

Building pangenome graphs

Erik Garrison^{1*†}, Andrea Guarracino^{1,2†}, Simon Heumos^{3,4}, Flavia Villani¹, Zhigui Bao⁵, Lorenzo Tattini^{6,7}, Jörg Hagmann⁸, Sebastian Vorbrugg⁹, Santiago Marco-Sola^{10,11}, Christian Kubica⁹, David G. Ashbrook¹, Kaisa Thorell¹², Rachel L. Rusholme-Pilcher¹³, Gianni Liti^{6,7}, Emilio Rudbeck¹⁴, Sven Nahnsen^{3,4}, Zuyu Yang¹⁵, Mwaniki N. Moses¹⁶, Franklin L. Nobrega¹⁷, Yi Wu¹⁷, Hao Chen¹⁸, Joep de Ligt¹⁵, Peter H. Sudmant¹⁹, Nicole Soranzo^{20,21,22,23,2}, Vincenza Colonna^{1,24}, Robert W. Williams¹ and Pjotr Prins¹

¹Department of Genetics, Genomics and Informatics, University of Tennessee Health Science Center, 71 S Manassas St, Memphis, 38163, Tennessee, USA. ²Fondazione Human Technopole, Viale Rita Levi Montalcini, 20157 Milan, Italy. ³Quantitative Biology Center (QBC) Tübingen, University of Tübingen, Tübingen, Germany. ⁴Biomedical Data Science, Dept. of Computer Science, University of Tübingen, Tübingen, Germany. ⁵Shenzhen Branch, Guangdong Laboratory of Lingnan Modern Agriculture, Genome Analysis Laboratory of the Ministry of Agriculture and Rural Affairs, Agricultural Genomics Institute at Shenzhen, Chinese Academy of Agricultural Sciences, Buxin Road 97, Shenzhen, 518120, Guangdong, China. ⁶Institute for Research on Cancer and Aging, Nice (IRCAN), Nice, France. ⁷CNRS UMR 7284, INSERM U 1081, Université Côte d'Azur (UCA), Nice, France. ⁸Computomics GmbH, Eisenbahnstr. 1, 72072 Tübingen, Baden-Württemberg, Germany. ⁹Department of Molecular Biology, Max Planck Institute for Biology, Max-Planck-Ring 9, 72076 Tübingen, Baden-Wuerttemberg, Germany. ¹⁰Computer Sciences Department, Barcelona Supercomputing Center, Barcelona 08034, Spain. ¹¹Departament d'Arquitectura de Computadors i Sistemes Operatius, Universitat Autònoma de Barcelona, Barcelona 08193, Spain. ¹²Chemistry and Molecular Biology, Faculty of Science, University of Gothenburg, Sweden. ¹³Earlham Institute, Norwich Research Park, Colney Lane, Norwich, Norfolk, NR4 7UZ, UK. ¹⁴Clinical Genomics Gothenburg, Bioinformatics and Data Centre, University of Gothenburg, Sweden. ¹⁵The Institute of Environmental Science and Research, New Zealand. ¹⁶Department of Computer Science, University of Pisa. ¹⁷School of Biological Sciences, Faculty of Environmental and Life Sciences, University of Southampton, Southampton, UK. ¹⁸Department of Pharmacology, Addiction Science and Toxicology, University of Tennessee Health Science Center, Memphis, TN. ¹⁹Department of Integrative Biology, University of California Berkeley, Berkeley, CA. ²⁰Wellcome Sanger Institute, Genome Campus, Hinxton CB10 1SA, UK. ²¹National Institute for Health Research Blood and Transplant Research Unit in Donor Health and Genomics, University of Cambridge, Cambridge, UK. ²²Department of Haematology, Cambridge Biomedical Campus, Cambridge CB2 0AW, UK. ²³British Heart Foundation Centre of Research Excellence, University of Cambridge, Cambridge, UK. ²⁴Institute of Genetics and Biophysics, National Research Council, Naples 80111, Italy.

*Corresponding author. E-mail: egarris5@uthsc.edu;

†These authors contributed equally to this work.

Abstract

Pangenome graphs can represent all variation between multiple genomes, but existing methods for constructing them are biased due to reference-guided approaches. In response, we have developed PanGenome Graph Builder (PGGB), a reference-free pipeline for constructing unbiased pangenome graphs. PGGB uses all-to-all whole-genome alignments and learned graph embeddings to build and iteratively refine a model in which we can identify variation, measure conservation, detect recombination events, and infer phylogenetic relationships.

Keywords: pangenomes, genome alignment, variant detection, comparative genomics, chromosome evolution, phylogenetics, population genetics

Pangenome graphs are powerful models that can represent all genetic variation between multiple genomes in a population [1–4]. They allow us to identify variation, measure conservation, detect recombination events, and infer phylogenetic relationships, making them valuable tools for studying sequence evolution and variation in diverse species [5, 6]. However, existing methods for constructing pangenome graphs [7, 8] are biased due to their reference and tree-guided approaches [5, 9], which can lead to incomplete and unstable representations of genetic variation [10]. Meanwhile, although approaches for unbiased pangenome graph construction have been proposed [10, 11], these have been limited to the graph induction step, while experience shows that specialized techniques for pangenome alignment and refinement are required to obtain high-quality pangenome builds [12]. Bias results from attempts to tackle the inevitable computational complexity that arises when building pangenome graphs, which imply all-to-all comparisons that scale quadratically with the number of included genomes [5, 7], or from a goal to structure the resulting graphs so that they are easier to use during read alignment [8].

To overcome these limitations, we propose the PanGenome Graph Builder (PGGB), a reference-free pipeline to construct unbiased pangenome graphs. Its output presents a base-level representation of the pangenome, including variants of all scales from single nucleotide polymorphisms (SNPs) to structural variants (SVs). The graph is unbiased—all genomes are treated equivalently, without input order or phylogenetic dependencies, and any genome may be used as a frame of reference in downstream analysis. PGGB makes no assumptions about phylogenetic relationships, orthology groups, or evolutionary histories, allowing data to speak for itself without risk of implicit bias that may affect inference made on the graph. PGGB is implemented as a modular shell script, integrating independent components via standard text-based file formats, and in this provides a template for future pangenome construction methods. The method is practical, scalable to hundreds of genomes, and has been proven accurate through years of development in the Human Pangenome Reference Consortium (HPRC) [12, 16] and in the broader bioinformatics community [17–19]. Here, we describe the specific innovations in the three main phases of the algorithm—alignment, graph creation, and graph normalization—which unlock this result. We then use cross-validation with MUMMER4 [20] to demonstrate the accuracy of our approach across a wide range of species and problem sizes.

The first step in the PGGB pipeline is sequence alignment. To avoid reference and order bias, PGGB uses an all-to-all alignment of the input sequences. This approach aligns sequences directly to each other, enabling each sequence in the pangenome to serve as a potential reference for exposing all related variation. To obtain alignments, PGGB applies WFMASH [21], which applies a generalization of the bidirectional Wavefront Algorithm (BiWFA) [22, 23] to generate base-level alignments for homology mappings obtained with an extension of MASHMAP [24]. MASHMAP provides highly efficient, accurate detection of homologies among genomes [24] and even across whole




Fig. 1 PGGB and its immediate downstream applications. (A) Visualization of PGGB’s algorithms and data flows. The primary data flow (red) progresses from FASTA to alignment, graph induction, and graph “smoothing”, then normalization with GFAFFIX, yielding the final graph (orange). Optional downstream outputs (blue) include statistics, variant calls (on multiple references), and graph visualizations in 1D and 2D. (B) A 1D visualization of a pangenome graph built from 14 haplotype-resolved, complete (T2T) primate assemblies homologous to human chromosome 6 (all-vs-all alignments shown above in A). Regions of paths that are oriented in the same direction as T2T-CHM13 shown in black, while those in the reverse complement orientation are shown in red. T2T-CHM13 annotations of the MHC, p-arm, centromere, q-arm have been injected into the graph as paths, highlighting features of the structural evolution of this chromosome. The p-arm region containing the MHC is inverted in Gibbon relative to other species, with the exception of chimpanzee and bonobo and between the orangutans. (C) A 2D visualization, rendered with the same human chromosomal annotations in GFAESTUS [13], suggests that structural variation involving subtelomeric ends has caused the graph to circularize. (D) Using ODGI [6], we extract a pairwise distance matrix based on in-graph Jaccard metrics over shared base pairs. This distance matrix yields a phylogenetic tree that matches previous results based on SNPs [14]. We posit that the greater phylogenetic distances reflect the inclusion of the centromeres—which tend to evolve rapidly by near-clonal evolution [15]—in our distance computation.

pangenomes [16]. However, the use of WFMASH is not required, and PGGB supports the use of any user-defined input alignment set in PAF format.

Pangenome graph induction takes a collection of genomes and alignments between them and converts them into an equivalent variation graph. We achieve this with SEQWISH [10], an approach that scales graph induction to terabase-pair scales via a series of disk-backed processing steps. Any single input genome is faithfully and fully embedded in a graph and can be completely extracted by tracing labeled paths through the nodes. The SEQWISH graph recovers transitive homology relationships that may not be present in the initial alignment set. This property allows us to apply random sparsification to reduce the computational complexity of very large alignment problems.

For large inputs, PGGB can use a heuristic based on the Erdős–Rényi model of random graphs to set a sparsification threshold for initial mappings. This model leads us to expect a giant component—or connected subgraph that contains a significant portion of the nodes in the graph—to arise in a random graph of N nodes when the probability of edges between two nodes is $P_{critical} = 1/(N-1)$ [25]. Considering the SEQWISH alignment graph (Figure 1A), where nodes correspond to subsequences and edges to mappings, we seek to ensure that a giant component exists for all homologous collections of nodes in all regions of the pangenome. This will let us reconstruct all transitive relationships in the variation graph without needing to directly compute all pairwise alignments. We thus set a sparsification parameter that ignores mappings with a probability $P_{sparse} \gg P_{critical}$, allowing us to avoid the expected $O(N^2)$ costs implied when $P = 1$. This allows us to dramatically reduce the runtime of alignment and graph induction with negligible effect on accuracy (Table 1), e.g. $10\times$ increase in genomes requires only $20\times$ increase in runtime.

| Pangenome | Size (bp) | Compr. | Time (m) | Mem. (GB) | SNVs | F1-score |
|-----------------|-------------|--------|----------|-----------|---------|----------|
| athaliana7.chr1 | 210174177 | 5.12 | 28.51 | 9.71 | 129374 | 0.877267 |
| ecoli50 | 249520474 | 12.56 | 89.35 | 12.97 | 56915 | 0.947041 |
| ecoli500* | 2572341327 | 23.99 | 1816.66 | 134.59 | 58259 | 0.936551 |
| hsapiens90.chr6 | 15508376475 | 81.17 | 1183.33 | 135.52 | 143972 | 0.971475 |
| mouse17.chr19 | 994731502 | 11.52 | 203.80 | 29.48 | 223951 | 0.907288 |
| primate14.chr6 | 2635610277 | 6.18 | 1742.37 | 61.38 | 2886064 | 0.909077 |
| scerevisiae8 | 96255507 | 6.47 | 8.78 | 3.53 | 53742 | 0.968729 |
| scerevisiae142 | 1702093905 | 55.29 | 1021.81 | 112.98 | 62796 | 0.955988 |
| scerevisiae142* | 1702093905 | 41.41 | 562.89 | 75.91 | 62580 | 0.955650 |
| soy37.chr18 | 2240871558 | 20.50 | 599.66 | 29.17 | 101907 | 0.907878 |
| tomato23.chr2 | 1280460312 | 20.69 | 78.53 | 43.84 | 39243 | 0.948173 |

Table 1 Performance of PGGB with pangenomes of difference species.

For each pangenome, we report its size, the compression ratio (pangenome sequence length divided by graph size), the PGGB runtime, the maximum memory usage of PGGB, the average number of SNVs (across all haplotypes except the one used as reference) identified with MUMMER4 that we used to evaluate SNVs identified with PGGB, and the average F1-score (across all haplotypes except the one used as reference) computed using MUMMER4’s SNVs as ground truth. The name of each pangenome indicates the species and the number of haplotypes. All runs were performed on machines equipped with AMD EPYC 7402P 24-Core, 378 GB of RAM, and a 1 TB Solid-State Drive. All PGGB runs were executed with 48 threads. *Erdős–Rényi random sparsification activated

We finish our graph building with SMOOTHXG, an iterative post-processing step that locally compresses and simplifies the pangenome graph, a new tool specifically designed for PGGB. Although the SEQWISH graph presents a complete, lossless model of the input genomes and their homologies, in our experience it often presents spurious local complexity that can cause problems for diverse types of downstream analysis. A key issue is that pairwise alignments derived across our input are not mutually normalized, leading to different representations of small variants like indels in low-complexity sequences, which in turn generate complex looping motifs that are difficult to process. We partly mitigate this issue by removing short matches from SEQWISH's input alignments. This reduces complexity, but also creates a graph that can be locally "under-aligned" and does not represent all local pairwise relationships. To resolve this, we apply a local realignment kernel, partial order alignment (POA) [26], across all parts of the graph. We do so at a scale of around 1 kbp, which is smaller than the size of most nonlinear patterns of structural variation found in genomes [12, 27]. This allows the PGGB graph to represent complex structurally-variable loci as simple loops through a single copy of duplicated sequences [12]. The kernel is applied to regions that are extracted from a 1-dimensional graph embedding [6]. This embedding orders nodes in the graph so that their distance in the order best-approximates their distance in the genomic paths of the graph. SMOOTHXG first learns this embedding, then obtains partially overlapping segments of the graph (blocks) to which it then applies POA. The realigned blocks are "laced" back into a complete variation graph. We iterate the entire SMOOTHXG step multiple times (3 by default) to limit edge effects that can occur near block boundaries, progressively refining the learned graph embedding. As a final normalization step, we apply GFAFFIX to compress redundant nodes [28], and sort the resulting graph using ODGI.

Downstream applications of pangenome graphs are diverse. PGGB integrates common steps that help to provide immediate feedback on graph build quality. Using ODGI [6], it produces basic graph statistics, such as size, node count, and base content. ODGI creates 1D and 2D visualizations that provide intuition about the structure of the entire graph, with the 1D view showing the relative alignment of paths into the graph structure, and the 2D view showing high-level features of the graph topology. Both can be applied at the scale of multi-gigabasepair graphs. Optionally, the graph statistics and diagnostic plots are summarized in a MultiQC [29] report. We also provide an option to call variants from the graph in the variant call format (VCF), and downstream normalization can be applied to decompose complex nested variation into a minimal reference-relative representation using BiWFA [30]. This allows PGGB to provide input to analyses based on small variants, leading to compatibility with virtually all downstream biological applications. PGGB is thus a multi-sample variant caller for whole genome assemblies.

PGGB has been applied and validated at large scale in projects in the HPRC [12], where it additionally has provided the first sequence-based evidence for systematic recombination between heterologous acrocentric human chromosomes [16]. Here, we present results from its application to a variety of diverse pangenome and comparative genomic contexts (Table 24). We provide information on runtime and resource requirements, showing that even for hundreds of (small) eukaryotic whole genomes, PGGB can provide a variation graph within hours. Due to lack of ground truth, quality evaluation on real data can be difficult. Here, we compare PGGB’s output with SNPs detected by MUMMER4 [20]. These show cross validation F-scores >90% across all tested contexts, indicating that the method performs equivalently to existing standards. However, while MUMMER4 provides only pairwise comparisons with a target reference, PGGB yields a full all-to-all comparison between genomes that leads to completely new bioinformatic analysis modalities.

As a demonstration of the transformative utility of our approach, we note that many downstream applications that are typically based on polarization of variants (e.g. SNPs) relative to a single reference genome may be directly implemented in the space of variation graphs built with PGGB and related methods. This follows from two basic concepts: in the variation graph, nodes are alleles, while genomes can be simply understood as vectors of allele counts. Methods based in this vector space allow us to simultaneously consider *all* classes of variation in downstream analyses, without reference bias, an objective, which to our knowledge has never before been achieved in bioinformatics with the practical generality provided by PGGB. As proof of principle, we put forward a phylogenetic tree constructed directly from distances measured within a pangenome graph of 14 complete assemblies of chromosome six from the great ape family (Figure 1D), which matches established phylogenies of the *Hominoidea* clade based on manually curated sets of SNPs that exclude structurally variable regions [14].

In summary, we present PGGB, a new, modular, and conceptually straightforward approach to understanding sequence relationships between many complete genomes in both pangenomic and comparative genomics settings. Our approach provides a general framework for genome graph building techniques which we expect researchers will upgrade and extend in the future. By making it easy to build variation graphs, PGGB opens the door to diverse downstream population and evolutionary genetic methods that can consider all classes of sequence variation simultaneously. This will allow us to develop conclusive understanding of the links between genome, phenotype, and evolution in an era where the complete assembly of genomes becomes routine.

Online content. PGGB is available at <https://github.com/pangenome/pggb>. Code used for experiments can be accessed at <https://github.com/pangenome/pggb-paper>.

Acknowledgments

The authors thank members of the HPRC Pangenome Working Group for their insightful discussion and feedback, and members of the HPRC production teams for their development of resources used in our exposition.

Funding

The authors gratefully acknowledge support from National Institutes of Health/NIDA U01DA047638 (E.G.), National Institutes of Health/NIGMS R01GM123489 (E.G. and P.P.), National Institutes of Health/NIGMS R35GM142916 (P.H.S), and NSF PPOSS Award #2118709 (E.G. and P.P.), and the Center for Integrative and Translational Genomics (E.G.). S.H. acknowledges funding from the Central Innovation Programme (ZIM) for SMEs of the Federal Ministry for Economic Affairs and Energy of Germany. This work was supported by the BMBF-funded de.NBI Cloud within the German Network for Bioinformatics Infrastructure (de.NBI) (031A532B, 031A533A, 031A533B, 031A534A, 031A535A, 031A537A, 031A537B, 031A537C, 031A537D, 031A538A).

Conflict of interest

Author J.H. is employed by Computomics GmbH.

Author contributions

Project conception: Erik Garrison. *Experimental execution:* Andrea Guarracino. *Project guidance:* Erik Garrison, Sven Nahnsen, Nicole Soranzo, Vincenza Colonna, Robert W. Williams, Pjotr Prins. *Paper writing and editing:* Erik Garrison, Andrea Guarracino, Simon Heumos, Sven Nahnsen, Vincenza Colonna, Robert W. Williams, Pjotr Prins. *Software development:* Erik Garrison, Andrea Guarracino, Simon Heumos, Santiago Marco-Sola, Mwaniki N. Moses. *Documentation:* Andrea Guarracino, Simon Heumos. *Experimental design:* Erik Garrison. *Quality evaluation:* Erik Garrison, Andrea Guarracino, Lorenzo Tattini. *Testing:* Erik Garrison, Andrea Guarracino, Simon Heumos, Flavia Villani, Zhigui Bao, Lorenzo Tattini, Jörg Hagmann, Sebastian Vorbrugg, Christian Kubica, Kaisa Thorell, Rachel L. Rusholme-Pilcher, Sven Nahnsen, Zuyu Yang, Mwaniki N. Moses, Franklin L. Nobrega, Hao Chen, Joep de Ligt, Peter H. Sudmant. *Mus musculus and Rattus Norvegicus:* Flavia Villani, David G. Ashbrook, Hao Chen, Vincenza Colonna. *Tomato pangenome:* Zhigui Bao. *S. cerevisiae and S. paradoxus:* Lorenzo Tattini, Gianni Liti. *Soy G.max (Cannabis C.sativa):* Jörg Hagmann. *A. thaliana:* Sebastian Vorbrugg, Christian Kubica. *Algorithm development:* Santiago Marco-Sola. *Helicobacter pylori:* Kaisa Thorell, Emilio Rudbeck. *Neisseria meningitidis:* Zuyu Yang, Joep de Ligt. *SARS-CoV-2 dataset:* Mwaniki N. Moses. *E. coli and Coliphages:* Franklin L. Nobrega, Yi Wu. *Primate pangenome:* Peter H. Sudmant. *High Performance Computing management:* Pjotr Prins.

Appendix A Methods

Here we provide details about components which are not described in other publications. Our primary focus is on SMOOTHXG, which is the main “glue” that ties together the PGGB pipeline into a coherent system. Through a series of passes over the pangenome, SMOOTHXG reshapes the graph to reduce local complexity and underalignment. This resolves key problems encountered in earlier attempts to implement all-vs-all alignment based graph construction [10, 31], which typically resulted in very complex, looping, graph motifs at small scales, and redundancy caused by match filtering. We additionally provide a description of the evaluation method we use in our cross-validation experiments where PGGB graphs are compared with SNPs determined by MUMMER4.

A.1 SMOOTHXG

SMOOTHXG requires a GFA pangenome graph as input, for example output from SEQWISH. The raw SEQWISH graph is globally unsorted and might be locally unaligned. SMOOTHXG sorts and normalizes the graph preserving nonlinear complex global structural variation. Detailed steps are described subsequently.

Preprocessing. A Path-Guided Stochastic Gradient Descent (PG-SGD) algorithm optimizes the one-dimensional (1D) node order of the graph to best match the nucleotide positions in the embedded paths. A grooming step ensures that for each node, the node orientation follows the consensus node orientation of all path steps visiting the node. A 1D topological sorting of the graph completes the overall sorting steps. Finally, the graph is chopped so that each node does contain a relatively little number of nucleotides (SMOOTHXG default: 100), preserving node topology and order. This prepares the graph for more exact cut points during the block creation process described in the next section.

Create blocks. The smoothable blocks are discovered by iterating over all nodes following the previously calculated order. A node is added to a block if its addition does not exceed the 1. total path length limit of a block, 2. the maximum edge jump limit of a block, or 3. the maximum block length. Blocks are broken at likely Variable Number of Tandem Repeat (VNTR) boundaries and to ensure that the path ranges within each block do not exceed the maximum sequence input size for the POA step described in the next section.

Smooth each block. For each block, padding extends each block to the left and right. This improves the local alignment at the boundaries of each block. The k-mer plus min-hash approach ensures that only unique sequences are passed to the POA step, which can significantly reduce runtime. POA is applied to each block. Optionally, this step generates a multiple sequence alignment in MAF format for each block. The padding is removed, and the block is saved for later integration into a full graph.

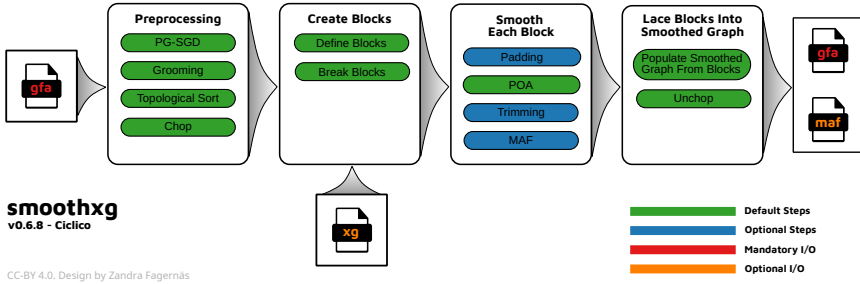


Fig. A1 Overview of the algorithmic steps in SMOOTHXG.

Lace blocks into smoothed graph. The smoothed blocks are laced together to the final pangenome graph following their initial input order. As a final step, the graph is unchopped, preserving the maximum possible node lengths in the graph.

A.2 Validation experiments

To evaluate the accuracy and reliability of our pangenome graph construction and variant calling methods, we designed a cross-validation approach that allowed us to compare the results obtained from the graph-based method (PGGB) against those generated by the widely-used pairwise alignment tool, NUCMER, in MUMMER4 [20].

The cross-validation process begins with the extraction of FASTA sequences from the pangenome graph GFA and preparation of reference sequences. Next, variants are identified using both PGGB (with VG) and nucmer (via a MUMMER4 script), generating a VCF file for each haplotype to ease comparison using the RealTime Genomics toolkit.

These variant files are then compared and evaluated, focusing on regions where both methods are able to call variants—an aspect that we found to be important in the HPRC cross-validation studies, wherein DIPCALL was used to find consistently-alignable regions in which comparisons were conceptually sound [12]. Finally, we collect metrics and statistics for further analysis and visualization. To simplify reproducibility, here we provide a detailed summary of the evaluation process:

1. **Extract FASTA file:** The script starts by extracting the FASTA sequences from the pangenome graph (GFA format) using the `odgi paths` tool.
2. **Take reference sequences:** The script extracts the reference paths in the GFA file and creates a new FASTA file containing these sequences.
3. **Identify variants with PGGB:** The script then identifies variants in the pangenome graph using the `vg deconstruct` tool with appropriate options for haplotype-based variant calling from the graph and complex allele decomposition with BiWFA and VCFLIB. The final variants are saved in a VCF format file.

4. **Pre-process the PGGB-based VCF files:** The script pre-processes the VCF files, including normalizing alleles, removing insertions and deletions larger than 1 base pair, and removing the ALT allele if it is not present in the haplotype.
5. **Identify variants with NUCmer:** The script performs a pairwise sequence alignment between the reference and each contig in the pangenome using NUCmer. The script extracts SNPs from the NUCmer delta file using the `show-snps` command and generates VCF files for each aligned contig.
6. **Merge variants by haplotype:** The script then merges all VCF files for each haplotype generated by NUCmer to create a single VCF file per haplotype.
7. **Variant evaluation:** RTG Tools' `vcfeval` is used to evaluate the performance of PGGB-based variants and NUCmer-based variants by comparing true positives, false positives, and false negatives in shared callable regions. This is done for both non-waved and waved PGGB-based VCF files, allowing for a direct comparison of the performance of these variant calling methods.
8. **Collect statistics:** The script computes summary statistics, such as precision, recall, and F1 scores for each haplotype and writes them to TSV files. It also calculates the total number of variants called and the ratio of evaluated variants for both NUCmer and PGGB-based methods.
9. **Organize output:** Finally, the script organizes the output data, including VCF files, evaluation results, and statistics, into a specified output directory.

Although imperfect due to our lack of ground truth in the context of whole genome alignment, this method provides a way to approximately compare the existing standard for whole genome pairwise alignment, MUMMER4, with PGGB. We focus on SNPs and omit comparison of structural variation for diverse reasons. First, we found extracting SVs from MUMMER4 output to be problematic and poorly-supported. Second this issue remains difficult in genomics due to the multiple near-equivalent representations that a given structural variant allele may have. However, we have addressed these topics in the context of the HPRC paper [12], where significant resources were available to drive an independent evaluation of PGGB and other graph building methods.

References

- [1] Garrison, E., Sirén, J., Novak, A.M., Hickey, G., Eizenga, J.M., Dawson, E.T., Jones, W., Garg, S., Markello, C., Lin, M.F., Paten, B., Durbin, R.: Variation graph toolkit improves read mapping by representing genetic variation in the reference. *Nature Biotechnology* **36**(9), 875–879 (2018). <https://doi.org/10.1038/nbt.4227>
- [2] Paten, B., Novak, A.M., Eizenga, J.M., Garrison, E.: Genome graphs and

- the evolution of genome inference. *Genome Res.* **27**(5), 665–676 (2017). <https://doi.org/10.1101/gr.214155.116>
- [3] Garrison, E.: Graphical pangenomics (2019). <https://doi.org/10.17863/CAM.41621>
- [4] Eizenga, J.M., Novak, A.M., Sibbesen, J.A., Heumos, S., Ghaffari, A., Hickey, G., Chang, X., Seaman, J.D., Rounthwaite, R., Ebler, J., Rautiainen, M., Garg, S., Paten, B., Marschall, T., Sirén, J., Garrison, E.: Pangenome graphs. *Annual Review of Genomics and Human Genetics* **21**(1), 139–162 (2020). <https://doi.org/10.1146/annurev-genom-120219-080406>
- [5] Armstrong, J., Hickey, G., Diekhans, M., Fiddes, I.T., Novak, A.M., Deran, A., Fang, Q., Xie, D., Feng, S., Stiller, J., Genereux, D., Johnson, J., Marinescu, V.D., Alföldi, J., Harris, R.S., Lindblad-Toh, K., Haussler, D., Karlsson, E., Jarvis, E.D., Zhang, G., Paten, B.: Progressive cactus is a multiple-genome aligner for the thousand-genome era. *Nature* **587**(7833), 246–251 (2020). <https://doi.org/10.1038/s41586-020-2871-y>
- [6] Guarracino, A., Heumos, S., Nahnsen, S., Prins, P., Garrison, E.: ODGI: understanding pangenome graphs. *Bioinformatics* (2022). <https://doi.org/10.1093/bioinformatics/btac308>
- [7] Li, H., Feng, X., Chu, C.: The design and construction of reference pangenome graphs with minigraph. *Genome Biology* **21**(1) (2020). <https://doi.org/10.1186/s13059-020-02168-z>
- [8] Hickey, G., Monlong, J., Ebler, J., Novak, A., Eizenga, J.M., Gao, Y., Marschall, T., Li, H., and, B.P.: Pangenome graph construction from genome alignment with minigraph-cactus (2022). <https://doi.org/10.1101/2022.10.06.511217>
- [9] Noll, N., Molari, M., Shaw, L.P., Neher, R.A.: PanGraph: scalable bacterial pan-genome graph construction (2022). <https://doi.org/10.1101/2022.02.24.481757>
- [10] Garrison, E., Guarracino, A.: Unbiased pangenome graphs. *Bioinformatics* **39**(1) (2022). <https://doi.org/10.1093/bioinformatics/btac743>
- [11] Minkin, I., Pham, S., Medvedev, P.: TwoPaCo: an efficient algorithm to build the compacted de bruijn graph from many complete genomes. *Bioinformatics* **33**(24), 4024–4032 (2016). <https://doi.org/10.1093/bioinformatics/btw609>
- [12] Liao, W.-W., Asri, M., Ebler, J., Doerr, D., Haukness, M., Hickey, G., Lu, S., Lucas, J.K., Monlong, J., Abel, H.J., Buonaiuto, S., Chang, X.H.,

- Cheng, H., Chu, J., Colonna, V., Eizenga, J.M., Feng, X., Fischer, C., Fulton, R.S., Garg, S., Groza, C., Guarracino, A., Harvey, W.T., Heumos, S., Howe, K., Jain, M., Lu, T.-Y., Markello, C., Martin, F.J., Mitchell, M.W., Munson, K.M., Mwaniki, M.N., Novak, A.M., Olsen, H.E., Pesout, T., Porubsky, D., Prins, P., Sibbesen, J.A., Tomlinson, C., Villani, F., Vollger, M.R., Bourque, G., Chaisson, M.J., Flicek, P., Phillippy, A.M., Zook, J.M., Eichler, E.E., Haussler, D., Jarvis, E.D., Miga, K.H., Wang, T., Garrison, E., Marschall, T., Hall, I., Li, H., and, B.P.: A draft human pangenome reference (2022). <https://doi.org/10.1101/2022.07.09.499321>
- [13] Fischer, C., Garrison, E.: chfi/gfaestus: a pangenome graph browser. Zenodo (2022). <https://doi.org/10.5281/ZENODO.6954035>. <https://zenodo.org/record/6954035>
- [14] Cagan, A., Theunert, C., Laayouni, H., Santpere, G., Pybus, M., Casals, F., Prüfer, K., Navarro, A., Marques-Bonet, T., Bertranpetit, J., Andrés, A.M.: Natural selection in the great apes. *Mol Biol Evol* **33**(12), 3268–3283 (2016). <https://doi.org/10.1093/molbev/msw215>
- [15] Logsdon, G.A., Vollger, M.R., Hsieh, P., Mao, Y., Liskovych, M.A., Koren, S., Nurk, S., Mercuri, L., Dishuck, P.C., Rhie, A., de Lima, L.G., Dvorkina, T., Porubsky, D., Harvey, W.T., Mikheenko, A., Bzikadze, A.V., Kremitzki, M., Graves-Lindsay, T.A., Jain, C., Hoekzema, K., Murali, S.C., Munson, K.M., Baker, C., Sorensen, M., Lewis, A.M., Surti, U., Gerton, J.L., Larionov, V., Ventura, M., Miga, K.H., Phillippy, A.M., Eichler, E.E.: The structure, function and evolution of a complete human chromosome 8. *Nature* **593**(7857), 101–107 (2021). <https://doi.org/10.1038/s41586-021-03420-7>
- [16] Guarracino, A., Buonaiuto, S., de Lima, L.G., Potapova, T., Rhie, A., Koren, S., Rubinstein, B., Fischer, C., Gerton, J.L., Phillippy, A.M., Colonna, V., and, E.G.: Recombination between heterologous human acrocentric chromosomes (2022). <https://doi.org/10.1101/2022.08.15.504037>
- [17] Crysanto, D., Leonard, A., Pausch, H.: Comparison of methods for building pangenome graphs. In: *Proceeding of 12th World Congress on Genetics Applied to Livestock Production (WCGALP) Technical and Species Orientated Innovations in Animal Breeding, and Contribution of Genetics to Solving Societal Challenges*, pp. 1066–1069 (2022). Wageningen Academic Publishers
- [18] Zhang, H., Wu, S., Aluru, S., Li, H.: Fast sequence to graph alignment using the graph wavefront algorithm. arXiv preprint arXiv:2206.13574 (2022)
- [19] Leonard, A.S., Crysanto, D., Mapel, X.M., Bhati, M., Pausch, H.: Graph

- construction method impacts variation representation and analyses in a bovine super-pangenome (2022). <https://doi.org/10.1101/2022.09.17.508368>
- [20] Marçais, G., Delcher, A.L., Phillippy, A.M., Coston, R., Salzberg, S.L., Zimin, A.: MUMmer4: A fast and versatile genome alignment system. *PLoS Computational Biology* **14**(1), 1005944 (2018). <https://doi.org/10.1371/journal.pcbi.1005944>
- [21] Guarracino, A., Mwaniki, N., Marco-Sola, S., Garrison, E.: wfmash: whole-chromosome pairwise alignment using the hierarchical wavefront algorithm (2021). <https://github.com/ekg/wfmash>
- [22] Marco-Sola, S., Moure, J.C., Moreto, M., Espinosa, A.: Fast gap-affine pairwise alignment using the wavefront algorithm. *Bioinformatics* (2020). <https://doi.org/10.1093/bioinformatics/btaa777>
- [23] Marco-Sola, S., Eizenga, J.M., Guarracino, A., Paten, B., Garrison, E., Moreto, M.: Optimal gap-affine alignment in $o(s)$ space. *Bioinformatics* **39**(2) (2023). <https://doi.org/10.1093/bioinformatics/btad074>
- [24] Jain, C., Koren, S., Dilthey, A., Phillippy, A.M., Aluru, S.: A fast adaptive algorithm for computing whole-genome homology maps. *Bioinformatics* **34**(17), 748–756 (2018). <https://doi.org/10.1093/bioinformatics/bty597>
- [25] Bollobás, B.: The evolution of random graphs—the giant component. In: *Random Graphs*, vol. 184, pp. 130–59 (2001). Cambridge university press Cambridge
- [26] Lee, C., Grasso, C., Sharlow, M.F.: Multiple sequence alignment using partial order graphs. *Bioinformatics* **18**(3), 452–464 (2002). <https://doi.org/10.1093/bioinformatics/18.3.452>
- [27] Vollger, M.R., DeWitt, W.S., Dishuck, P.C., Harvey, W.T., Guitart, X., Goldberg, M.E., Rozanski, A.N., Lucas, J., Asri, M., Munson, K.M., Lewis, A.P., Hoekzema, K., Logsdon, G.A., Porubsky, D., Paten, B., Harris, K., Hsieh, P., and, E.E.E.: Increased mutation rate and inter-locus gene conversion within human segmental duplications (2022). <https://doi.org/10.1101/2022.07.06.498021>
- [28] Doerr, D., Marijon, P., Marschall, T.: GFAffix identifies walk-preserving shared affixes in variation graphs and collapses them into a non-redundant graph structure (2023). <https://github.com/marschall-lab/GFAffix>
- [29] Ewels, P., Magnusson, M., Lundin, S., Käller, M.: Multiqc: summarize analysis results for multiple tools and samples in a single report. *Bioinformatics* **32**(19), 3047 (2016). <https://doi.org/10.1093/bioinformatics/btw354>

[btw354](#)

- [30] Garrison, E., Kronenberg, Z.N., Dawson, E.T., Pedersen, B.S., Prins, P.: A spectrum of free software tools for processing the VCF variant call format: vcfliib, bio-vcf, cyvcf2, hts-nim and slivar. *PLoS Computational Biology* **18**(5), 1009123 (2022). <https://doi.org/10.1371/journal.pcbi.1009123>

- [31] Llamas, B., Narzisi, G., Schneider, V., Audano, P.A., Biederstedt, E., Blauvelt, L., Bradbury, P., Chang, X., Chin, C.-S., Fungtammasan, A., Clarke, W.E., Cleary, A., Ebler, J., Eizenga, J., Sibbesen, J.A., Markello, C.J., Garrison, E., Garg, S., Hickey, G., Lazo, G.R., Lin, M.F., Mahmoud, M., Marschall, T., Minkin, I., Monlong, J., Musunuri, R.L., Sagayaradj, S., Novak, A.M., Rautiainen, M., Regier, A., Sedlazeck, F.J., Siren, J., Souilmi, Y., Wagner, J., Wrightsman, T., Yokoyama, T.T., Zeng, Q., Zook, J.M., Paten, B., Busby, B.: A strategy for building and using a human reference pangenome. *F1000Res* **8**, 1751 (2021). <https://doi.org/10.12688/f1000research.19630.2>