict of Interest: none declared.

References

- 1000 Genomes Project Consortium. (2015) A global reference for human genetic variation. *Nature*, **526**, 68–74.
- Alikhan, N.-F. *et al.* (2018) A genomic overview of the population structure of *Salmonella*. *PLOS Genet.*, **14**, e1007261.
- Almodaresi, F. *et al.* (2017) Rainbowfish: a succinct colored de Bruijn graph representation. In *17th International Workshop on Algorithms in Bioinformatics (WABI 2017)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, Article No. 18.
- Altschul, S.F. *et al.* (1990) Basic local alignment search tool. *J. Mol. Biol.*, **215**, 403–410.
- Altschul, S.F. *et al.* (2001) The estimation of statistical parameters for local alignment score distributions. *Nucleic Acids Res.*, **29**, 351–361.
- Antipov, D. *et al.* (2016) hybridSPAdes: an algorithm for hybrid assembly of short and long reads. *Bioinformatics*, **32**, 1009–1015.
- Brandt, D.Y. *et al.* (2015) Mapping bias overestimates reference allele frequencies at the HLA genes in the 1000 Genomes Project Phase I data. *G3: Genes, Genomes, Genetics*, **5**, 931–941.
- Buchfink, B. *et al.* (2015) Fast and sensitive protein alignment using DIAMOND. *Nat. Methods*, **12**, 59–60.
- Chikhi, R. *et al.* (2016) Compacting de Bruijn graphs from sequencing data quickly and in low memory. *Bioinformatics*, **32**, i201–i208.
- Degner, J.F. *et al.* (2009) Effect of read-mapping biases on detecting allele-specific expression from RNA-sequencing data. *Bioinformatics*, **25**, 3207–3212.
- Dilthey, A. *et al.* (2015) Improved genome inference in the MHC using a population reference graph. *Nat. Genet.*, **47**, 682–688.
- Dilthey, A.T. *et al.* (2016) High-accuracy HLA type inference from whole-genome sequencing data using population reference graphs. *PLoS Comput. Biol.*, **12**, e1005151.
- Edgar, R.C. (2010) Search and clustering orders of magnitude faster than BLAST. *Bioinformatics*, **26**, 2460–2461.
- Edgar, R.C. and Batzoglou, S. (2006) Multiple sequence alignment. *Curr. Opin. Struct. Biol.*, **16**, 368–373.

- Frith, M.C. and Shrestha, A.M. (2018) A simplified description of child tables for sequence similarity search. *IEEE/ACM Trans. Comput. Biol. Bioinform.*, **15**, 2067–2073.
- Garrison, E. et al. (2018) Variation graph toolkit improves read mapping by representing genetic variation in the reference. *Nat. Biotechnol.*, **36**, 875–879.
- Holley, G. and Melsted, P. (2020) Bifrost: highly parallel construction and indexing of colored and compacted de Bruijn graphs. *Genome Biol.*, **21**, 1–20.
- Holley, G. et al. (2016) Bloom Filter Trie: an alignment-free and reference-free data structure for pan-genome storage. *Algorithms Mol. Biol.*, **11**, 3.
- Iqbal, Z. et al. (2012) De novo assembly and genotyping of variants using colored de Bruijn graphs. *Nat. Genet.*, **44**, 226–232.
- Karaolis, D.K. et al. (2001) Comparison of *Vibrio cholerae* pathogenicity islands in sixth and seventh pandemic strains. *Infect. Immun.*, **69**, 1947–1952.
- Karlin, S. and Altschul, S.F. (1990) Methods for assessing the statistical significance of molecular sequence features by using general scoring schemes. *Proc. Natl. Acad. Sci. USA*, **87**, 2264–2268.
- Kavya, V.N.S. et al. (2019) Sequence alignment on directed graphs. *J. Comput. Biol.*, **26**, 53–67.
- Kent, W.J. (2002) BLAT—the BLAST-like alignment tool. *Genome Res.*, **12**, 656–664.
- Lamason, R.L. et al. (2005) SLC24A5, a putative cation exchanger, affects pigmentation in zebrafish and humans. *Science*, **310**, 1782–1786.
- Lee, C. et al. (2002) Multiple sequence alignment using partial order graphs. *Bioinformatics*, **18**, 452–464.
- Limasset, A. et al. (2016) Read mapping on de Bruijn graphs. *BMC Bioinformatics*, **17**, 237.
- Luhmann, N. et al. (2021) BlastFrost: fast querying of 100,000s of bacterial genomes in Bifrost graphs. *Genome Biol.*, **22**, 30.
- Marcus, S. et al. (2014) SplitMEM: a graphical algorithm for pan-genome analysis with suffix skips. *Bioinformatics*, **30**, 3476–3483.
- Marschall, T. et al. (2016) Computational pan-genomics: status, promises and challenges. *Brief. Bioinform.*, **19**, 118–135.
- Minkin, I. et al. (2017) TwoPaCo: an efficient algorithm to build the compacted de Bruijn graph from many complete genomes. *Bioinformatics*, **33**, 4024–4032.
- Muggli, M.D. et al. (2017) Succinct colored de Bruijn graphs. *Bioinformatics*, **33**, 3181–3187.
- Myers, E.W. and Miller, W. (1989) Approximate matching of regular expressions. *Bull. Math. Biol.*, **51**, 5–37.
- Navarro, G. (2000) Improved approximate pattern matching on hypertext. *Theor. Comput. Sci.*, **237**, 455–463.
- Notredame, C. (2007) Recent evolutions of multiple sequence alignment algorithms. *PLoS Comput. Biol.*, **3**, e123.
- Pearson, W.R. (1998) Empirical statistical estimates for sequence similarity searches. *J. Mol. Biol.*, **276**, 71–84.
- Rautiainen, M. et al. (2019) Bit-parallel sequence-to-graph alignment. *Bioinformatics*, **35**, 3599–3607.
- Shapiro, B.J. et al. (2017) Origins of pandemic *Vibrio cholerae* from environmental gene pools. *Nat. Microbiol.*, **2**, 16240.
- Soejima, M. and Koda, Y. (2006) Population differences of two coding SNPs in pigmentation-related genes SLC24A5 and SLC45A2. *Int. J. Legal Med.*, **121**, 36–39.
- Steingger, M. and Söding, J. (2017) MMseqs2 enables sensitive protein sequence searching for the analysis of massive datasets. *Nat. Biotechnol.*, **35**, 1026–1028.
- Stephens, Z.D. et al. (2015) Big data: astronomical or genomic? *PLoS Biol.*, **13**, e1002195.
- Suzuki, S. et al. (2015) Faster sequence homology searches by clustering subsequences. *Bioinformatics*, **31**, 1183–1190.
- Van Nguyen, H. and Lavenier, D. (2009) Plast: parallel local alignment search tool for database comparison. *BMC Bioinformatics*, **10**, 329.
- Vaser, R. et al. (2016) SWORD—a highly efficient protein database search. *Bioinformatics*, **32**, i680–i684.
- Vernikos, G. et al. (2015) Ten years of pan-genome analyses. *Curr. Opin. Microbiol.*, **23**, 148–154.
- Waterman, M.S. and Vingron, M. (1994) Rapid and accurate estimates of statistical significance for sequence data base searches. *Proc. Natl. Acad. Sci. USA*, **91**, 4625–4628.
- Wolfsheimer, S. et al. (2011) Accurate statistics for local sequence alignment with position-dependent scoring by rare-event sampling. *BMC Bioinformatics*, **12**, 47.
- Zhao, Y. et al. (2012) RAPSearch2: a fast and memory-efficient protein similarity search tool for next-generation sequencing data. *Bioinformatics*, **28**, 125–126.
- Zhou, Z. et al. (2018) GrapeTree: visualization of core genomic relationships among 100,000 bacterial pathogens. *Genome Res.*, **28**, 1395–1404.