

UC Santa Cruz

UC Santa Cruz Previously Published Works

Title

Pangenome Graphs

Permalink

<https://escholarship.org/uc/item/1jf167k7>

Journal

Annual Review of Genomics and Human Genetics, 21(1)

ISSN

1527-8204

Authors

Eizenga, Jordan M

Novak, Adam M

Sibbesen, Jonas A

et al.

Publication Date

2020-08-31

DOI

10.1146/annurev-genom-120219-080406

Peer reviewed



Published in final edited form as:

Annu Rev Genomics Hum Genet. 2020 August 31; 21: 139–162. doi:10.1146/annurev-genom-120219-080406.

Pangenome graphs

Jordan M Eizenga¹, Adam M Novak¹, Jonas A Sibbesen¹, Simon Heumos², Ali Ghaffaari^{5,6,7}, Glenn Hickey¹, Xian Chang¹, Josiah D Seaman^{3,4}, Robin Rounthwaite¹, Jana Ebler^{5,6,7}, Mikko Rautiainen^{5,6,7}, Shilpa Garg^{8,9,10}, Benedict Paten¹, Tobias Marshall^{5,6}, Jouni Sirén¹, Erik Garrison^{*,1}

¹Genomics Institute, University of California Santa Cruz, Santa Cruz, CA, USA, 95064

²Quantitative Biology Center, University of Tübingen, Tübingen, Germany, 72076

³Royal Botanic Gardens Kew, Richmond, UK, TW9 3AB

⁴Queen Mary University of London, London, UK, E1 4NS

⁵Center for Bioinformatics, Saarland University, Saarbrücken, Germany, 66123

⁶Max Planck Institute for Informatics, Saarbrücken, Germany, 66123

⁷Saarbrücken Graduate School for Computer Science, Saarbrücken, Germany, 66123

⁸Department of Genetics, Harvard Medical School, Boston, MA 02215

⁹Department of Data Sciences, Dana-Farber Cancer Institute, Boston, MA 02215

¹⁰Department of Biomedical Informatics, Harvard Medical School, Boston, MA 02215

Abstract

Low-cost whole genome assembly has enabled the collection of haplotype-resolved pangenomes for numerous organisms. In turn, this technological change is encouraging the development of methods that can precisely address the sequence and variation described in large collections of related genomes. These approaches often use graphical models of the pangenome to support algorithms for sequence alignment, visualization, functional genomics, and association studies. The additional information provided to these methods by the pangenome allows them to achieve superior performance on a variety of bioinformatic tasks, including read alignment, variant calling, and genotyping. Pangenome graphs stand to become a ubiquitous tool in genomics. Although it is unclear if they will replace linear reference genomes, their ability to harmoniously relate multiple

*corresponding author: erik.garrison@ucsc.edu.

AUTHOR CONTRIBUTIONS

JME wrote §4.5, §6.2, §6.3, and §8.1, made Table 2, and revised the paper; AMN wrote §6.1, contributed to §6.3, and revised the paper; JAS wrote §8.2 and §8.3, and contributed to section 6.3; SH made Table 1 and contributed to §4.4, §6.1, and Figure 2. AG made Figure 1; GH contributed to §8.1; XC contributed to §6.3; JDS contributed to Table 1 and contributed significantly to §6.1; RR contributed to §4; JE contributed to §8.1; MR contributed to §8; SG contributed to §8.2; BP provided guidance and opinion; TM contributed to §1, §2 and §6.2; JS wrote §5; EG organized the work, wrote §1, §2, §3, §7, §9, made Figures 2 and 3, and revised the paper.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

sequence and coordinate systems will make them useful irrespective of which pangenomic models become most common in the future.

Keywords

pangenome; genome graph; variation graph

1. Introduction

A *pangenome* models the full set of genomic elements in a given species or clade. Pangenomics thus stands in contrast to reference-based genomic approaches which relate sequences to a particular consensus model of the genome (figure 1). Genomes that are reconstructed with the aid of a reference genome can appear more similar to the reference than they actually are. Pangenomic reference systems can reduce this bias by enabling the direct relationship of new genome to all those represented in the pangenome.

Pangenomics has been important to microbiology, where genomic plasticity and diversity has made it indispensable (141), and it has increasingly seen application to eukaryotic genomes (22, 46, 107). Standard pangenomic analyses focus on the presence or absence of genes from given strains and the determination of a core (commonly present) and accessory (frequently absent) pangenome (108). However, they have tended to pay less attention to variation between these sequences, and they do not typically attempt to provide a precise model relating many genomes to each other at the base level.

In contrast, most “high-throughput” analyses of large genomes depend on comparison to a single reference genome. However, in recent years, reduced sequencing and *de novo* assembly costs have supported the discovery of significant levels of large-scale genomic variation in many eukaryotic species, including humans (135, 58, 24, 8), arabidopsis (3), brewer’s yeast (145), and the fruit fly (25). These observations have generated wide interest in extending bioinformatic operations to use a pangenomic reference model (32). The availability of true pangenomic references for humans (28) and other model organisms will increasingly render the use of a single reference genome suboptimal. However, using pangenomes effectively requires the development of new bioinformatic methods capable of constructing, querying, and operating on them.

In this review, we consider an emerging class of bioinformatic methods that let us consider genetic diversity at every stage of analysis. We focus primarily on those methods which achieve this result by replacing linear reference genome system with a graphical, pangenomic one. To provide background, we first consider how limitations of the current dominant genome inference paradigm motivate these methods (§2). Then, we describe pangenomic models, with a particular focus on graphical ones (§3). This provides a foundation to understand how they can be constructed from sequencing data or assembled genomes (§4). We then survey index data structures that allow us to interact with them efficiently (§5). These index structures allow us to interrogate, visualize, and relate new information to pangenomes, such as in the process of read alignment (§6), yielding results in a variety of new data formats (§7). Finally, we examine a number of downstream

applications of these models (§8), and reflect on the future impact of these methods on genomics (§9).

2. Resequencing reconsidered

Biological research frequently depends on our ability to infer the sequence level relationships between genomes. In an earlier era, when many sequencing studies focused on small regions or single genes, it was often feasible to relate all relevant sequences to each other. Although limited in scope, these analyses were effectively pangenomic. They were based on many-to-many relationships between sequences, typically derived by multiple sequence alignment (MSA). Precise MSA methods are expensive, with popular algorithms scaling cubically with the number of input sequences (103). It is infeasible to apply such computationally demanding methods to the data scales obtained with modern high-throughput sequencing. Instead, high-quality genome assemblies and high-throughput sequencing have encouraged *resequencing* methodologies, wherein reads from each sample are aligned to a single reference genome. This approach is practical and scalable. State of the art resequencing pipelines can jointly analyze tens of thousands of genomes (112) at a cost per genome that is only a small fraction of the total sequencing cost.

Although efficient and conceptually simple, resequencing has a significant limitation. Precise genomic relationships are only visible for those sequences that are similar enough to the reference genome to be alignable (figure 1). This effect is known as *reference bias*. It is strongest for structural variation or sequences that are absent from the reference system (135), but it can be relevant even for SNPs, which causes problems in allele specific expression (ASE) quantification (23) and in the analysis of ancient DNA (147). Given that this bias shapes the methods that we use to establish models of the truth (148), it is even difficult to evaluate without paradigmatic change in our analysis techniques.

Estimates based on short read sequencing data have placed the human pangenome at between 1% (85) and 10% (127) larger than the GRCh38 human reference assembly. Others have demonstrated up to several Mbp of sequence are present in each new individual and not in the reference (58, 8). We expect that whole genome telomere-to-telomere assemblies based on long single-molecule sequencing will provide greater insight into the extent, placement, and significance of these novel sequences (92, 77). Even if we know of their existence, the reference bias inherent in resequencing will continue to limit the ability of researchers to relate new sequences to these regions. Resequencing methods that use a pangenomic reference system should not be subject to this limitation.

3. Pangenomic models

A *pangenomic model* (figure 1) is a data structure that represents the genomic sequences of a population, a species, a clade, or even a metagenome (32). The model serves as a central coordinating entity to describe the collection of sequences and genomes in the pangenome. Pangenomic models may take many forms, including collections of unaligned sequences or learned sequence models, but here we will focus mostly on graphical ones.

3.1. Graphical models

The *sequence graph* serves to compress many redundant input sequences into a smaller data structure that is still representative of the full set (59). Sequence graphs may have their nodes or edges labeled with DNA sequences, but for simplicity we will focus on the node-labeled case. In a node-labeled sequence graph, edges indicate when concatenations of the nodes they connect occur in the sequences modeled by the graph. Walks through a sequence graph thus include the set of sequences from which it was built. We call them *bidirected* when they represent both strands of DNA and inversions between them. Sequence graphs were first used to represent multiple sequence alignments (59, 78). In assembly, they have been applied to represent the full information in a set of sequencing reads (as in the *string graph*) (98), or fixed-length *k*-mers (as in the *de Bruijn graph* (DBG)) (111).

Genome graphs are sequence graphs used to represent whole genome relationships (110). Walks through these graphs represent recombinations of the genomes included in the model. Regions of the graph where multiple paths connect a common head and tail node, often referred to as *bubbles* (109), represent variation. *Variation graphs* further structure this model by embedding the linear sequences of the pangenome as *paths* (51)¹. Paths provide a stable coordinate system that is unaffected by the manner in which the graph was built, thus supporting the coordination of positions, annotations, and alignments between variation graphs and linear reference genomes.

4. Building a pangenome

Methods to construct pangenomic data structures mirror the classes of pangenomic models. A pangenome may simply be a collection of sequences, in which case construction is similar to the genome or metagenome assembly problem. Or, it may include information about the alignment of sequences or genomes within it. This alignment could be compressed into a set of variants found against a set of reference sequences. If this alignment is based on *k*-mers, then it implies a de Bruijn graph. If it is a complete, gapped alignment, covering small and large variation, then the pangenome model can be thought of as a whole genome alignment.

4.1. Collecting sequences

A pangenome can be represented as a collection of sequences. Several approaches support the construction, annotation, and interrogation of these pangenomic sequence collections. PANSEQ (76) finds novel regions, determines the core and accessory genome, finds SNPs within the core pangenome, and then determines a subset of loci useful for molecular fingerprinting. PGAP (146) extends PANSEQ's approach with modules for evolutionary and functional analysis, and is implemented as a single standalone executable. Recent work has focused on scaling these techniques to ever larger genomes. PANTOOLS (126) detects and annotates homology groups in large collections of large genomes using a *k*-mer based approach. Its detailed graph database model connects the panproteome defined by homology groups to genomic annotations and sequences. HUPAN (38) extends the sequence collection

¹ *Variation graphs* are similar to *variant graphs* used in textual research to model a collection of revisions of the same text (124)

model to human and large eukaryotic genomes, taking assembled genomes as input and finding non-reference sequences within them by comparison to a reference genome.

4.2. Adding variation

Rather than collecting unique sequences that represent a collection of genomes, we can consider small variants between the collection and a reference genome. Such a model directly implies a directed acyclic graph, ordered along the reference genome, with bubbles at the sites of variation. This pangenome construction approach is used in diverse graph genome read mappers, including GENOMEMAPPER (125), Seven Bridges' GRAPH GENOME PIPELINE, and (116), PANVC (140), and GRAMTOOLS (88). The variation graph toolkit, specifically VG CONSTRUCT (51), can be applied to transform VCF files and reference sequences into genome graphs. Some methods, like the journaled string tree (115), and methods based on elastic degenerate texts (12) such as SOPANG (29), transform the variant set and reference into a structure optimized for online sequence queries of the pangenome. Deciding which variation should be added to a graph is non-trivial, and has encouraged studies of graph utility (105) and algorithms to determine which variation is helpful (113).

4.3. Colored, linked, and compacted de Bruijn graphs

De Bruijn graph-based assemblers can be given a pangenomic quality through the addition of “colors” to their nodes (k -mers) or unitigs (unbranching components in the graph). Each color provides a mapping between a specific biosample and a subset of the graph. CORTEX first demonstrated that colored DBGs could perform population scale analyses with an efficient graph implementation (66). Recent improvements to colored DBG construction, such as BIFROST (63), allow the construction of colored DBGs from very large sequence sets (the authors build a pangenome of 118,000 *Salmonella* genomes), and further support efficient updates of these pangenomic models. The feature which makes these methods efficient, the fixed k on which they are based, also limits their resolution of repetitive genomic features. It is not feasible to build them from noisy third-generation sequencing reads. Addressing these limitations, several methods embed linking information within the DBG that can be used to reconstruct embedded haplotypes or reads (15, 138).

A number of methods use compacted DBG² construction to elaborate pangenome graphs. SPLITMEM (90) uses a suffix tree with suffix skips to derive the set of maximal exact matches k between a set of genomes. Improving on this result in both time and space efficiency, (11) demonstrated two similar pangenome graph induction algorithms based on succinct representations of the suffix tree and the Burrows–Wheeler transform (BWT) (20). TwoPACo (95) applies a probabilistic data structure to narrow the set of candidate vertexes in the compacted de Bruijn graph of a set of genomes, supporting the efficient generation of a pangenome graph from larger genomes than previous methods.

²Compacted here means that chains of k -mers which contain no internal furcations are merged into a single node in the graph representation.

4.4. Alignment-based sequence graphs

Sequence graphs (59) can be understood as representations of the mutual alignment of a set of sequences. Alignment-based pangenome structures form the basis of a number of pangenomic approaches. They have found use in the construction of acyclic multiple sequence alignments. The partial order aligner POA (78, 54) uses an acyclic, directed sequence graph model to build multiple sequence alignments. PROGRESSIVECACTUS (7) produces whole genome alignments that can be rendered as sequence graphs. SIBELIAZ (94) finds collinear blocks within TwoPACO's compacted DBG and applies POA to each to yield a whole genome alignment graph. The multiple sequence/graph aligner (VG MSGA) (105, 51, 48) generalizes the progressive approach of POA to build generic variation graphs that include cycles and inversions.

Not all researchers have focused on generic graphs, with several arguing that completely generic models are either computationally intractable or not relevant to important practical analyses. REVEAL (86) builds a pangenome graph from a syntenic set of maximal exact unique matches of decreasing size between a pair of sequences (or graphs), and later adds inversions detected by alignment of paths through bubbles in this graph. NOVOGRAPH (14) follows a reference-guided approach, breaking a set of genomes into syntenic alignable blocks and deriving a multiple sequence alignment for each, and yielding a VCF file as its output. Similarly, SEQ-SEQ-PAN (69) employs existing whole genome alignment methods to find a set of locally collinear blocks, in which it compacts into a sequence graph that respects the synteny of the input genomes. GENGRAPH (4) realigns previously identified collinear blocks yielding a genome graph from an MSA.

Recent, unpublished methods explore two new alternatives to alignment-based pangenome construction. MINIGRAPH³ extends the MINIMAP2 (84) alignment chaining model to work on graphs. It applies this alignment model to progressively build out a pangenome graph from a series of genomes that contains large (>250bp) sequences that were not previously seen in other genomes. The resulting pangenome does not contain all input sequences and variation between them, but a representative subset and large structural variants. In contrast, SEQWISH⁴ (48) generates the full variation graph implied by a collection of sequences and alignments between them. The paths embedded in its output graph precisely and completely reconstruct the input sequences, while the topology of the graph describes all variants represented in the input alignments.

4.5. Positional systems in pangenomes

Reference genome sequences provide a coordinate system to catalog and exchange information about genes, protein binding sites, epigenetic profiles, variants, and homologies. In linear references, genomic coordinates are easily interpretable, and they unambiguously indicate both the layout of the sequence and the distance between bases, but this is not the case when these coordinates are embedded within a graph (117).

³<https://github.com/lh3/minigraph>

⁴<https://github.com/ekg/seqwish>

It is possible to use reference coordinates in a graphical pangenome by embedding reference sequences inside the graph and labeling graph nodes with their relative positions in these paths (51, 48). This has been extensively explored in variation graph based tools. However, several problems remain. The embedded coordinate systems may be incomplete, in that they may not fully cover the graph. Also, particular graph instantiations may induce ambiguity in reference positions. For instance, a copy number variant that is collapsed in the graph will contain multiple overlapping reference path coordinate ranges.

These limitations have driven the development of complete coordinate systems for genome graphs. One solution is to build a hierarchy of graph components, based on a starting reference sequence, adding a new name and coordinate range for each non-reference sequence that is included (117)⁵. Another technique is to build positional systems based solely on the topology in the graph (109). Genomic variation creates a system of nested bubble structures that can be used to spatially organize graph elements. This approach has a rigorous, if complex, mathematical basis. Similar decompositions of the graph topology have been used in assembly-based variant detection (66, 106).

5. Indexing pangenomes

Index data structures for pangenome graphs support efficient random access to elements and features of the graph. Attention must be given to ensure that these index structures do not require significant overheads relative to the information content of the graph. Naive implementations of sequence and structural indexes of the graph can incur significant runtime and memory costs, which can become problematic as graph sizes increase. Succinct data structures and careful encoding of these data are thus required to reliably fit large graphs into the main memory of commodity computing systems. Particular index models lie at the core of the highest-performing graph based visualization (§6.1), read mapping (§6.3), and variant calling systems (§8.1).

Building *text indexes* for sequences encoded in a graph is more involved than for linear references. In graphs with regions of dense variation, the number of k -bp paths can grow exponentially in k , often rendering their complete enumeration intractable even for low values of k . In order to limit the exponential growth, the index may only support relatively short query strings. Some indexes (132) support longer queries by doing extensive preprocessing. In others (137, 64, 88), queries mapping to complex graph regions can be slow. Instead of indexing the entire graph, the index may only contain k -mers from a simplified graph, or from specific paths of the graph.

5.1. Indexing sequences using a graph

The FM-index (42) is a text index, based on the Burrows-Wheeler transform (BWT) (20), that is frequently used with DNA sequences. One variant of the FM-index, the RLCSA (99), run-length encodes the BWT, allowing it to store and index a collection of similar sequences space-efficiently. If we know a good global alignment of the sequences, we can use that information to make the index both smaller and faster (65). This approach was developed

⁵MINIGRAPH uses a similar model

further in the FM-index of alignment (101, 100). Both (65) and (101) use the graph induced by the alignment as a space-efficient representation of the sequences.

5.2. Indexing acyclic graphs

One class of graph indexing methods supports only acyclic graphs, often represented as directed acyclic graphs (DAGs). This constraint can exist either because the acyclicity of the graph provides guarantees that simplify the problem, or because incidental features of the method's software implementation preclude use on cyclic graphs.

GENOMEMAPPER (125), the first graph-based read aligner, was limited to such graphs. Its indexing was also relatively simple. GENOMEMAPPER uses a simple hash-based k -mer index, with $k = 13$ to limit memory usage.

GCSA (132) was the first attempt to generalize the BWT for graphs. It applies a number of graph transformations that preserve the graph's sequence space while creating an unambiguous ordering for nodes. When the graph's complexity is low, these transformations are reasonably fast and do not increase the size of the graph significantly. However, at a certain threshold of variant density the transformed graph quickly becomes too large to handle.

BWBBLE (64) is a BWT-based representation for VCF-based pangenome graphs. Simple substitutions are encoded in the sequence using IUPAC codes, and the sequence is indexed using a normal FM-index. Because each base can be encoded using 8 different characters, the search branches at every base to cover all possible characters which admit the base searched. In practice, most branches quickly run out of results and can be pruned from the search. BWBBLE represents insertions and deletions with extra sequences, including a given amount of context around the variant. The length of this context is an effective upper bound for query length.

The vBWT (88) took another approach to using the BWT for indexing VCF-based pangenome graphs. It encodes variants as (ref|alt1|alt2|...) in the sequence. When the search encounters a variant, it must branch to handle each allele separately. Both BWBBLE and vBWT trade faster index construction for slower queries. However, a combination of IUPAC codes for substitutions, the vBWT approach for other variants, and a k -mer index for matching the first 5–10 bases, is faster than either of the originals (21).

5.3. General graphs

Some text indexes are based on Lempel–Ziv parsing or context-free grammars. These indexes first find partial matches between the query string and the indexed phrases and then combine the partial matches into full matches using two-dimensional range queries. In the hypertext index (137), each node is a separate phrase. Queries mapping to a single node or crossing a single edge can be matched efficiently, while finding mappings to complex graph regions can be slow.

Techniques similar to the GCSA can be used to represent de Bruijn graphs (16). If the graph transformations used in GCSA construction are stopped after i steps, the resulting graph is

equivalent to an order- 2^i de Bruijn graph. This de Bruijn graph can be used to approximate the original graph. Queries of this index yield no false negatives, but matches longer than 2^i may be false positives. By using this approach, GCSA2 (130) attempts to support fast queries in arbitrary graphs.

GCSA2 faces the same issues with complex graphs as GCSA. In practice, most graphs must be simplified before they can be indexed. Typical simplifications include removing high-degree nodes and complex regions from the graph and replacing them with the reference sequence. If a collection of haplotypes is available, the removed regions can be replaced with new subgraphs that contain separate paths for each distinct local haplotype (131). This way, the index contains all k -mers from the haplotypes, while usually missing some k -mers from their recombinations.

5.4. Indexing graphs using sequences

Instead of attempting to index the entire graph, it is often sufficient to index only selected paths in it. CHOP (96) takes the paths corresponding to haplotypes and breaks them into smaller pieces. The distinct pieces form an artificial linear reference, which can be used with any read aligner. The process guarantees that any substring of the haplotypes of length k is also a substring of one of the pieces. As with BWBBLE, k represents an effective upper bound for query length.

The PSI (52) follows a similar approach with artificial paths. Instead of using haplotypes, PSI uses a greedy algorithm to find a set of paths that covers as many k bp windows in the graph as possible. When a fully sensitive index is needed, PSI can reverse the role of the query strings and the graph. While complex graph regions may contain an excessive number of k -mers, the reads mapping to them only contain a limited number of k -mers. By indexing a batch of reads and searching for the complex regions in that index, all mappings of the query strings to the graph can be found with reasonable resources.

5.5. Indexing haplotypes and genomes in variation graphs

Haplotypes in related individuals are typically highly similar, and thus compressible. A series of results lead from a compact haplotype index for biallelic SNPs to a generic haplotype index for complex variation graphs. The positional BWT (PBWT) (39) provides an efficient compressed representation of a set of haplotypes over biallelic variable sites. The PBWT, like the BWT, supports efficient haplotype matching queries such as maximal exact match finding. Later work (45) showed the PBWT to be equivalent to the wavelet matrix (30), which is to date the most efficient known encoding of strings with large alphabets supporting a variety of important random access queries. The graph positional BWT (gPBWT) (104) extends the PBWT model to haplotype walks embedded in complex sequence graphs. The graph BWT (GBWT) (131) builds on a number of assumptions that tend to hold for sequence variation graphs to improve runtime and memory costs relative to the gPBWT. Provided the variation graph on which they are built encodes the full complement of variation in the pangenome, both the gPBWT and GBWT are excellent, and in principle lossless, compressors of genomes. A GBWT for the 5008 haplotypes in the 1000 Genomes Project required 14.6 GB, or ≈ 1 bit per 1 kbp of encoded genomic sequence.

6. Relating new information to the pangenome

Here we focus on two major avenues to interact with pangenome graphs. By *visualizing* these graphs, we can gain insight into the relationship between genomes, learning about variation small and large. Pangenome models are often used as a reference system for *sequence alignment*, which enables the relation of a new biosample to the pangenome to support a wide variety of downstream applications.

6.1. Visualization

The visualization of pangenome graphs presents significant challenges that have proven to be difficult to resolve in a single system. We present an overview of such tools in Table 1, while Figure 2 compares different visualization methods applied to the same graph.

Complex, nonlinear graph structures are difficult to present in a convenient number of dimensions. Traditional genome browsers organize information in a 2D rectangle, with a linear reference's sequence space providing the horizontal dimension, and a stack of annotations plotted in the vertical dimension (57). When visualizing a graph, there are two basic approaches: attempt to mimic linear browsers and compress the graph into one dimension, or give the graph two dimensions and represent annotations using interactivity, labels, or color. Either option is challenging, as pangenome graphs are often *nonplanar*, requiring at least three dimensions to draw without overlapping lines. Force-directed layout, an optimization-based graph layout method, is a popular approach in the second category, but displaying annotations in this layout is challenging.

Visualization tools be categorized by the kinds of graphs they are intended to display, in addition to whether they can linearize a graph. Many tools, including BANDAGE (142) and GFA VIZ (53), are designed for interpreting assembly graphs. These tools tend to focus on displaying the overall, rather than base-level, structure of the graph, and have no visual features to reflect the pangenomic aspect of the graph.

Tools initially designed for variation graph visualization, on the other hand, tend to focus more on base-level structure and pangenomic relationships. For example, SEQUENCE TUBE MAP (13) displays precise base-scale variation, haplotypes, and short read mapping locations, using a visual language inspired by transit system maps.

There is a difference in scale when moving from an assembly or scaffold graph to a comprehensive variation graph of a species pangenome, and this scale difference also separates different visualization tools. High-level assembly-graph tools like ASSEMBLY GRAPH BROWSER (AGB) (93) struggle to show fine details in the graph, while low-level variation graph tools like VG VIEW (51) and SEQUENCE TUBE MAP, cannot scale to large graphs. One tool which works well at both large scale and high detail is MoMI-G (144). Designed as a "multi-scale" graph browser, it presents both a SEQUENCE TUBE MAP of base-level differences and a CIRCOS (73) plot of chromosomal-scale connections, and can uniquely visualize long reads in the context of pangenomic haplotypes (144). ODGI VIZ⁶, uses binning and direct

⁶<https://github.com/vgteam/odgi>

rendering to a raster image to generate visualizations representing gigabase scale pangenomes. The approach is a rasterized version of the linear layout technique of VG viz (48).

Interactively visualizing human-genome-scale, human-pangenome-detail graphs coherently across zoom levels remains an open problem.

6.2. Graph alignment algorithms

Sequence comparison is at the core of many genomic analyses, and sequence alignment is the essential method for doing so. Classic algorithms like Smith-Waterman (133) do not directly apply to genome graphs. However, the recurrence relations that drive their scoring and traceback routines can be extended to allow the alignment of sequences to acyclic sequence graphs, as popularized by the partial order aligner POA (78). Further generalizations support the alignment of sequences graphs to sequence graphs (54), sequences to cyclic graphs (102), and even cyclic sequence graphs to cyclic sequence graphs (97, 5). It is notable that many of these findings have been independently rediscovered or refined by contemporary researchers (6, 119, 68, 75). Some earlier algorithms require restricted scoring functions to achieve efficiency (119), but recent contributions have used less restricted functions that produce more biologically meaningful alignments in some contexts (68).

Graph alignment algorithms have also become faster. POA had equivalent asymptotic run time to linear alignment but required acyclic graphs (78). Later optimizations simply ran slower on general graphs (70). Algorithms are now known with equivalent run time even on general graphs (68). In addition, researchers have developed modified algorithms that run quickly in the practical context of real-world computer architectures (136, 118, 67).

6.3. Genome graph mapping

Efficient methods to map reads to large pangenome graphs have been developed in recent years. Many draw on recent research in alignment (§6.2), and advances in pangenome indexing (§5).

Although these mapping tools all target sequence graphs, there are significant differences in the types of graphs that they handle. Several tools apply only to acyclic variation graphs formed by adding variants to a linear reference. Examples include GENOMEMAPPER (125), Seven Bridges' GRAPH GENOME ALIGNER (116), HISAT2 (72), V-MAP (139). In contrast, VG (48) and GRAPHALIGNER (120) appear to be the only tools with open ambitions of mapping to arbitrary variation graphs, including complex local and global topologies. GRAPHALIGNER can also align to generic overlap graphs and DBGs, a feature which it uses to drive the error correction of long reads using DBGs (61, 44).

The majority of these tools emphasize mapping short-read *next-generation sequencing* (NGS) data. To our knowledge, GRAPHALIGNER and V-MAP are the only graph mapping tools designed long read sequencing data (120, 139). While V-MAP also supports NGS reads, GRAPHALIGNER's seeding strategy limits it to long reads. VG also supports long-read alignment (51), but this is based on a hierarchical approach that applies the alignment

algorithm for short reads to chunks of long reads (48). The approach is accurate but nearly an order of magnitude slower than `GRAPHALIGNER` (120).

For indexing (§5), most graph mapping tools have opted for some variation of a k -mer table. `GRAPHALIGNER`, `GenomeMapper`, Seven Bridges' mapper, and `V-MAP` all use this strategy (120, 125, 116, 139). The remaining mappers use succinct text indexes. `VG` uses the `GCSA2` (130) and a longest-common-prefix array, which enable very specific queries at the expense of high memory utilization (48). `HISAT2` uses a modified `GCSA` (132) that also encodes the graph structure itself. This helps give `HISAT2` an impressively low memory footprint but a somewhat more limited set of queries (72). `GRAPHALIGNER` also has the option of seeding with a full text index of the node sequences of the graph, but this is not enabled by default (120).

Most graph mappers employ graph-based alignment algorithms. The exceptions are `GENOME_MAPPER`, which aligns to all paths out from a seed, and `HISAT2`. The `HISAT2` alignment algorithm relies on a complex set of heuristics that depend heavily on its exact match index. This makes it exceptionally fast, although it also can hurt alignment quality around indels. `VG` and `V-MAP` both employ some version of partial order alignment (48, 139). Seven Bridges first searches for a near-exact match using an exponential depth-first search, applying partial order alignment if this fails (116).

Due to the recent development of these methods, there have been few independent comparative studies of their performance and accuracy. In general, `VG` compares favorably to other tools in terms of accuracy on NGS data (117). However, it requires more memory, has slower indexing, and often maps more slowly than the alternatives (72, 139). `V-MAP`'s fast clustering heuristics allow it to align long reads faster than `VG`, but it was not compared to `GRAPHALIGNER` (139). `GRAPHALIGNER` is the only mapper to incorporate the most recent research into graph alignment algorithms. It uses a banded alignment algorithm to achieve impressive speed aligning long reads to genome graphs (120).

RNA splicing can be represented directly in a genome graph model (78). It follows that graph mappers can be applied to RNA sequencing data. `HISAT2` can map RNA-seq data in addition to genomic DNA (72). It is based on the RNA mapper `HISAT` (71), and it retains the capacity for spliced alignment. The ability to create spliced variation graphs has also been added to `VG` (51). In these variation graphs, known splice junctions are added as edges, similar to the addition of a deletion event. `VG` supports any type of variation, but its splicing-awareness is limited to splice junctions represented in the graph. Thus, reads that span a novel splice junction will only map partially. We further discuss applications of graph mapping to functional genomics in §8.3.

7. Pangenomic data formats

A number of common data formats are used to exchange pangenomic models. Pangenomes can be stored as collections of sequences in `FASTA` format. Variant calls in `VCF` format (33) may be added to such a collection to describe small or structural variants found in the pangenome. However, to exchange graphical pangenomes, the community frequently uses a

subset of version 1 of the Graphical Fragment Assembly format (GFAv1)⁷. Only a small subset of GFA is required to represent pangenome graphs, but using this format allows pangenomic analyses to use many genome assembly tools.

To represent read alignments to pangenome graphs, the VG toolkit has developed the GAM format (51), which generalizes the SAM/BAM (60) data model to pangenome graphs. GAM is produced by several other alignment tools (120, 68), and consumed by numerous downstream applications. The Graph Alignment Format (GAF)⁸ generalizes the text-based Pairwise Alignment Format (PAF)⁹ to work on graphs encoded in GFA. GAF can describe mappings to graphs encoded in rGFA¹⁰, which is a specialization of GFA for reference pangenome graphs.

8. Applications of pangenomic models

Although graphical pangenomic techniques can be applied throughout biology, most recent work has focused on a handful of applications where they can provide substantial benefits. Reduced reference bias and coherent representations of alleles yields significant improvements in structural variation detection and can decrease the runtime costs of genotyping. They have enabled haplotype reconstruction of quasispecies in metagenomic studies and transcripts in functional genomics, and supported pangenomic association studies.

8.1. Variant calling and genotyping

Typically, *variant calling* and *genotyping* indicate different aspects of a reference-guided genome inference process. Genotyping consists of determining whether a previously observed variant is present in a new sample, whereas variant calling involves detecting yet unobserved variation. When our reference system is a linear genome, these two steps are often merged. A single process will detect candidate variation and infer a sample genotype at each putatively variable locus (82, 49).

These methods are often Bayesian, and combine a model for observation and error with a simple *a priori* model of the pattern of variation that we expect to see, typically based on the expected rate of heterozygosity or mutation. By using *multisample* variant calling, genotype phasing, and genotype imputation (1, 18), we can share information between samples to improve the accuracy of our reconstruction of genomes from short reads. However, this joint calling approach is expensive, and not applicable when we only have a few new genomes to reconstruct. Furthermore, it does not help our primary interpretation of new sequencing data during read mapping.

Alignment to any allele in a pangenome reference graph is as efficient as alignment to the reference allele in a linear reference sequence. Including known genetic variation in our pangenomic reference system can thus reduce bias towards the reference allele when

⁷<https://github.com/GFA-spec/GFA-spec/blob/master/GFA1.md>

⁸<https://github.com/lh3/gfatools/blob/master/doc/rGFA.md#the-graph-alignment-format-gaf>

⁹<https://github.com/lh3/miniasm/blob/master/PAF.md>

¹⁰<https://github.com/lh3/gfatools/blob/master/doc/rGFA.md>

genotyping heterozygous variants (Figure 3). This effect is strongest for alleles that are highly divergent from the reference, suggesting that the impact of these approaches on variant calling and genotyping will be strongest for large indels and structural variants.

8.1.1. Sample-specific references.—Rather than implementing genotyping or variant calling over a pangenome model, the model can be used to infer a likely haploid version of a new sample's sequence, which is used as a reference genome for variant calling. These methods show incremental improvements in accuracy over conventional methods. GRAMTOOLS (88) and PANVC (140) both use specialized pangenome indexes to map reads for this purpose. PRG uses read k -mers from reads to choose likely haplotype paths through a graph (35).

8.1.2. Small variants.—In many ways, small variants stand to benefit the least from pangenomic variant calling and genotyping. NGS read lengths are sufficient to span their entirety, and the associated variant calling algorithms are quite mature. However, reference bias in mapping can be a source of small variant calling error, particularly for indels.

One strategy consists of realigning reads to graphs of known variation after mapping to a linear reference. This approach was pioneered in the 1000 Genomes Project, which applied GLIA to establish genotype likelihoods for indels and complex alleles (1). GRAPH TYPER refines this approach to achieve competitive accuracy in joint genotyping large cohorts with very low computational costs, providing improved genotyping performance at known variable sites described in its graph (40).

Colored de Bruijn graphs support pangenomic, reference-free variant calling. CORTEX calls variants based on coverage in bubble structures in a colored de Bruijn graph, which is constructed from multiple read sets and/or reference genomes (66). BUBBLEPARSE extends CORTEX's model to improve SNP discovery (81).

8.1.3. Genotyping structural variation.—The study of structural variants (SVs)—typically defined as variants affecting at least 50 bp—has more to gain from using genome graphs. SVs are difficult to call with NGS reads because they are large relative to the read length. Long read sequencing does not share this difficulty, but this technology remains prohibitively expensive for population-scale studies or routine use.

BAYESTYPER (129) compares the distribution of k -mers from sequencing reads to the distribution of k -mers along paths in the graph. It calls structural variants with high accuracy almost irrespective of the size of the variant. However, it has a high memory footprint, and later analysis has also suggested that its reliance on long exact matches makes it fragile to breakpoint uncertainty (62).

In contrast, PARAGRAPH (26), GRAPH TYPER2 (41), and VG CALL (62) use genome graphs to genotype SVs. The largest difference is that PARAGRAPH and GRAPH TYPER2 first map to a linear reference, then locally realign to regional graphs, whereas VG maps reads directly to a whole genome graph. These methods all use read coverage to determine the genotype, and significantly outperform competing reference-based methods.

Some pangenomic methods have sought improved accuracy and efficiency by focusing on specific regions where reference bias makes inference challenging. The highly polymorphic and medically important human leukocyte antigen (HLA) genes have received an especially large amount of attention. HLA*LA (36), KOURAMI (79), and HISAT-GENOTYPE (72) have all demonstrated techniques for genotyping HLA genes by aligning NGS reads to a graph encoding various HLA alleles. Their results rival gold-standard Sanger sequencing methods in accuracy. EXPANSIONHUNTER (37) and HISAT-GENOTYPE (72) used similar methodologies to achieve comparable or better accuracy than existing methods for short tandem repeats (STRs).

8.2. Inferring precise haplotypes from pangenomes and metagenomes

Current assembly approaches often produce a result that mixes both haplotypes of a diploid genome together. This inaccurate representation has led to the development of diploid assembly methods. WHDENOVO (47) addresses this by using long reads to infer phase within the pangenomic space of the assembly graph. In an alternative approach to haplotype inference, pangenomic error correction methods can be used to clean small errors from long reads, rendering them precise observations of long haplotypes (122, 120). These methods build an assembly from accurate reads (NGS or PacBio circular consensus), to which reads can be aligned. The path of the alignment through the graph is taken as the corrected read.

Several methods have used pangenome graphs to support haplotype reconstruction, but in the context of a mixed population of related genomes sampled from a metagenome or quasispecies mixture. MYKROBE predictor (17) and GROOT (121) use graph-based structures of bacterial genomes and gene sets to predict antibody resistance in sequencing samples. META KALLISTO (123) performs taxonomic classification and quantification of metagenomic sequencing data using a database of known sequences represented as colored de Bruijn graphs. VIRUS-VG (10) builds a variation graph from assembled viral contigs in order to construct haplotypes and predict associated abundances in viral quasispecies from sequencing reads. This method was later improved in VG-Flow (9) which can scale to much larger genomes, such as bacteria.

8.3. Functional pangenomics

The effects of reference bias on analyses of allele-specific protein binding and transcription have led to ample interest in pangenomic techniques. Reference bias may also affect our ability to associate genome with phenotype in association mapping. Here, we consider ways in which pangenomic models can improve the accuracy of such assays into genome function.

8.3.1. CHiP-seq peak calling.—CHiP-seq data is mapped back to the reference genome in order to locate protein binding sites. GRAPH PEAK CALLER is based on VG and is the first tool to use a genome graph for this process (56). It was shown to find binding sites more enriched for known DNA binding motifs than linear methods on *A. thaliana*. It was also applied to human data to discover novel sites for enhancers (55).

8.3.2. Transcriptomics.—Some transcriptomic analyses are strongly affected by reference bias. Chief among these is allele-specific expression (ASE) (34, 134, 23). ASE analysis estimates the expression levels of genes or transcripts on each allele separately by comparing the number of RNA sequencing (RNA-seq) reads mapped to the two different alleles of heterozygous variants. A mapping bias in favor of one of the alleles can therefore create illusory differences in expression between the alleles. Using variation information during mapping can help ameliorate this and improve estimates of ASE (23, 91).

The simplest approach to using variation data in mapping involves creating a personalized diploid genome or transcriptome, which is then used as the reference for a standard linear mapping method (114). Methods using this approach have been shown to reduce reference bias, but they require diploid reconstruction of the genome in question. Variant-aware mappers like GSNAP (143), iMapSPlice (87), ASELux (91), and HISAT2 (72) remove this necessity, and have been shown to reduce reference bias at known SNVs during mapping (23, 87). Variation-aware analysis of RNA-seq data is also important for accurately analyzing highly polymorphic regions, such as HLA. ALTHAPALIGNR (80) and HLAPEPERS (2) compare reads against a collection of known HLA haplotypes yielding improved estimation of HLA expression.

8.3.3. Pangenomic association studies.—Pangenome-wide association studies (PWAS) generalize the genome-wide association (GWAS) concept to pangenomes. This new area has primarily used traditional pangenome definitions from microbiology (19). Recent work applies PWAS specifically to pangenome graphs. In the frequented region (FR) technique, a syntenic region finding algorithm similar to those used in whole genome alignment (§4.4) detects regions of a compacted DBG that are shared between many individuals (31), and these are employed as features in PWAS (89). When tested in 100 yeast strains, this approach marginally improves on standard GWAS techniques, but for some phenotypes provides a dramatic improvement in performance.

9. Discussion

In the near future, we expect complete, haplotype-resolved, “telomere-to-telomere” assemblies of large genomes to be readily obtained at low cost (92). The impeding resolution of the genome assembly problem raises new issues. To make full use of genome assemblies, we must relate them to each other. And, to maximize their value, we must make it possible to use the prior information contained in them to guide subsequent genomic analyses.

These goals drive us to work with the pangenome implied by a collection of whole genome assemblies. Pangenomes can be modeled as simple collections of DNA sequences, but this can obscure variation between genomes that is essential to unlocking insight into biology. Increasingly, researchers have explored pangenome graphs which represent both sequences and variation between them. These methods are flexible. In representing the mutual alignment of many genomes, a graphical pangenome can contain, and relate between many linear reference systems. They are also scalable. Recently-developed methods support the

compact storage and query of collections of tens of thousands of genomes. And, they can improve alignment and genotyping accuracy in the context of known variation.

However, adding variation to the reference system is not without potential drawbacks. Model construction, indexing, and alignment steps typically require more time for pangenome graphs than linear reference genomes. Additional information can increase ambiguity, and care must be taken to build models that improve utility by including relevant variation. Working with a graphical reference system necessitates knowledge of graph-theoretic concepts that may be unfamiliar to many biologists. Users also must consider that, pangenome graphs are not observable in the same sense that a given genome is. Their construction is often guided more by application than a clear ground truth.

Due to these issues, some argue that it is likely that linear genomic models will remain important into the future (128). Our survey does not disagree with this possibility. Many of the works we have considered foresee a future in which reference systems are graphical, but only a handful (primarily those based on variation graphs) produce alignments or genotype calls in the context of a pangenome graph. Linear or hierarchical coordinate systems for the pangenome may be preferred by the genomics community. If so, these reference systems are likely to proliferate as we explore the pangenome of humans and other species.

Of course, a future full of many reference genomes is essentially a pangenomic one. Whether or not the community fosters the development of standardized pangenomic reference models, the proliferation of whole genome sequences will only increase the importance of the methods that we have considered here. Pangenome graphs provide a distributed framework which we can use to bring many reference systems into the same analytical context. We can use them to build reference models optimized for particular research or clinical settings, potentially mixing public and private sources of data, without sacrificing our ability to relate our findings to standard reference models. Provided their continued improvement, graphical pangenomic methods will be well-suited to the pluralistic, decentralized attributes of a future in which genomes are easily sequenced and assembled.

ACKNOWLEDGMENTS

Research reported in this publication was supported the National Institutes of Health under Award Numbers U54HG007990, U01HL137183 and 2U41HG007234. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of National Institute of Health. The research was also made possible by the generous financial support of the W.M. Keck Foundation (DT06172015). T.M. acknowledges funding from the German Federal Ministry for Research and Education (BMBF 031L0184). The work of J.A.S. was supported by the Carlsberg Foundation. S.H. acknowledges funding from the central innovation programme (ZIM) for SMEs of the Federal Ministry for Economic Affairs and Energy of Germany. J.D.S. thanks BBSRC for funding BB/S004661/1. We would also like to thank all the attendees of the joint NCBI/UCSC Pangenomics Hackathon which took place at UCSC in the spring of 2019 and which spurred many conversations that contributed to this review.

LITERATURE CITED

1. 1000 Genomes Project Consortium, et al. 2015. A global reference for human genetic variation. *Nature* 526:68 [PubMed: 26432245]
2. Aguiar VRC, César J, Delaneau O, Dermitzakis ET, Meyer D. 2019. Expression estimation and eQTL mapping for HLA genes with a personalized pipeline. *PLoS Genet.* 15:e1008091 [PubMed: 31009447]

3. Alonso-Blanco C, Andrade J, Becker C, Bemm F, Bergelson J, et al. 2016. 1,135 genomes reveal the global pattern of polymorphism in *Arabidopsis thaliana*. *Cell* 166:481–491 [PubMed: 27293186]
4. Ambler JM, Mulaudzi S, Mulder N. 2019. GenGraph: a python module for the simple generation and manipulation of genome graphs. *Bioinformatics* 20
5. Amir A, Lewenstein M, Lewenstein N. 1997. Pattern matching in hypertext. In *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 160–173
6. Antipov D, Korobeynikov A, McLean JS, Pevzner PA. 2015. hybridSPAdes: an algorithm for hybrid assembly of short and long reads. *Bioinformatics* 32:1009–1015 [PubMed: 26589280]
7. Armstrong J, Hickey G, Diekhans M, Deran A, Fang Q, et al. 2019. Progressive alignment with cactus: a multiple-genome aligner for the thousand-genome era. *bioRxiv preprint* :730531
8. Audano PA, Sulovari A, Graves-Lindsay TA, Cantsilieris S, Sorensen M, et al. 2019. Characterizing the major structural variant alleles of the human genome. *Cell* 176:663–675.e19 [PubMed: 30661756]
9. Baaijens JA, Stougie L, Schönhuth A. 2019. Strain-aware assembly of genomes from mixed samples using variation graphs. *bioRxiv* :645721
10. Baaijens JA, Van der Roest B, Köster J, Stougie L, Schönhuth A. 2019. Full-length de novo viral quasispecies assembly through variation graph construction. *Bioinformatics* :287177
11. Baier U, Beller T, Ohlebusch E. 2015. Graphical pan-genome analysis with compressed suffix trees and the burrows–wheeler transform. *Bioinformatics* 32:497–504 [PubMed: 26504144]
12. Bernardini G, Pisanti N, Pissis SP, Rosone G. 2019. Approximate pattern matching on elastic-degenerate text. *Theoretical Computer Science*
13. Beyer W, Novak AM, Hickey G, Chan J, Tan V, et al. 2019. Sequence tube maps: making graph genomes intuitive to commuters. *Bioinformatics*
14. Biederstedt E, Oliver JC, Hansen NF, Jajoo A, Dunn N, et al. 2018. NovoGraph: Human genome graph construction from multiple long-read de novo assemblies. *F1000Research* 7:1391 [PubMed: 30613392]
15. Bolger A, Denton A, Bolger M, Usadel B. 2017. Logan: A framework for lossless graph-based analysis of high throughput sequence data. *bioRxiv* :175976
16. Bowe A, Onodera T, Sadakane K, Shibuya T. 2012. Succinct de Bruijn graphs. In *Proceedings of the 12th Workshop on Algorithms in Bioinformatics (WABI 2012)*, vol. 7534 of LNCS. Springer
17. Bradley P, Gordon NC, Walker TM, Dunn L, Heys S, et al. 2015. Rapid antibiotic-resistance predictions from genome sequence data for *staphylococcus aureus* and *mycobacterium tuberculosis*. *Nature Communications* 6:10063
18. Browning SR, Browning BL. 2011. Haplotype phasing: existing methods and new developments. *Nature Reviews Genetics* 12:703
19. Brynildsrud O, Bohlin J, Scheffer L, Eldholm V. 2016. Rapid scoring of genes in microbial pan-genome-wide association studies with Scoary. *Genome Biology* 17:238 [PubMed: 27887642]
20. Burrows M, Wheeler DJ. 1994. A block sorting lossless data compression algorithm. *Tech. Rep.* 124, Digital Equipment Corporation
21. Büchler T, Ohlebusch E. 2019. An improved encoding of genetic variation in a Burrows-Wheeler transform. *bioRxiv* :658716
22. Cao J, Schneeberger K, Ossowski S, Günther T, Bender S, et al. 2011. Whole-genome sequencing of multiple *Arabidopsis thaliana* populations. *Nature Genetics* 43:956–963 [PubMed: 21874002]
23. Castel SE, Levy-Moonshine A, Mohammadi P, Banks E, Lappalainen T. 2015. Tools and best practices for data processing in allelic expression analysis. *Genome Biology* 16:195 [PubMed: 26381377]
24. Chaisson MJP, Sanders AD, Zhao X, Malhotra A, Porubsky D, et al. 2019. Multi-platform discovery of haplotype-resolved structural variation in human genomes. *Nature Communications* 10:1784
25. Chakraborty M, VanKuren NW, Zhao R, Zhang X, Kalsow S, Emerson J. 2018. Hidden genetic variation shapes the structure of functional elements in *drosophila*. *Nature Genetics* 50:20–25 [PubMed: 29255259]

26. Chen S, Krusche P, Dolzhenko E, Sherman RM, Petrovski R, et al. 2019. Paragraph: A graph-based structural variant genotyper for short-read sequence data. *bioRxiv* :635011
27. Chimani M, Gutwenger C, Jünger M, Klau G, Klein K, Mutzel P. 2012. The Open Graph Drawing Framework (OGDF). *Handbook of Graph Drawing and Visualization*
28. Church DM, Schneider VA, Steinberg KM, Schatz MC, Quinlan AR, et al. 2015. Extending reference assembly models. *Genome Biology* 16:13 [PubMed: 25651527]
29. Cislak A, Grabowski S, Holub J. 2018. SOPanG: online text searching over a pan-genome. *Bioinformatics* 34:4290–4292 [PubMed: 29939210]
30. Claude F, Navarro G, Ordóñez A. 2015. The wavelet matrix: An efficient wavelet tree for large alphabets. *Information Systems* 47:15–32
31. Cleary A, Ramaraj T, Kahanda I, Mudge J, Mumeby B. 2018. Exploring frequented regions in pan-genomic graphs. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 16:1424–1435 [PubMed: 30106690]
32. Computational pan-genomics consortium. 2016. Computational pan-genomics: status, promises and challenges. *Briefings in Bioinformatics* 19:118–135
33. Danecek P, Auton A, Abecasis G, Albers CA, Banks E, et al. 2011. The variant call format and VCFtools. *Bioinformatics* 27:2156–2158 [PubMed: 21653522]
34. Degner JF, Marioni JC, Pai AA, Pickrell JK, Nkadori E, et al. 2009. Effect of read-mapping biases on detecting allele-specific expression from RNA-sequencing data. *Bioinformatics* 25:3207–3212 [PubMed: 19808877]
35. Dilthey A, Cox C, Iqbal Z, Nelson MR, McVean G. 2015. Improved genome inference in the MHC using a population reference graph. *Nature Genetics* 47:682 [PubMed: 25915597]
36. Dilthey AT, Mentzer AJ, Carapito R, Cutland C, Cereb N, et al. 2019. HLA* LA—HLA typing from linearly projected graph alignments. *Bioinformatics* 35:4394–4396 [PubMed: 30942877]
37. Dolzhenko E, Deshpande V, Schlesinger F, Krusche P, Petrovski R, et al. 2019. Expansion-Hunter: A sequence-graph based tool to analyze variation in short tandem repeat regions. *bioRxiv* :572545
38. Duan Z, Qiao Y, Lu J, Lu H, Zhang W, et al. 2019. HUPAN: a pan-genome analysis pipeline for human genomes. *Genome Biology* 20:149 [PubMed: 31366358]
39. Durbin R. 2014. Efficient haplotype matching and storage using the positional burrows-wheeler transform (pbwt). *Bioinformatics* 30:1266–1272 [PubMed: 24413527]
40. Eggertsson HP, Jonsson H, Kristmundsdóttir S, Hjartarson E, Kehr B, et al. 2017. Graph typer enables population-scale genotyping using pangenome graphs. *Nature Genetics* 49:1654 [PubMed: 28945251]
41. Eggertsson HP, Kristmundsdóttir S, Beyter D, Jonsson H, Skuladóttir A, et al. 2019. Graph typer2 enables population-scale genotyping of structural variation using pangenome graphs. *Nature Communications* 10
42. Ferragina P, Manzini G. 2005. Indexing compressed text. *Journal of the ACM* 52:552–581
43. Franz M, Lopes C, Huck G, Dong Y, Sumer O, Bader G. 2016. Cytoscape.js: a graph theory library for visualization and analysis. *Bioinformatics* 32:309–311 [PubMed: 26415722]
44. Fu S, Wang A, Au KF. 2019. A comparative evaluation of hybrid error correction methods for error-prone long reads. *Genome Biology* 20:26 [PubMed: 30717772]
45. Gagie T, Manzini G, Sirén J. 2017. Wheeler graphs: A framework for bwt-based data structures. *Theoretical Computer Science* 698:67–78 [PubMed: 29276331]
46. Gao L, Gonda I, Sun H, Ma Q, Bao K, et al. 2019. The tomato pan-genome uncovers new genes and a rare allele regulating fruit flavor. *Nature Genetics* 51:1044–1051 [PubMed: 31086351]
47. Garg S, Rautiainen M, Novak AM, Garrison E, Durbin R, Marschall T. 2018. A graph-based approach to diploid genome assembly. *Bioinformatics* 34:i105–i114 [PubMed: 29949989]
48. Garrison E. 2019. Graphical pangenomics. Ph.D. thesis, University of Cambridge
49. Garrison E, Marth G. 2012. Haplotype-based variant detection from short-read sequencing. *arXiv*:1207.3907
50. Garrison E, Marth G. 2012. Haplotype-based variant detection from short-read sequencing. *arXiv*:1207.3907

51. Garrison E, Sirén J, Novak AM, Hickey G, Eizenga JM, et al. 2018. Variation graph toolkit improves read mapping by representing genetic variation in the reference. *Nature Biotechnology* 36:875–879
52. Ghaffaari A, Marschall T. 2019. Fully-sensitive seed finding in sequence graphs using a hybrid index. *Bioinformatics* 35:i81–i89 [PubMed: 31510650]
53. Gonnella G, Niehus N, Kurtz S. 2018. GfaViz: flexible and interactive visualization of gfa sequence graphs. *Bioinformatics* 35:2853–2855
54. Grasso C, Lee C. 2004. Combining partial order alignment and progressive multiple sequence alignment increases alignment speed and scalability to very large alignment problems. *Bioinformatics* 20:1546–1556 [PubMed: 14962922]
55. Groza C, Kwan T, Soranzo N, Pastinen T, Bourque G. 2019. Personalized and graph genomes reveal missing signal in epigenomic data. *bioRxiv* :457101
56. Grytten I, Rand KD, Nederbragt AJ, Storvik GO, Glad IK, Sandve GK. 2019. Graph Peak Caller: Calling ChIP-seq peaks on graph-based reference genomes. *PLOS Computational Biology* 15:e1006731 [PubMed: 30779737]
57. Haeussler M, Zweig AS, Tyner C, Speir ML, Rosenbloom KR, et al. 2018. The UCSC Genome Browser database: 2019 update. *Nucleic Acids Research* 47:D853–D858
58. Hehir-Kwa JY, Marschall T, Kloosterman WP, Francioli LC, Baaijens JA, et al. 2016. A high-quality human reference panel reveals the complexity and distribution of genomic structural variants. *Nature Communications* 7:12989
59. Hein J 1989. A new method that simultaneously aligns and reconstructs ancestral sequences for any number of homologous sequences, when the phylogeny is given. *Molecular Biology and Evolution* 6:649–668 [PubMed: 2488477]
60. Heng L, Handsaker B, Wysoker A, Fennell T, Ruan J, et al. 2009. The Sequence Alignment/Map format and SAMtools. *Bioinformatics* 25:2078–2079 [PubMed: 19505943]
61. Heydari M, Miclotte G, Van de Peer Y, Fostier J. 2018. BrownieAligner: accurate alignment of Illumina sequencing data to de Bruijn graphs. *BMC Bioinformatics* 19
62. Hickey G, Heller D, Monlong J, Sibbesen JA, Siren J, et al. 2019. Genotyping structural variants in pangenome graphs using the vg toolkit. *bioRxiv* :654566
63. Holley G, Melsted P. 2019. Bifrost - Highly parallel construction and indexing of colored and compacted de Bruijn graphs. *bioRxiv*
64. Huang L, Popic V, Batzoglou S. 2013. Short read alignment with populations of genomes. *Bioinformatics* 29:i361–i370 [PubMed: 23813006]
65. Huang S, Lam T, Sung W, Tam S, Yiu S. 2010. Indexing similar DNA sequences. In *Proceedings of the The Sixth International Conference on Algorithmic Aspects in Information and Management (AAIM 2010)*, vol. 6124 of LNCS. Springer
66. Iqbal Z, Caccamo M, Turner I, Flicek P, McVean G. 2012. De novo assembly and genotyping of variants using colored de Bruijn graphs. *Nature Genetics* 44:226–232 [PubMed: 22231483]
67. Jain C, Misra S, Zhang H, Dilthey A, Aluru S. 2019. Accelerating sequence alignment to graphs. In *2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE
68. Jain C, Zhang H, Gao Y, Aluru S. 2019. On the complexity of sequence to graph alignment. In *International Conference on Research in Computational Molecular Biology*. Springer
69. Jandrasits C, Dabrowski PW, Fuchs S, Renard BY. 2018. seq-seq-pan: Building a computational pan-genome data structure on whole genome alignment. *BMC Genomics* 19:47 [PubMed: 29334898]
70. Kavya VNS, Tayal K, Srinivasan R, Sivadasan N. 2019. Sequence alignment on directed graphs. *Journal of Computational Biology* 26:53–67 [PubMed: 30204489]
71. Kim D, Langmead B, Salzberg SL. 2015. HISAT: a fast spliced aligner with low memory requirements. *Nature Methods* 12:357–360 [PubMed: 25751142]
72. Kim D, Paggi JM, Park C, Bennett C, Salzberg SL. 2019. Graph-based genome alignment and genotyping with HISAT2 and HISAT-genotype. *Nature Biotechnology* 37:907–915
73. Krzywinski M, Schein J, Birol I, Connors J, Gascoyne R, et al. 2009. Circos: an information aesthetic for comparative genomics. *Genome Res* 19:1639–1645 [PubMed: 19541911]

74. Kunyavskaya O, Prjibelski AD. 2018. SGTk: a toolkit for visualization and assessment of scaffold graphs. *Bioinformatics* 35:2303–2305
75. Kural D 2014. Methods for inter-and intra-species genomics for the detection of variation and function. Ph.D. thesis, Boston College. Graduate School of Arts and Sciences
76. Laing C, Buchanan C, Taboada EN, Zhang Y, Kropinski A, et al. 2010. Pan-genome sequence analysis using Panseq: an online tool for the rapid analysis of core and accessory genomic regions. *BMC Bioinformatics* 11:461 [PubMed: 20843356]
77. Langley SA, Miga KH, Karpen GH, Langley CH. 2019. Haplotypes spanning centromeric regions reveal persistence of large blocks of archaic DNA. *eLife* 8:e42989 [PubMed: 31237235]
78. Lee C, Grasso C, Sharlow MF. 2002. Multiple sequence alignment using partial order graphs. *Bioinformatics* 18:452–464 [PubMed: 11934745]
79. Lee H, Kingsford C. 2018. Kourami: graph-guided assembly for novel human leukocyte antigen allele discovery. *Genome Biology* 19:16 [PubMed: 29415772]
80. Lee W, Plant K, Humburg P, Knight JC. 2018. AltHapAlignR: improved accuracy of RNA-seq analyses through the use of alternative haplotypes. *Bioinformatics* 34:2401–2408 [PubMed: 29514179]
81. Leggett RM, Ramirez-Gonzalez RH, Verweij W, Kawashima CG, Iqbal Z, et al. 2013. Identifying and classifying trait linked polymorphisms in non-reference species by walking coloured de Bruijn graphs. *PLoS One* 8:e60058 [PubMed: 23536903]
82. Li H 2011. A statistical framework for SNP calling, mutation discovery, association mapping and population genetical parameter estimation from sequencing data. *Bioinformatics* 27:2987–2993 [PubMed: 21903627]
83. Li H 2013. Aligning sequence reads, clone sequences and assembly contigs with BWA-MEM. arXiv
84. Li H 2018. Minimap2: pairwise alignment for nucleotide sequences. *Bioinformatics* 34:3094–3100 [PubMed: 29750242]
85. Li R, Li Y, Zheng H, Luo R, Zhu H, et al. 2010. Building the sequence map of the human pan-genome. *Nature Biotechnology* 28:57
86. Linthorst J, Hulsman M, Holstege H, Reinders M. 2015. Scalable multi whole-genome alignment using recursive exact matching. *bioRxiv* :022715
87. Liu X, MacLeod JN, Liu J. 2018. iMapSplice: Alleviating reference bias through personalized RNA-seq alignment. *PLOS ONE* 13:e0201554 [PubMed: 30096157]
88. Maciua S, del Ojo Elias C, McVean G, Iqbal Z. 2016. A natural encoding of genetic variation in a Burrows-Wheeler transform to enable mapping and genome inference. *Algorithms in Bioinformatics* :222–233
89. Manuweera B, Mudge J, Kahanda I, Mumey B, Ramaraj T, Cleary A. 2019. Pangenome-wide association studies with frequented regions. *Proceedings of the 10th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics - BCB '19*
90. Marcus S, Lee H, Schatz MC. 2014. SplitMEM: a graphical algorithm for pan-genome analysis with suffix skips. *Bioinformatics* 30:3476–3483 [PubMed: 25398610]
91. Miao Z, Alvarez M, Pajukanta P, Ko A. 2018. ASElux: an ultra-fast and accurate allelic reads counter. *Bioinformatics* 34:1313–1320 [PubMed: 29186329]
92. Miga KH, Koren S, Rhie A, Vollger MR, Gershman A, et al. 2019. Telomere-to-telomere assembly of a complete human X chromosome. *bioRxiv* :735928
93. Mikheenko A, Kolmogorov M. 2019. Assembly Graph Browser: interactive visualization of assembly graphs. *Bioinformatics* 35:3476–3478 [PubMed: 30715194]
94. Minkin I, Medvedev P. 2019. Scalable multiple whole-genome alignment and locally collinear block construction with SibeliaZ. *BioRxiv* :548123
95. Minkin I, Pham S, Medvedev P. 2016. TwoPaCo: an efficient algorithm to build the compacted de Bruijn graph from many complete genomes. *Bioinformatics* :btw609
96. Mokveld TO, Linthorst J, Al-Ars Z, Reinders M. 2018. CHOP: Haplotype-aware path indexing in population graphs. *bioRxiv* :305268

97. Myers E, Miller W. 1989. Approximate matching of regular expressions. *Bulletin of Mathematical Biology* 51:5–37 [PubMed: 2706401]
98. Myers EW. 2005. The fragment assembly string graph. *Bioinformatics* 21:ii79–ii85 [PubMed: 16204131]
99. Mäkinen V, Navarro G, Sirén J, Välimäki N. 2010. Storage and retrieval of highly repetitive sequence collections. *Journal of Computational Biology* 17:281–308 [PubMed: 20377446]
100. Na JC, Kim H, Min S, Park H, Lecroq T, et al. 2018. FM-index of alignment with gaps. *Theoretical Computer Science* 710:148–157
101. Na JC, Kim H, Park H, Lecroq T, Léonard M, et al. 2016. FM-index of alignment: A compressed index for similar strings. *Theoretical Computer Science* 638:159–170
102. Navarro G. 2000. Improved approximate pattern matching on hypertext. *Theoretical Computer Science* 237:455–463
103. Notredame C, Higgins DG, Heringa J. 2000. T-Coffee: A novel method for fast and accurate multiple sequence alignment. *Journal of molecular biology* 302:205–217 [PubMed: 10964570]
104. Novak AM, Garrison E, Paten B. 2017. A graph extension of the positional Burrows–Wheeler transform and its applications. *Algorithms for Molecular Biology* 12
105. Novak AM, Hickey G, Garrison E, Blum S, Connelly A, et al. 2017. Genome graphs. *bioRxiv* :101378
106. Onodera T, Sadakane K, Shibuya T. 2013. Detecting superbubbles in assembly graphs. In *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 338–348
107. Ou L, Li D, Lv J, Chen W, Zhang Z, et al. 2018. Pan-genome of cultivated pepper (*Capsicum*) and its use in gene presence-absence variation analyses. *New Phytologist* 220:360–363
108. Page AJ, Cummins CA, Hunt M, Wong VK, Reuter S, et al. 2015. Roary: rapid large-scale prokaryote pan genome analysis. *Bioinformatics* 31:3691–3693 [PubMed: 26198102]
109. Paten B, Eizenga JM, Rosen YM, Novak AM, Garrison E, Hickey G. 2018. Superbubbles, ultrabubbles, and cacti. *Journal of Computational Biology* 25:649–663 [PubMed: 29461862]
110. Paten B, Novak AM, Eizenga JM, Garrison E. 2017. Genome graphs and the evolution of genome inference. *Genome Research* 27:665–676 [PubMed: 28360232]
111. Pevzner PA, Tang H, Waterman MS. 2001. An Eulerian path approach to DNA fragment assembly. *Proceedings of the National Academy of Sciences* 98:9748–9753
112. Poplin R, Ruano-Rubio V, DePristo MA, Fennell TJ, Carneiro MO, et al. 2018. Scaling accurate genetic variant discovery to tens of thousands of samples. *bioRxiv* :201178
113. Pritt J, Chen NC, Langmead B. 2018. FORGe: prioritizing variants for graph genomes. *Genome Biology* 19
114. Raghupathy N, Choi K, Vincent MJ, Beane GL, Sheppard KS, et al. 2018. Hierarchical analysis of RNA-seq reads improves the accuracy of allele-specific expression. *Bioinformatics* 34:2177–2184 [PubMed: 29444201]
115. Rahn R, Weese D, Reinert K. 2014. Journalized string tree—a scalable data structure for analyzing thousands of similar genomes on your laptop. *Bioinformatics* 30:3499–3505 [PubMed: 25028723]
116. Rakocevic G, Semenyuk V, Lee WP, Spencer J, Browning J, et al. 2019. Fast and accurate genomic analyses using genome graphs. *Nature Genetics* 51:354–362 [PubMed: 30643257]
117. Rand KD, Grytten I, Nederbragt AJ, Storvik GO, Glad IK, Sandve GK. 2017. Coordinates and intervals in graph-based reference genomes. *BMC Bioinformatics* 18:263 [PubMed: 28521770]
118. Rautiainen M, Mäkinen V, Marschall T. 2019. Bit-parallel sequence-to-graph alignment. *Bioinformatics* 35:3599–3607 [PubMed: 30851095]
119. Rautiainen M, Marschall T. 2017. Aligning sequences to general graphs in $O(V + mE)$ time. *bioRxiv* :216127
120. Rautiainen M, Marschall T. 2019. GraphAligner: Rapid and versatile sequence-to-graph alignment. *bioRxiv* :810812
121. Rowe WPM, Winn MD. 2018. Indexed variation graphs for efficient and accurate resistome profiling. *Bioinformatics* 34:3601–3608 [PubMed: 29762644]

122. Salmela L, Walve R, Rivals E, Ukkonen E. 2016. Accurate self-correction of errors in long reads using de Bruijn graphs. *Bioinformatics* 33:799–806
123. Schaeffer L, Pimentel H, Bray N, Melsted P, Pachter L. 2017. Pseudoalignment for metagenomic read assignment. *Bioinformatics* 33:2082–2088 [PubMed: 28334086]
124. Schmidt D, Colomb R. 2009. A data structure for representing multi-version texts online. *International Journal of Human-Computer Studies* 67:497–514
125. Schneeberger K, Hagmann J, Ossowski S, Warthmann N, Gesing S, et al. 2009. Simultaneous alignment of short reads against multiple genomes. *Genome Biology* 10:R98 [PubMed: 19761611]
126. Shekhezadeh Anari S, de Ridder D, Schranz ME, Smit S. 2018. Efficient inference of homologs in large eukaryotic pan-proteomes. *BMC Bioinformatics* 19
127. Sherman RM, Forman J, Antonescu V, Puiu D, Daya M, et al. 2019. Assembly of a pan-genome from deep sequencing of 910 humans of African descent. *Nature Genetics* 51:30 [PubMed: 30455414]
128. Sherman RM, Salzberg SL. 2020. Pan-genomics in the human genome era. *Nature Reviews Genetics*
129. Sibbesen JA, Maretty L, Krogh A. 2018. Accurate genotyping across variant classes and lengths using variant graphs. *Nature Genetics* 50:1054 [PubMed: 29915429]
130. Sirén J. 2017. Indexing variation graphs. In *Proceedings of the Nineteenth Meeting on Algorithm Engineering and Experiments (ALENEX 2017)*. SIAM
131. Sirén J, Garrison E, Novak AM, Paten B, Durbin R. 2019. Haplotype-aware graph indexes. *Bioinformatics*
132. Sirén J, Välimäki N, Mäkinen V. 2014. Indexing graphs for path queries with applications in genome research. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 11:375–388 [PubMed: 26355784]
133. Smith TF, Waterman MS. 1981. Comparison of biosequences. *Advances in Applied Mathematics* 2:482–489
134. Stevenson KR, Coolon JD, Wittkopp PJ. 2013. Sources of bias in measures of allele-specific expression derived from RNA-seq data aligned to a single reference genome. *BMC Genomics* 14:536 [PubMed: 23919664]
135. Sudmant PH, Rausch T, Gardner EJ, Handsaker RE, Abyzov A, et al. 2015. An integrated map of structural variation in 2,504 human genomes. *Nature* 526:75–81 [PubMed: 26432246]
136. Suzuki H. 2018. dozeu. <https://github.com/ocxtal/dozeu>
137. Thachuk C. 2013. Indexing hypertext. *Journal of Discrete Algorithms* 18:113–122
138. Turner I, Garimella KV, Iqbal Z, McVean G. 2018. Integrating long-range connectivity information into de Bruijn graphs. *Bioinformatics* 34:2556–2565 [PubMed: 29554215]
139. Vaddadi K, Srinivasan R, Sivadasan N. 2019. Read mapping on genome variation graphs :17
140. Valenzuela D, Norri T, Välimäki N, Pitkänen E, Mäkinen V. 2018. Towards pan-genome read alignment to improve variation calling. *BMC Genomics* 19:87 [PubMed: 29764365]
141. Vernikos G, Medini D, Riley DR, Tettelin H. 2015. Ten years of pan-genome analyses. *Current Opinion in Microbiology* 23:148–154 [PubMed: 25483351]
142. Wick RR, Schultz MB, Zobel J, Holt KE. 2015. Bandage: interactive visualization of de novo genome assemblies: Fig. 1. *Bioinformatics* 31:3350–3352 [PubMed: 26099265]
143. Wu TD, Nacu S. 2010. Fast and SNP-tolerant detection of complex variants and splicing in short reads. *Bioinformatics* 26:873–881 [PubMed: 20147302]
144. Yokoyama TT, Sakamoto Y, Seki M, Suzuki Y, Kasahara M. 2019. MoMI-G: Modular Multiscale Integrated Genome Graph Browser. *bioRxiv* :540120
145. Yue JX, Li J, Aigrain L, Hallin J, Persson K, et al. 2017. Contrasting evolutionary genome dynamics between domesticated and wild yeasts. *Nature Genetics* 49:913–924 [PubMed: 28416820]
146. Zhao Y, Wu J, Yang J, Sun S, Xiao J, Yu J. 2011. PGAP: pan-genomes analysis pipeline. *Bioinformatics* 28:416–418 [PubMed: 22130594]

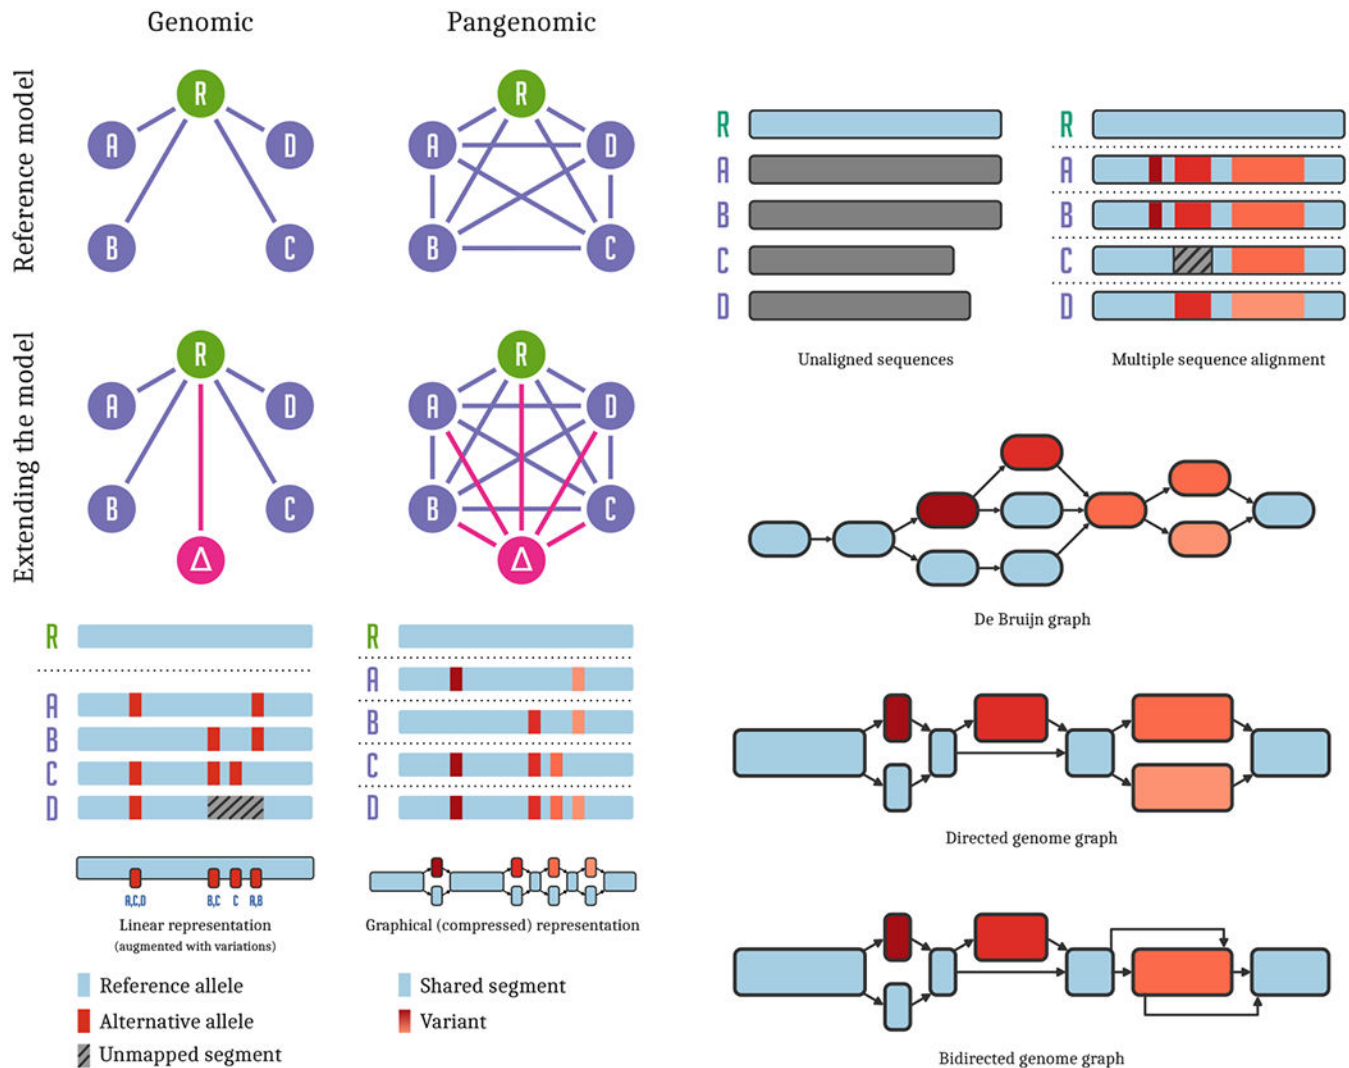
147. Zhou B, Wen S, Wang L, Jin L, Li H, Zhang H. 2017. AntCaller: an accurate variant caller incorporating ancient DNA damage. *Molecular Genetics and Genomics* 292:1419–1430 [PubMed: 28836000]
148. Zook JM, Chapman B, Wang J, Mittelman D, Hofmann O, et al. 2014. Integrating human sequence data sets provides a resource of benchmark SNP and indel genotype calls. *Nature Biotechnology* 32:246

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**Figure 1:**

Pangenomic models. *Left panel: top left:* In reference based genomic analyses, all genomes ($A \dots D$) are compared to each other via their relationship to the reference genome R . *top right:* In a pangenomic setting, we attempt to model direct relationships between all the genomes in our analysis, of which a particular reference R is chosen arbitrarily. *middle left:* When extending our analysis with a new genome, Δ , we add it to the genomic model by comparing it to reference R . *middle right:* In contrast, adding a new genome to a pangenomic analysis compares it directly with all other genomes in the model. *bottom left:* Regions of some genomes are unalignable against the reference, and cannot be represented in a list of variants. *bottom right:* A graphical model of the genomes allows direct all-to-all comparison, capturing all of their sequence relationships. *Right panel: top left:* A collection of sequences representing a pangenome. *top right:* Multiple sequence alignment of the sequences captures their mutual relationships. *middle top:* In a de Bruijn graph, sequences are represented without bias, but variants may correspond to larger graph structures. *middle bottom:* An acyclic sequence graph is equivalent to the multiple sequence alignment. *bottom:* A generic sequence graph can represent a structural variant (in orange, right)

compactly, using edges between the forward and reverse strands of the graph to indicate the presence of an inversion.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

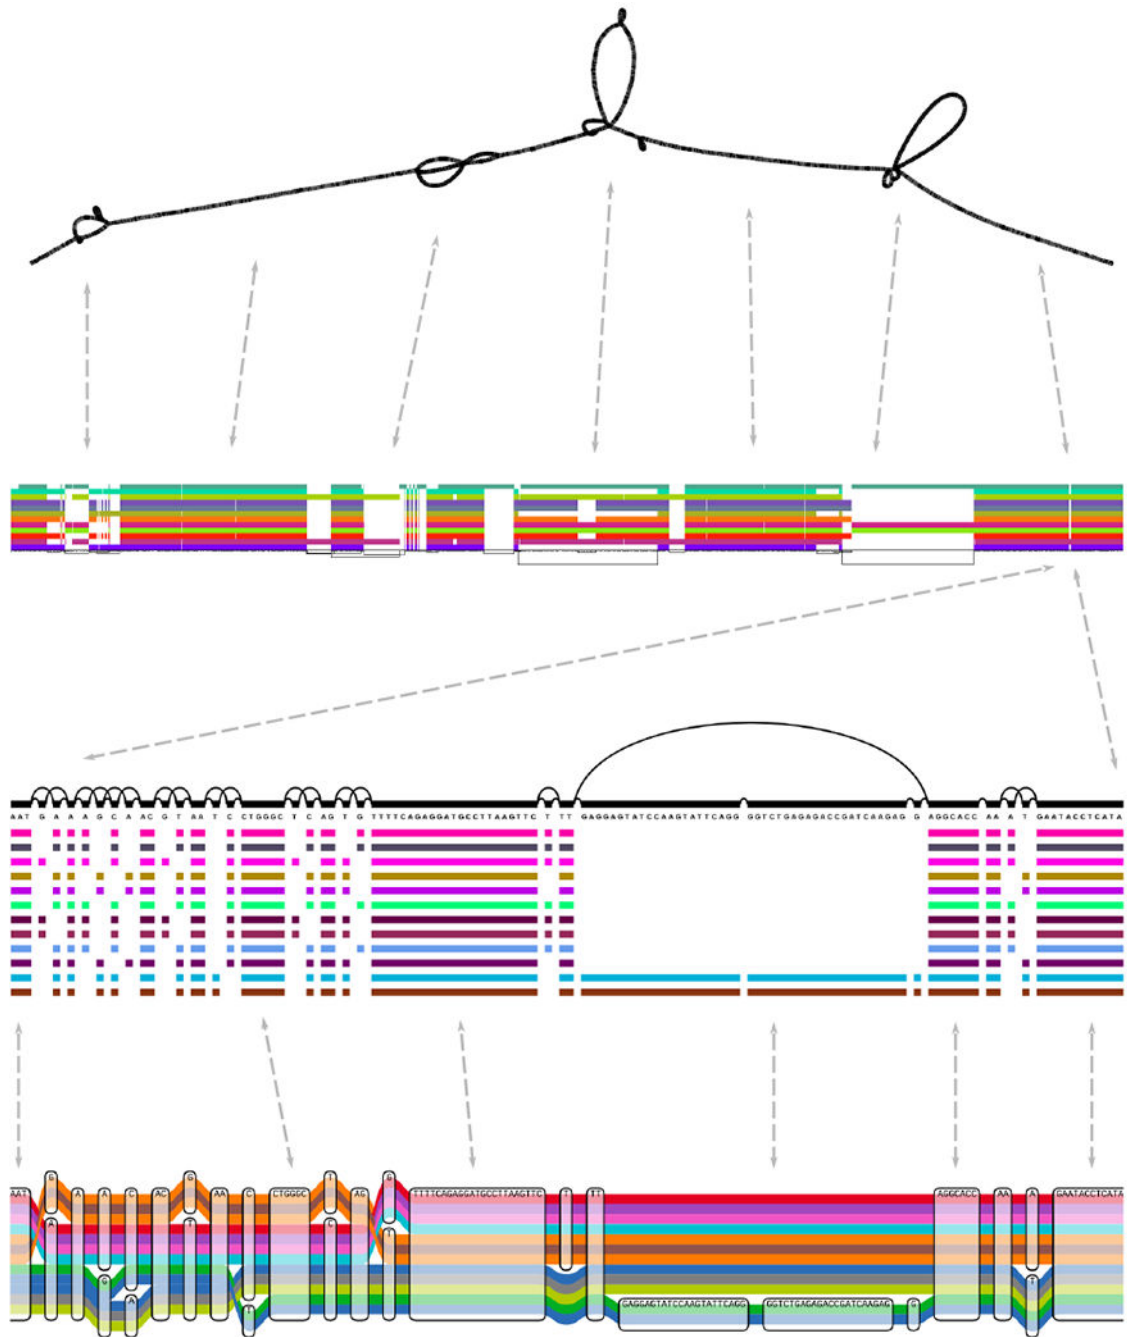


Figure 2:

Visualizing a graph of GRCh38 and its alternate sequences in the gene HLA-DRB1 built with VG MSGA (48). *Top*: BANDAGE's force directed layout reveals large scale structures (142). *Top middle*: ODGI viz's binned, linearized rendering of the paths (colored bars) versus the sequence and topology of the graph (below). *Bottom middle*: A fragment of VG viz's linearized rendering, showing base-level detail. *Bottom*: The same fragment rendered with the Sequence Tube Map (13). Dashed lines show correspondence between the visualizations. Path colors are assigned independently by each method.




Figure 3:
 The mean alternate allele fraction at heterozygous variants in HG002/NA24385 validated in the Genome in a Bottle truth set (148) as a function of deletion or insertion size (SNPs at 0). Error bars are ± 1 s.e.m. Blue points show the allele balance metric across allele lengths for alignments with BWA MEM (83) and variant calls made by FREEBAYES (50). Variant calls were made with alignment to a variation graph built from the 1000 Genomes Project variants (1), followed by variant calling in VG. These are divided into two groups: calls at variants in the graph used for the alignment (red) and those to variants that are private to HG002 (green). Reprinted from (51).

Table 1:

An overview of graph visualization tools. **FD** = Force-directed layout, **Rank** = GraphViz-style rank-based layout, **SM** = Sorted-matrix layout, **A** = Assembly graph, **S** = Scaffold graph, **V** = Variation graph, **B** = Base-level view, **L** = Linear view, **UI** = User interface, **App** = native application, **Web** = browser-based interface, **CLI** = command-line tool.

Tool	Layout	Graph Type	Proven Scale	Extra Views	UI	Backend
Bandage (142)	FD	A	100 Mbp		App	
GfaViz (53)	FD	A, V	1 Mbp		App	OGDF (27)
SGTK (74)	FD	A, S	1 Mbp	B, L	Web	cytoscape.js (43)
AGB (93)	Rank	A	10k edges ^a	L	Web	d3-graphviz ^b
Sequence Tube Map (13)	Tube	V	100 Mbp	B, L	Web	
MoMI-G (144)	Tube, Circos	V	1 Gbp	B, L	Web	Sequence Tube Map (13), Circos (73)
vg view (51)	Rank	V	10 Kbp	B	CLI	GraphViz
vg viz (48)	SM	V	100 Kbp	B, L	CLI	
odgi viz	SM	V	1 Gbp	B, L	CLI	

^ahttps://github.com/almiheenko/almiheenko.github.io/blob/8f4b2f8c/AGB/Flye_Human/data/repeat_graph.jsona

^b<https://github.com/magjac/d3-graphviz>

Table 2:

An overview of graph mapping tools

Tool	Graph Types	Sequencing Types	Other Notes
deBGA	De Bruijn graph	NGS	
BGREAT			Gapless alignment
BrownieAligner			
GenomeMapper	Acyclic variation graph	NGS	No longer maintained
Graph Genome Aligner			
HISAT2			Fast
V-MAP		NGS and long read	
VG	Variation graph	NGS and long read	Accurate, high memory usage
GraphAligner	Variation graph or overlap graph	Long read	Fast, high memory usage