



Fast and accurate mapping of long reads to complete genome assemblies with VerityMap

Andrey V. Bzikadze, Alla Mikheenko and Pavel A. Pevzner

Genome Res. published online November 15, 2022

Access the most recent version at doi:[10.1101/gr.276871.122](https://doi.org/10.1101/gr.276871.122)

P<P	Published online November 15, 2022 in advance of the print journal.
Accepted Manuscript	Peer-reviewed and accepted for publication but not copyedited or typeset; accepted manuscript is likely to differ from the final, published version.
Creative Commons License	This article is distributed exclusively by Cold Spring Harbor Laboratory Press for the first six months after the full-issue publication date (see https://genome.cshlp.org/site/misc/terms.xhtml). After six months, it is available under a Creative Commons License (Attribution-NonCommercial 4.0 International), as described at http://creativecommons.org/licenses/by-nc/4.0/ .
Email Alerting Service	Receive free email alerts when new articles cite this article - sign up in the box at the top right corner of the article or click here .

An advertisement banner with a teal background. On the left, the text reads 'CRISPR and RNAi Genetic Screening. Your new superpower.' in white. In the center, there is a white-bordered box containing the words 'LEARN MORE' in teal. On the right, there is a photograph of a woman wearing a red superhero mask and a red cape over a grey top. To the right of the photo is the Cellecta logo, which consists of a cluster of green dots of varying sizes, with the word 'CELLECTA' in white capital letters below it.

To subscribe to *Genome Research* go to:
<https://genome.cshlp.org/subscriptions>

Published by Cold Spring Harbor Laboratory Press

Fast and accurate mapping of long reads to complete genome assemblies with VerityMap

Andrey V. Bzikadze^{1*†}, Alla Mikheenko^{2*}, Pavel A. Pevzner³

¹ Graduate Program in Bioinformatics and Systems Biology, University of California, San Diego, CA, USA

² Center for Algorithmic Biotechnology, Saint Petersburg State University, Saint Petersburg, Russia

³ Department of Computer Science and Engineering, University of California, San Diego, CA, USA

† corresponding author: abzikadze@ucsd.edu, * these authors contributed equally

Running title: Fast and accurate mapping with VerityMap

Abstract

Recent advancements in long-read sequencing have enabled the telomere-to-telomere (complete) assembly of a human genome and are now contributing to the haplotype-resolved complete assemblies of multiple human genomes. Since the accuracy of read mapping tools deteriorates in highly-repetitive regions, there is a need to develop accurate, *error-exposing* (detecting potential assembly errors), and *diploid-aware* (distinguishing different haplotypes) tools for read mapping in complete assemblies. We describe the first accurate, error-exposing, and partially diploid-aware VerityMap tool for long-read mapping to complete assemblies.

Introduction

The initial draft sequence of the human genome contained thousands of gaps and scaffolds (International Human Genome Sequencing Consortium 2001; Venter et al. 2001). Even nowadays, the up-to-date human reference genome (GRCh38.p13) still contains 349 gaps and 472 scaffolds (Miga et al. 2014). These gaps contain regions that include large and biomedically important multi-megabase arrays of tandem satellite repeats that orchestrate chromosome segregation during cell division (Miga 2019, 2020). For two decades, centromeric and pericentromeric satellite arrays evaded all assembly efforts resulting in a limited understanding of their sequence organization. Recently, the Telomere-to-Telomere (T2T) Consortium generated the first nearly complete assembly of an effectively haploid human cell line CHM13 that, in particular, includes assemblies of all satellite repeats (M. Jain et al. 2018; Miga et al. 2020; Logsdon et al. 2021; Nurk et al. 2022).

Long-read technologies, such as the ones developed by Pacific Biosciences (PacBio) and Oxford Nanopore Technologies (ONT), have changed the landscape of computational methods for genome assembly and opened a possibility to generate haplotype-resolved complete assemblies (Chin et al. 2016; Li 2016; Lin et al. 2016; Koren et al. 2017; Wick et al. 2017; Kolmogorov et al. 2019; Ruan and Li 2020; Nurk et al. 2020; Kolmogorov et al. 2020; Cheng et al. 2021). Some of these assemblers utilize ultra-long error-prone reads in order to assemble the most repetitive genomic regions, such as satellite arrays (Bzikadze and Pevzner 2020; Miga et al. 2020; Shafin et al. 2020). The emergence of long high-fidelity (HiFi) PacBio reads has, once again, revolutionized the field of genome assembly (Wenger et al. 2019; Logsdon et al. 2021) and resulted in some of the most contiguous assemblies to-date (Nurk et al. 2020; Cheng et al. 2021; Nurk et al. 2022).

The emergence of “complete genomics” (Nurk et al. 2020; Cheng et al. 2021; Nurk et al. 2022) and “complete metagenomics” (Bickhart et al. 2022) is shifting the focus of read mapping from detecting mutations to detecting errors in the newly generated assemblies (see Supplementary Note “Similarities and differences between detecting mutations and detecting misassemblies”, Supplementary Figure 1). Although most previous read mapping efforts were aimed at detecting mutations by aligning reads against the reference genome instead of assembling them (since genome assemblers typically fail to assemble highly-repetitive regions), HiFi assemblers accurately assemble even the most complex genomic regions, and while a limited number of misassemblies persist, these assemblies provide a path to potentially eliminate the need for mapping reads against a reference. They also result in extremely low rates of single-nucleotide errors in non-repetitive regions (Nurk et al. 2020; Cheng et al. 2021; Nurk et al. 2022), largely eliminating the need for the follow-up read mapping to these regions: instead of mapping reads using external tools, existing HiFi assemblers automatically provide highly accurate alignments of reads to the assemblies they generate. However, HiFi assemblers still make assembly errors in highly-repetitive regions and thus necessitate development of new mapping tools for detecting these errors by identifying discordant reads.

The need for evaluation of the complete genome assemblies and the absence of ground truth datasets has resulted in the development of specialized pipelines for assembly validation (Mc Cartney et al. 2022). Even though the QUAST tools represent the state-of-the-art framework for quality assessment of genome assemblies (Gurevich et al. 2013; Mikheenko, Saveliev, and Gurevich 2016; Bushmanova et al. 2016; Mikheenko et al. 2018), these tools mostly rely on the availability of a *reference* genome and are thus inapplicable for evaluating complete assemblies that include megabases of novel sequence. In particular, QUAST uses BUSCO (Simão et al. 2015) to evaluate the gene content of assemblies and assess completeness of an assembly without a reference genome. However, BUSCO only analyzes conserved single-copy genes and is inapplicable for gene-poor regions such as centromeres. Moreover, QUAST identifies assembly errors only in comparison with a reference genome. Alternative *reference-free* methods can be divided into *alignment-free* methods that compare *k*-mer spectra between reads and an assembly (Mapleson et al. 2017; Rhie et al. 2020) and *alignment-based* methods that use the read mapping tools (Li and Durbin 2009; Langmead and Salzberg 2012; Li 2016, 2018; C. Jain et al. 2020) to generate read alignments, use them for assessing the quality of an assembly, and correct the assembly based on the identified discordant reads. However, although alignment-free methods can estimate assembly quality and completeness, they do not output information about the positions of assembly errors.

Mikheenko et al. (Mikheenko et al. 2020) demonstrated that performance of existing read-mapping tools deteriorates in complex regions, especially in the presence of assembly errors, and developed the TandemMapper tool for mapping long reads to extra-long tandem repeat (ETRs) in the constructed assemblies. Even though TandemMapper accurately maps reads to centromeric regions and enables evaluations of their assemblies (Bzikadze and Pevzner 2020; Miga et al. 2020; Logsdon et al. 2021), it falls short of mapping reads in even more challenging megabase-long genomic regions like HSAT2,3. Moreover, TandemMapper, designed specifically for ETRs, becomes prohibitively slow for evaluating complete genome assemblies.

Generating a complete assembly for a single *haploid* human genome represents a landmark achievement (Nurk et al. 2022). However, since a single genome does not represent the genomic diversity of the human population, the Human Pangenome Reference (HPR) Consortium is now generating high-quality *diploid* assemblies for hundreds of humans. This ambitious goal raises the bar for scalable *diploid-aware* quality assessment tools that can evaluate *entire* diploid assemblies rather than certain selected (albeit complex) loci and identify errors in these assemblies. To the best of our knowledge, no error-exposing tool for reference-free benchmarking of complete haplotype-resolved assemblies or even individual loci is currently available. Although Winnowmap2 (C. Jain et al. 2020, 2022) maps reads to complex repetitive regions of the *error-free* assemblies, it is neither error-exposing nor diploid-aware since its performance deteriorates in the case of *error-prone* assemblies and heterozygous sites (Table 2) (Mikheenko et al. 2020).

We present VerityMap, a fast, accurate, and error-exposing aligner for mapping long reads to the complete assemblies (see Table 1 for an informal comparison of various long-read aligners). VerityMap also represents the first step towards diploid-aware alignment by mapping reads to haplotype-resolved assemblies of individual loci (see Supplementary Note “The challenge of diploid-aware read mapping”). It was used for detecting and correcting misassemblies in the intermediate assemblies of the first complete human genome assembled by the T2T Consortium (Mc Cartney et al. 2022; Nurk et al. 2022; Miga et al. 2020). Section “Alpha-satellite and human satellite 1,2,3 validation” in (Nurk et al. 2022) illustrates how VerityMap contributed to verifying the complete human genome assembly generated by the T2T Consortium..

All state-of-the-art long-read aligners are based on finding some shared k -mers between a read and an assembly followed by sparse dynamic programming in a graph where each vertex represents a pair of the starting positions of a shared k -mer. Ideally, the parameter k (k -mer size) should be selected to be as large as possible (e.g., slightly below the read length) for accurate mapping of *error-free* reads. In practice, it is selected to be rather small to ensure that an *error-prone* read has some shared k -mer with an

assembly. This condition dictates selecting a rather small k in the case of long error-prone reads (minimap2 v2.24 uses default $k = 15$ for mapping ONT reads, and $k = 19$ for mapping PacBio CLR reads <https://lh3.github.io/minimap2/minimap2.html>) and makes it difficult to map reads in highly-repetitive regions. Accurate HiFi reads allow one to increase k by an order of magnitude thus facilitating read mapping in complex regions.

Although an increased k -mer size looks like a rather minor parameter change in read mapping tools, it actually represents a new challenge. We show that developing a fast, accurate, error-exposing, and diploid-aware long-read mapping tool requires new algorithmic ideas, not unlike constructing the de Bruijn graph of large genomes for large k -mer sizes (Bankevich et al. 2022). To achieve this goal, VerityMap first identifies all rare k -mers in the assembly, carefully selects a small subset of rare k -mers (*solid* k -mers), finds locations of solid k -mers in each read, constructs a *compatibility graph* with the vertex-set formed by all matches between the selected solid k -mers shared by a read and the assembly, finds an optimal path in this graph using sparse dynamic programming, and uses this path for read mapping. Since HiFi reads are accurate, VerityMap utilizes large k -mer sizes (e.g., $k=300$) for constructing the compatibility graph to achieve accurate read mapping. It thus faces the algorithmic challenge of identifying rare k -mers in a large genome for a large k .

Although efficient indexing (finding locations) of all k -mers in the genome (and thus identifying rare k -mers) represents a backbone of many bioinformatics algorithms (Marçais and Kingsford 2011), it remains an open problem in the case of large genomes and large k -mer sizes. Indeed, a naive indexing algorithm with running time $O(|Genome|^k)$ becomes prohibitively slow in the case of accurate HiFi reads since mapping these reads is based on large k -mer sizes (e.g., $k=300$). Jellyfish (Marçais and Kingsford 2011), KMC3 (Kokot, Długosz, and Deorowicz 2017), and more scalable GPU-based k -mer counting approaches (Nisa et al. 2021) generate a database of counts that allows a constant-time count query for any k -mer. However, even though one can rapidly generate a *counting* database, existing implementations for *indexing* all rare k -mers still require $O(|Genome|^k)$ time. On the other hand, the Meryl k -mer counting algorithm (Rhie et al. 2020) only works with $k \leq 64$, which is substantially lower than what is needed for

accurate mapping of HiFi reads (e.g., $k = 300$). Inspired by Jellyfish, VerityMap overcomes limitations of existing approaches for rapid indexing of rare k -mers in the case of large genomes and large k -mer sizes by using multiple *Bloom filters* (Bloom 1970) and the *count-min sketches* (Cormode and Muthukrishnan 2005).

VerityMap also addresses the challenge of identifying errors in genome assemblies in a reference-free mode. Even though the CHM13 cell line is effectively haploid, it features genomic instabilities between two haplotypes typically represented by insertions in one of the homologous chromosomes that are referred to as *Het sites* (Nurk et al. 2022). The ratio of two haplotypes in a Het site is not necessarily 1:1 suggesting that one of the haplotypes might be more prevalent. Automatic detection of Het sites in a haploid assembly and distinguishing them from misassemblies is an important prerequisite for validating diploid assemblies. In addition to its read-mapping module, VerityMap includes a misassembly detection module for identifying misassemblies, Het sites, collapsed haplotypes, and haplotype-switch errors. The T2T Consortium applied this module to verify and correct intermediate assemblies of the CHM13 cell line (Nurk et al. 2022).

Results

Limitations of existing read-mapping approaches. Since all existing assemblers generate some misassembled contigs, accurate mapping of reads that span the misassembly breakpoints is a critically important assembly validation step. Although minimap2 (Li 2018, 2021) and Winnowmap2 (C. Jain et al. 2020, 2022) accurately map reads to an error-free assembly (even in repetitive regions), their accuracy deteriorates in the case of error-prone assemblies or haplotype-resolved assemblies (since reads from highly similar regions often map to incorrect instances of these regions, albeit with several mismatches). Moreover, the dynamic programming algorithm for sequence alignment (with standard scoring) often fails to correctly align reads in highly-repetitive regions, motivating the need for a new scoring (Supplementary Note “Frequency-aware sequence alignment scoring”, Supplementary Figure 2 and 3).

Currently, TandemMapper is the only error-exposing aligner that accurately maps long reads to ETR assemblies that potentially contain misassemblies (Mikheenko et al. 2020). However, even though TandemMapper was used to validate the first centromere assemblies (Bzikadze and Pevzner 2020; Miga et al. 2020; Logsdon et al. 2021), it has limitations that prevent its applications to more complex regions of the genome, e.g., it is designed for analyzing various *higher-order repeats (HORs)* such as human centromeres and is not applicable to repeats without HORs or for analyzing non-repetitive parts of the genome.

Rare and solid k -mers. To speed-up the standard (time-consuming) dynamic programming algorithm for aligning long reads, many mappers (Li 2018; C. Jain et al. 2020; Mikheenko et al. 2020), as well as fast algorithms for detecting overlapping reads in genome assemblers (Chin et al. 2016; Li 2016; Lin et al. 2016; Koren et al. 2017; Wick et al. 2017; Kolmogorov et al. 2019; Ruan and Li 2020; Nurk et al. 2020; Kolmogorov et al. 2020; Cheng et al. 2021) construct a compatibility graph on a carefully selected small subset of k -mers in each read. The key differences between these tools are reflected in algorithms for selecting these k -mers and weighting the edges of the compatibility graph (rather than constructing the compatibility graph). VerityMap uses a new approach for selecting k -mers and weighting edges in the compatibility graph that addresses limitations of the previously developed long-read mappers and enables error-exposing and partially diploid-aware read mapping. VerityMap further finds a longest path in this graph using sparse dynamic programming, and uses this path for read mapping.

A k -mer from an assembly is *rare* if it appears at most *MaxRareOccurrences* times in this assembly (otherwise a k -mer is called *frequent*). A rare k -mer is *solid* if it appears in a single contig, and its reverse-complementary k -mer does not appear in any contig. Note that this definition of a solid k -mer differs from the one used by TandemMapper (Mikheenko et al. 2020). Although the highly-repetitive regions in the newly assembled complete genome constitute the biggest challenge for read mapping, they form less than 8% of the human genome (Nurk et al. 2022). Since nearly all k -mers in other regions are

solid, VerityMap downsamples them to reduce the running time and memory footprint (Supplementary Note “Downsampling solid k -mers in non-repetitive regions”).

VerityMap pipeline. Figure 1 illustrates the VerityMap pipeline. First, VerityMap stores all k -mers from the assembly in two Bloom filters (Bloom 1970) and count–min sketch (CMS) data structures (Cormode and Muthukrishnan 2005) for estimating the count of each k -mer in the assembly and generating the set of solid k -mers that are later used for read mapping (see subsection “Bloom filter and count-min sketch”).

For each read, VerityMap finds all solid k -mers in this read and builds a compatibility graph with the vertex-set formed by pairs of all solid k -mers shared between a read and the assembly. Since the downsampled set of solid k -mers is small, the compatibility graph is typically small; the longest path in this graph can be rapidly found using sparse dynamic programming. The careful definition of edge-weights in the compatibility graph helps to select the correct path (that represents the primary read alignment) even in the presence of assembly errors. Moreover, since the graph stores information about discrepancies in distances between solid k -mers occurring in a read and solid k -mers occurring in the assembly, it can be used to detect approximate locations of misassembly breakpoints without a base-level alignment that might be quite unreliable in highly-repetitive regions. VerityMap first attempts to align the read to the forward strand, and then — to the reverse strand. Defaults for all parameters of VerityMap were selected based on performance analysis (Supplementary Figure 5 and 6, Supplementary Note “VerityMap parameters”). The output of VerityMap is in the standard SAM format.




Figure 1. VerityMap pipeline. (Left) The input of VerityMap is an assembly (a set of contigs) and a set of reads that contributed to this assembly. VerityMap iterates through each contig twice in order to identify solid k -mers. At the first iteration, VerityMap stores k -mers that appear in multiple contigs and all reverse-complementary k -mers within the BanBloomFilter. For each contig, VerityMap constructs a CMS that counts occurrences of k -mers within this contig. Finally, VerityMap uses OnceBloomFilter (and BanBloomFilter) to distinguish between rare k -mers that appear within a single and multiple contigs. Both Bloom filters and the CMS corresponding to the current contig are being modified simultaneously during the first iteration through the assembly. At the second iteration, VerityMap queries the constructed data structures to identify the set of solid k -mers. (Right) Aligning a read GTTAGATAGATGGATT against a misassembled contig GTTGGATTGATAGATAGATG with an 8-nucleotide long deletion TAGATAGA (solid k -mers are shown in blue). The solid k -mer GT (TG) precedes (follows) the deletion breakpoint. The nucleotide-based fitting alignment fails to identify this deletion due to limitations of the standard scoring approaches in highly-repetitive regions. In contrast, VerityMap identifies this deletion using the k -mer-based sparse fitting alignment and a new scoring approach. To achieve this goal, it constructs a compatibility graph on all pairs of solid k -mers shared between a read and the assembly and finds a longest path in this graph. The new scoring reflects the discrepancies in distances between solid k -mers in the assembly (distance 2 between GT and TG in the assembly) and solid k -mers in the read (distance 10 between GT and TG in the read), resulting in $diff(GT, TG)=8$. VerityMap incorporates these discrepancies into the edge-weights of the compatibility graph and outputs a longest path in this graph as the primary read alignment.

Datasets. Recently, the T2T Consortium generated the first complete assembly of a human genome represented by the CHM13 cell line (referred to as **CHM13 assembly**; (Nurk et al. 2022), ChrY in this assembly is from a male sample HG002). In addition to analyzing the entire assembly using VerityMap,

we focused on a particularly complex centromeric satellite region in Chromosome 9. Supplementary Table 2 in Supplementary Note “Extended benchmarking of long-read mapping tools” describes benchmarking results on a multitude of even more challenging datasets. Below we focus on one of these datasets that was generated by extracting the centromeric region from Chromosome 9 of the CHM13 assembly (referred to as **Cen9**). This region contains a human satellite 3 (HSat3) array that represented some of the most difficult challenges for the assembly effort of the T2T Consortium (Supplementary Note “Resolution of Chr6 and Chr9” of Nurk et al. 2022). To analyze how VerityMap handles assembly errors we introduced an artificial deletion of length 10 kbp into the Cen9 region at position 20 Mbp (referred to as **Cen9_{del10}**). We further introduce a series of assemblies with various artificial misassemblies including a deletion, insertions, duplications, and a relocation in the Cen9 region: **Cen9_{del5}**, **Cen9_{ins5}**, **Cen9_{ins10}**, **Cen9_{tandem}**, **Cen9_{dup}**, **Cen9_{reloc}** (detailed description of all datasets is presented in Supplementary Table 1, Supplementary Note “Extended benchmarking of long-read mapping tools”).

It remains unclear how to benchmark VerityMap on haplotype-resolved diploid assemblies since the HPR Consortium is still validating these assemblies. We thus used the assembled centromere of Chromosome X (CenX) from the HG002 HiFi read-set (from a male) generated by the HPR Consortium and the assembled CenX from the haploid CHM13 HiFi read-set generated by the T2T Consortium to mimic a diploid assembly of CenX. To analyze how VerityMap handles diploid assemblies we took alpha-satellite regions from Chromosome X of CHM13 and HG002 assemblies (combined, we refer to these chromosomes as **ChrXDiploid** — a synthetic diploid chromosome) and merged them into one file (referred to as **CenX-Diploid**). Then we cut each assembly at one of the canonical HOR units and concatenated them to mimic a *haplotype-switch* error (referred to as **CenX-Diploid-Switch**). Afterward, we simulated reads from both HG002 and CHM13 assemblies of ChrX to check whether VerityMap can detect the haplotype-switch error. We refer to alpha satellite arrays in Chromosome X in CHM13 (HG002) genome as **ASat-X (ASat-X-HG002)**.

To illustrate that VerityMap identifies errors in real assemblies, we consider Cen10 in the interim version of **CHM13** that predates the earliest publicly released version v0.9 (referred to as

CHM13-Cen10-interim). **CHM13-Cen10-interim** contains a structural error that was detected by VerityMap and was later corrected by the T2T Consortium (**Cen10** in current **CHM13 assembly** does not contain this error).

Below we list some benchmarking datasets (see Supplementary Table 1 in Supplementary Note “Extended benchmarking of long-read mapping tools” for the list of all benchmarking datasets). All simulated datasets are generated with pbsim2 (Ono, Asai, and Hamada 2021).

- **CHM13-SimHiFi** dataset contains HiFi reads simulated from the CHM13 assembly.
- **Cen9Sim** dataset contains HiFi reads simulated from **Cen9**, while **Cen9Sim-Het** dataset contains HiFi reads simulated from both **Cen9** and **Cen9_{del10}** (1:1 ratio).
- **ChrxDiploid-Sim** dataset contains HiFi reads simulated from **ASat-X** and **ASat-X-HG002**.
- **CHM13-RealHiFi** dataset contains real HiFi reads for the CHM13 sample generated by the T2T Consortium (20 kbp library; accession numbers: [SRX7897685-8](https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=SRX7897685-8)). This dataset is extended by reads originating from the ChrY of HG002 sample.
- **ChrxDiploid-RealHiFi** dataset contains HiFi reads from **CHM13-RealHiFi** recruited to **ChrxDiploid** using Winnowmap2.

VerityMap maps nearly all reads with an extremely low number of incorrectly mapped reads. Analysis of the CHM13-SimHiFi dataset (Supplementary Table 2) illustrates that VerityMap mapped more than 97% of reads in all chromosomes (except for Chr21 where it mapped 94.25% of reads) with a low number of incorrectly mapped reads (76 out of 1.7M reads for the entire genome). Mapped reads cover 99.5% of the entire genome (only 14M uncovered bases).

VerityMap identifies assembly errors and heterozygous sites even in highly-repetitive genomic regions. Table 1 illustrates that VerityMap, Winnowmap2, and minimap2 accurately map long and accurate reads to error-free assemblies (see subsection “Datasets” and Supplementary Table 1 in Supplementary Note “Extended benchmarking of long-read mapping tools” for description of all datasets

and detailed benchmarking). However, assembly errors and Het variants often trigger incorrect alignments, especially in highly-repetitive regions. To benchmark the ability to map reads in a vicinity of the misassembly breakpoint, we aligned reads from **Cen9Sim** dataset to **Cen9_{del10}** assembly containing an artificial deletion of length 10 kbp (Table 2, top left). Although minimap2 and Winnommap2 map more reads than VerityMap, they have a high rate of incorrectly mapped reads while VerityMap yields few incorrectly mapped reads. Since incorrect read alignments may prevent the detection of errors and heterozygous sites in downstream analysis, it is preferable to classify some reads as “unalignable” instead of generating erroneous alignments. Both minimap2 and Winnommap2 failed to detect this deletion: minimap2 did not report any primary alignment in a 1 kb region before the breakpoint (all reads were incorrectly mapped somewhere else) while Winnommap2 incorrectly extended alignments through the deletion breakpoint in the **Cen9_{del10}** region (Figure 2). In addition, to reproduce a scenario of a heterozygous deletion (Table 2, top right), we aligned reads from the **Cen9Sim-Het** dataset containing HiFi reads simulated from both **Cen9** and **Cen9_{del10}** to the **Cen9_{del10}** assembly. Table 2 (top right) illustrates that both minimap2 and Winnommap2 incorrectly mapped reads simulated from **Cen9** near the deletion breakpoint. VerityMap either reports reads that are correctly spanning the deletion breakpoint or clips read alignments from **Cen9Sim** dataset at the breakpoint and thus indicates a putative heterozygous deletion. Table 2, bottom provides further benchmark of VerityMap, Winnommap2, and minimap2 of aligning simulated reads from the Cen9Sim read-set to the **Cen9_{del5}**, **Cen9_{ins5}**, **Cen9_{ins10}**, **Cen9_{tandem}**, **Cen9_{dup}**, **Cen9_{reloc}** assemblies (detailed description of these datasets, and benchmarking on more datasets is presented in Supplementary Tables 1–4, Supplementary Note “Extended benchmarking of long-read mapping tools”). In all cases, VerityMap has the minimal number of incorrectly mapped reads and either reports reads that are correctly spanning the breakpoint or clips read alignments at the breakpoint.

Supplementary Note “Validation of fragmented haploid and diploid assemblies” shows how VerityMap can be applied to fragmented rather than complete assemblies. Supplementary Note “Extended benchmarking of long-read mapping tools” (Supplementary Figure 4) applies VerityMap to the assembly

of the CHM13 cell line generated by LJA (Bankevich et al., 2022) and identifies putative misassemblies and heterozygous sites.

Table 1. Benchmarking VerityMap, TandemMapper, Winnowmap2, and minimap2 with reads simulated from Chromosome 1 and 9.

read mapping tool	accurate in error-free assemblies	error-exposing (accurate in error-prone assemblies)	diploid-aware	CPU time (minutes)		memory footprint (GB)	
				Chr1	cen9	Chr1	cen9
VerityMap	+	+	+/-	275	500	6	4
TandemMapper	+/-	+/-	-	7012	—	212	—
Winnowmap2	+	-	-	257	1720	8	32
minimap2	+	-	-	28	33	7	8.5

Aligning 137,507 simulated reads to Chromosome 1 (248 Mb length) and 40,272 simulated reads to the centromere of Chromosome 9 (42 Mb length) of the CHM13 assembly. Simulated reads are generated as described in Supplementary Note “Extended benchmarking of long-read mapping tools”. VerityMap, Winnowmap2, and minimap2 incorrectly mapped only 7, 20, and 22 reads to Chromosome 1, respectively. Since TandemMapper was designed for accurate read mapping to HOR arrays (rather than large regions without an HOR-like structure), it incorrectly mapped many reads in other regions, so we rate it “+/-” in terms of accuracy for error-free assemblies. VerityMap allows accurate mapping to haplotype-resolved assemblies of individual loci (Figure 3) and thus provides the initial step towards diploid-aware mappers for complete haplotype-resolved assemblies (we rate it “+/-”). The best value for each column is indicated in bold. Precise definitions of terms “accurate”, “error-exposing”, and “diploid-aware” are given in Supplementary Note “Summary of benchmarking results”.

Table 2. Benchmarking VerityMap, Winnowmap2, and minimap2 on data containing artificial misassembly breakpoints.

	Cen9Sim. Total reads: 40263			Cen9Sim-Het. Total reads: 80559		
	#correctly / incorrectly mapped reads	#correctly / incorrectly extended through misassembly breakpoint	# alignments from haplotype w/o deletion clipped at breakpoint	#correctly/ incorrectly mapped reads	#correctly / incorrectly extended through misassembly breakpoint	# alignments from haplotype w/o deletion clipped at breakpoint
VerityMap	37498/33	6/0	17	747945/73	2/0	14
Winnowmap2	39180/1083	0/9	0	78368/2176	0/9	0

minimap2	39140/1123	0/0	0	78286/2230	0/0	0
-----------------	------------	-----	---	------------	-----	---

Dataset	# alignments correctly / incorrectly extended through breakpoint			# alignments clipped at breakpoint		
	VerityMap	Winnomap2	minimap2	VerityMap	Winnomap2	minimap2
Total reads: 40263						
Cen9_{del5}	8 / 0	0 / 0	0 / 0	6	9	6
Cen9_{tandem}	6 / 0	0 / 9	0 / 0	5	0	0
Cen9_{dup}	19 / 0	4 / 11	15 / 4	2	0	0
Cen9_{ins5}	14 / 0	9 / 0	14 / 0	1	1	1
Cen9_{ins10}	12 / 0	0 / 0	2 / 0	7	0	0
Cen9_{reloc}	5+0 / 0+0	0+0 / 8+0	0+0 / 0+0	10+8	0+2	0+0

The concept of a correctly/incorrectly mapped read is defined in Supplementary Note “Measuring performance of mapping software”. Only primary alignments were taken into account. The majority of reads incorrectly mapped by minimap2 and Winnomap2 have secondary alignments to the correct positions: in cases of nearly perfect duplications, minimap2 and Winnomap2 might incorrectly choose primary alignments, while VerityMap classifies a read as “unalignable”. Aligning simulated reads (top left) from the Cen9Sim read-set to the Cen9_{del10} region with an artificial deletion, (top right) from the Cen9Sim-Het dataset to the Cen9_{del10} sequence. VerityMap reports 2741 (5692) “unalignable” reads in Cen9Sim (Cen9Sim-Het) dataset. (Top left) The total number of reads spanning the deletion breakpoint is 9. VerityMap correctly identifies 6 reads that span the breakpoint and reports 17 read alignments that are clipped at the breakpoint site (note that some of these reads did not span the breakpoint during simulation). (Top right) The deletion breakpoint in this assembly is spanned by 24 simulated reads (15 reads simulated from a haplotype with deletion and 9 reads from a haplotype without deletion). Even though all mappers correctly map all 15 reads from the haplotype with the deletion, only VerityMap correctly identifies two reads from a haplotype without the deletion spanning the breakpoint. Additionally, it reports 14 alignments for reads from haplotype without the deletion clipped at the breakpoint site. The best value for each column is indicated in bold. (Bottom) Benchmarking VerityMap, Winnomap2, and minimap2 when aligning simulated reads from the Cen9Sim read-set to the Cen9_{del5}, Cen9_{ins5}, Cen9_{ins10}, Cen9_{tandem}, Cen9_{dup}, Cen9_{reloc} assemblies. The number of correctly (incorrectly) mapped reads ranges from 37467 to 37477 (31 to 47) for VerityMap, 39167 to 39218 (1047 to 1096) for Winnomap2, 39139 to 39191 (1072 to 1124) for minimap2. The best value for each column group is indicated in bold.



Figure 2. Alignments of reads from **Cen9Sim** dataset to **Cen9_{del10}** generated by minimap2, Winnowmap2, and VerityMap shown using IGV browser (Robinson et al. 2011). minimap2 does not report any primary alignments (shown in gray) spanning the deletion breakpoint and all secondary alignments (shown in white) indicate systematic mismatches indicating mapping to an incorrect copy of the repeat. Winnowmap2 reports mappings that are incorrectly spanning the breakpoint. VerityMap either reports an alignment correctly spanning the deletion breakpoint (indicates a 10kb deletion), or clips read alignment at the breakpoint.

VerityMap correctly distinguishes haplotypes in diploid assemblies and identifies haplotype-switch errors. We aligned reads from the CenXDiploid-Sim dataset to **CenXDiploid-Switch** assembly containing a haplotype-switch error. We launched VerityMap in the special *DiploidVerityMap* mode that utilizes a more permissive strategy for selecting solid k -mers. Figure 3 illustrates that VerityMap does not incorrectly extend alignments of any reads through the haplotype-switch breakpoint (and thus detects this haplotype-switch error), while Winnowmap2 and minimap2 extend them through it.



Figure 3. Visualization of VerityMap, Winnowmap2, and minimap2 alignments of simulated reads from the CenX-Diploid dataset to the CenX-Diploid-Switch sequence using the IGV browser (Robinson et al. 2011). VerityMap, in difference from Winnowmap2 and minimap2, reveals the haplotype-switch breakpoint.

We also aligned reads from the **ChrX-RealHiFi** dataset (total number of reads 621,522) to **ChrXDiploid** assembly. VerityMap (Winnowmap2, minimap2) mapped 4474 (10,190, 10,255) reads to the incorrect haplotype and 493,012 (553,953, 552,051) — to the correct haplotype. Here, we filtered all secondary alignments, as well as all alignments with mapping quality 0. VerityMap uses a more conservative strategy for mapping reads to the diploid reference to achieve higher accuracy (that is critically important for detecting misassemblies) at the cost of mapping fewer reads. Even though VerityMap presents an advancement in accurate mapping of reads to diploid assemblies, this problem remains a difficult challenge.

VerityMap identifies assembly errors and Het sites using real HiFi datasets. The **CHM13-Cen10-interim** dataset contains an interim version of Cen10 in **CHM13** that predates the earliest publicly released CHM13 assembly by the T2T Consortium. VerityMap was used to detect a

deletion of length ~ 2.4 kbp in this version, an error that was later corrected by the T2T Consortium and is fixed in the current version of Cen10 in the **CHM13 assembly** (Figure 4).

We also aligned the CHM13-RealHiFi dataset to CHM13 assembly to reveal possible assembly errors and Het sites (Supplementary Table 4). Overall, we found 86 heterozygous sites with *Concordance* = 0.2-0.8, and one site with *Concordance* = 1. The manual validation of this site confirmed a likely 2.4 kbp insertion in the Chr19 (Figure 5, right; (Nurk et al. 2022)). Importantly, this site was not detected by the state-of-the-art methods for detecting structural variations (Mc Cartney et al., 2022).

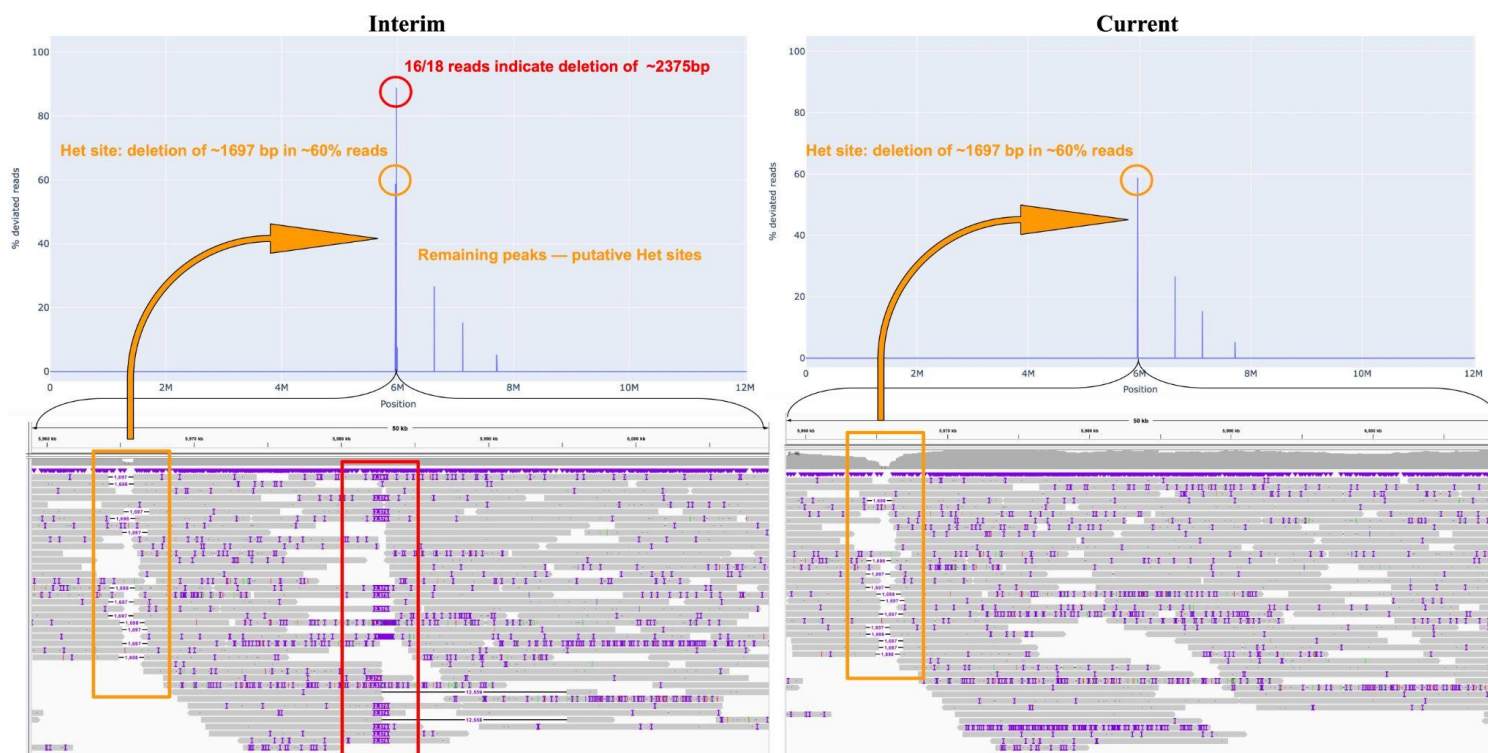


Figure 4. VerityMap detects a deletion in CHM13-Cen10-interim that is absent from the publicly released CHM13 assembly. The top panel shows the Distance Concordance test applied to **CHM13-Cen10-interim** (left) and **Cen10** in **CHM13 assembly** (right). The bottom panel shows read alignments around ~ 6 Mbp in two versions of the assembly. (Bottom left) VerityMap reveals a deletion of length ~ 2.4 kbp in **CHM13-Cen10-interim**. The remaining peaks correspond to the Het sites with various multiplicities in the assembly. In particular, one heterozygous deletion of ~ 1.7 kbp that is supported by $\sim 60\%$ of reads is only 20 kbp upstream to the misassembly site. (Bottom right) The misassembly is corrected in the current version of **Cen10** assembly, while the putative Het sites remain.

Discussion

We presented the VerityMap tool — the first accurate, error-exposing and partially diploid-aware tool for long-read mapping to complete assemblies. We utilized it to validate and improve the recently published complete human assembly during several rounds of its evaluations (Nurk et al. 2022; Mc Cartney et al. 2022). Even though (Mc Cartney et al. 2022) used VerityMap for validating this assembly, they refer to TandemMapper since VerityMap was still in development and did not even have a name at the time when this paper was submitted.

Below we discuss some limitations of VerityMap that we plan to address. All state-of-the-art read-mappers use multiple (often poorly justified) parameters for selecting a small set of k -mers in each read, weighting edges in the compatibility graph, etc. For example, minimap2, the workhorse of long-read mapping, has many parameters (including some hidden parameters) and it remains unclear how they were optimized (no description in (Li 2018)). Even though we used the default parameters for Winnowmap2 and minimap2, we do not rule out that other parameter values could improve the performance of these tools. VerityMap also suffers from the “curse of multi-parameters”: even though we carefully selected its parameters using alignments of reads to the human genome, a systematic approach using diverse organisms for learning these parameters is needed. In particular, scoring of the edges in the compatibility graph might suffer overfit to the human genome k -mer distribution.

The output of VerityMap is in standard SAM/BAM format and is, in principle, compatible with downstream variant calling tools for detection of Het variants or misassemblies. However, VerityMap detects approximate locations of a structural event rather than its exact coordinates which may present a challenge for variant calling tools especially in highly-repetitive regions such as centromeres.

The definition of a solid k -mer requires that it appears in a single contig of the assembly. Even though solid k -mers are not uniformly distributed over the human genome, our benchmarking revealed that VerityMap detected many true positive and never reported false positive misassemblies. However, this definition might be too restrictive for a highly inbred organism with long stretches of homologous

chromosomes sharing the same sequence. A highly-fragmented assembly presents a similar challenge since rare k -mers within a single chromosome might be shared between several contigs rendering them non-solid. As a result, little to no solid k -mers will be selected in such regions of the genome leading to gaps in read coverage. Allowing a solid k -mer to be shared by multiple contigs, albeit potentially more error-prone, might mitigate this issue. Similarly, the requirement that the reverse-complement of a solid k -mer does not appear in any contig, needs to be relaxed. We selected a conservative strategy for the initial VerityMap development since our primary goal was to evaluate the complete haploid complete assembly. Our next goal is to overcome these limitations and apply VerityMap for validating haplotype-resolved assemblies generated by the HPR Consortium. The described benchmarking suggests that VerityMap has a potential for effective evaluation of complete haplotype-resolved assemblies. Supplementary Note “Validation of fragmented haploid and diploid assemblies” introduces a more general k -mer indexing scheme as well as modifications to the compatibility-graph that allow validation of more general fragmented haploid or diploid assemblies (rather than limited to only complete haploid assemblies). Supplementary Note “Extended benchmarking of long-read mapping tools” presents an example of application of VerityMap to a fragmented assembly. Accurate estimation of VerityMap’s statistical power for detecting misassemblies in fragmented assemblies remains an open problem.

Even though HiFi reads allow validation of the most regions of the CHM13 assembly, ONT reads provide important complementary quality assessment (Nurk et al. 2022; Bzikadze and Pevzner 2020). Even though this paper focuses on HiFi reads, VerityMap has an ONT mode for ONT reads. We have chosen to focus on HiFi reads since ONT reads currently represent a moving target: over the last year, the error-profile of ONT reads has been rapidly improving with the latest chemistry approaching error-rate as low as 1% (Sereika et al., n.d.). While HiFi reads are more accurate (hence they allow selection of longer solid k -mers), they are shorter than ONT reads. As a result, the length of the longest paths in the compatibility graph for ONT reads can exceed those lengths for HiFi reads and thus connect solid k -mers surrounding “solid k -mer deserts” — long stretches of DNA without any solid k -mers. Together with

rapidly improving chemistry, ONT reads might become a self-sufficient alternative for quality assessment of genome assembly.

VerityMap has been designed primarily to detect errors in the assemblies rather than identify structural variations (SVs) in genomes. Even though, many misassemblies and SVs can be represented as indels, the distribution of misassemblies and SVs over the universe of all possible indels are substantially different. First, a prominent source of genomic diversity is presented by short SVs and copy number variants (CNVs) in short tandem repeats (STRs). Such differences can hardly ever be observed as misassemblies since they do not present challenge for long-read assemblers. Previous studies demonstrated that HiFi assemblers are very accurate with respect to single-nucleotide substitutions and short indels (Nurk et al., 2020, Cheng et al., 2021, Bankevich et al., 2022) and these inaccuracies in assemblies typically present polishing artifacts rather than imperfections of an assembly algorithm itself. Second, misassemblies typically involve large (multi-kbp) long indels (or inversions) that correspond to the incorrect traversal of the assembly graph (or, worse, its faulty topology). While it is possible that an alternative (and incorrect) traversal of the graph might result in a haplotype that is present in the population (the authors are aware of a single case that involves a large 600 kbp inversion on ChrY recently assembled by the T2T Consortium), it remains unclear whether this is a widespread phenomenon. Finally, some of the variations in the genome (for example, in human centromeres) are much more complex than simple CNVs of the tandem repeat units and it seems infeasible that an assembly algorithm might produce such a complicated misassembly as a result of an incorrect graph traversal.

When developing VerityMap, we have concentrated on simulated examples that resemble misassemblies that we have observed in practice of developing assembly algorithms. Specifically, (1) the indels that were introduced in the early versions of centromere assemblies (Miga et al., 2020) that were subsequently detected by an early version of TandemMapper (Mikheenko et al., 2020) and were later corrected in Bzikadze and Pevzner, 2020, (2) misassemblies that were present in interim versions of CHM13 assemblies generated by the T2T Consortium, (3) misassemblies generated by LJA assembler (Bankevich et al., 2022). Supplementary Note “Similarities and differences between detecting mutations

and detecting misassemblies” demonstrates that VerityMap can be used as an input to various SV detection tools and that VerityMap’s own DistanceDiscordant test can complement these tools.

Methods

Selecting solid k -mers. Similar to TandemMapper (Mikheenko et al. 2020), VerityMap uses rare k -mers as anchors for the read mapping. Selection of the k -mer size is dictated by the error rate in long reads. Since long error-prone reads have a high error rate, long genomic k -mers rarely “survive” without errors in such reads. For example, the probability that a 19-mer “survives” in an error-prone ONT read (generated with basecaller Guppy flip-flop 2.3.1) is only 0.34 (Bzikadze and Pevzner 2020). However, since long and accurate reads have much smaller error rates (0.7% for ONT reads generated with flow cell R10.4 <https://nanoporetech.com/accuracy> and 0.2% for HiFi reads (Nurk et al. 2020)), one can set the large default k -mer size to accurately map such reads even in highly-repetitive regions such as **Cen9** (see Supplementary Note “The choice of k -mer size and the shortest unique substrings”, Supplementary Figure 7). VerityMap sets the default value $k = 301$ and $MaxRareOccurrences = 10$ (see Supplementary Note “VerityMap parameters”).

HSAT2,3 array in the **Cen9** region is an ETR containing long inverted repeats. Importantly, for 6,344 out of ≈ 22 million rare k -mers ($k = 301$) in this region, their reverse-complementary k -mer is frequent in the genome. Since the strandedness information is not available, using such k -mers for read mapping might reduce the mapping accuracy. Similarly, utilizing a rare k -mer that is shared between two contigs (for example, two homologous or non-homologous chromosomes) might also lead to incorrect read mapping. Since VerityMap aims primarily to reduce the number of incorrect alignments, it implements a rather conservative strategy and uses solid rather than rare k -mers for read mapping to complete assemblies.

Sparse dynamic programming for read mapping. A standard dynamic programming approach for computing the *fitting alignment* of a query onto a target finds a longest path in the grid-like directed

acyclic graph (DAG) where vertices correspond to all pairs of positions (one in the target and one in the query). Analysis of the CHM13 assembly by the T2T Consortium revealed that, although the fitting alignment of a read usually finds its correct position in the assembly, it may misplace the reads that originated from highly-repetitive regions (e.g., centromeres and other ETRs), particularly in regions with assembly errors (see Supplementary Note “Frequency-aware sequence alignment scoring”, Supplementary Figure 2 and 3). Thus, mapping reads to ETRs requires a new algorithm and a new scoring approach.

Sparse dynamic programming (Gusfield 1997) is a common way to speed-up sequence alignment by selecting a small subset of points in a grid that are likely traversed by an optimal path, constructing an edge-weighted DAG on these points (rather than on all points in the entire grid), and finding a longest path in the resulting smaller graph. We note that although sparse dynamic programming approximates rather than reconstructs the optimal sequence alignment, this approximation is usually sufficient for read mapping. However, to achieve an accurate approximation, one has to (i) carefully select a vertex-set of the DAG, and (ii) carefully define the edge-weights in the sparse DAG to address the ETR mapping challenge. E.g., a k -mer-based *sparse fitting alignment* amounts to finding a longest path in the DAG where vertices correspond to all pairs of k -mer-matching positions from the target and the query (positions are k -mer-matching if k -mers starting at these positions are identical) and directed edges connects all pairs of vertices (i,j) and (i',j') with $i < i'$ and $j < j'$. However, the standard scoring approach for the sparse fitting alignment, that scores all edges with unit weights, fails to address the ETR mapping challenge since it does not account for greatly varying frequencies of k -mers in some targets.

Although the sparse fitting alignment is faster than the standard fitting alignment, it still may be computationally infeasible in long highly-repetitive regions due to a quadratic running time in the number of matching positions. Below we describe how VerityMap addresses the mapping challenge by considering matches of (undersampled) solid k -mers rather than all k -mers.

Compatible k -mers. As solid k -mers are not necessarily unique in the assembly, we consider each occurrence of each solid k -mer separately. Let a_R and b_R (a_A and b_A) be occurrences of solid k -mers a and b in the read R (contig A) such as a_R precedes b_R (a_A precedes b_A). To make a_R and b_R uniquely defined for each read, we limit attention to solid k -mers that appear exactly once in this read. We define $d(a_R, b_R)$ ($d(a_A, b_A)$) as the distance between a_R and b_R in R (a_A and b_A in A). We refer to the pair of a_A and a_R (b_A and b_R) as a *match* a_M (b_M) and write $a_M < b_M$ if a_A precedes b_A and a_R precedes b_R . We define $distance(a_M, b_M) = \min\{d(a_R, b_R), d(a_A, b_A)\}$, $diff(a_M, b_M) = |d(a_R, b_R) - d(a_A, b_A)|$, and $bias(a_M, b_M) = \ln(diff(a_M, b_M))$ (Supplementary Note “Frequency-aware sequence alignment scoring”, Supplementary Figure 2 and 3).

We call matches a_M and b_M *close* if $distance(a_M, b_M) < MaxJump$ (the default value $MaxJump = 100,000$) to exclude k -mers locating far apart from each other but preserve consecutive rare k -mers even if they are separated by a long stretch of frequent k -mers. Matches a_M and b_M are *overlapping* if $distance(a_M, b_M) < k$. Close matches a_M and b_M are *synced* if either (i) a_M and b_M are overlapping with $diff(a_M, b_M) = 0$, or (ii) they are not overlapping. Finally, we call a_M and b_M *compatible* if (i) $a_M < b_M$ and (ii) a_M and b_M are synced (see Supplementary Note “Frequency-aware sequence alignment scoring”, Supplementary Figure 2 and 3). Note that VerityMap modifies the notion of compatible k -mers introduced in (Mikheenko et al. 2020) to detect assembly errors that evade detection by TandemMapper.

Compatibility graph. Given a read R , an assembly A , and a set S of solid k -mers, VerityMap constructs a weighted *compatibility graph* $G(R, A, S)$ on the vertex-set of all pairs of identical solid k -mers from R and A . Vertices $a_M = (a_A, a_R)$ and $b_M = (b_A, b_R)$ are connected by an edge if a_M and b_M are compatible. Edge-weights in this graph should be carefully chosen since a match of a unique solid k -mer is more valuable than a match of a non-unique solid k -mer and since non-overlapping matches are more valuable than overlapping matches. VerityMap thus defines $weight(a_M, b_M) = premium(a_M, b_M) - penalty(a_M, b_M)$ in such a way that (i) $premium(a_M, b_M)$ is larger in the case when b_M represents a unique k -mer as compared to the case when it represents a non-unique k -mer, (ii) $premium(a_M, b_M)$ is larger for non-overlapping matches a_M and b_M as compared to overlapping matches, and (iii) if a_M and b_M represent a pair of k -mers

that flank misassembly breakpoints, $penalty(a_M, b_M)$ should not be too large, thus allowing VerityMap to detect a misassembly.

To address the goal (i), VerityMap assigns the score *UniqueScore* to all unique k -mers and the smaller score *RareScore* to all non-unique solid k -mers in the assembly (default values *UniqueScore* = 3 and *RareScore* = 0.1). The score $Score(b_M)$ of a match b_M is defined as the score of the k -mer it represents.

To address the goal (ii), VerityMap defines the score $EdgeScore(a_M, b_M)$ of the edge (a_M, b_M) in the compatibility graph as 1 for non-overlapping matches and as a smaller $distance(a_M, b_M) / k$ for overlapping matches. Finally, it defines $premium(a_M, b_M) = Score(b_M) * EdgeScore(a_M, b_M)$.

To ensure that the longest path in the compatibility graph correctly aligns a read against the assembly, it is important to penalize edges with large bias, e.g., to define $penalty(a_M, b_M) = bias(a_M, b_M)$. However, to address the goal (iii) and to design an error-exposing read-mapper, it is important that the penalties of edges connecting k -mers that represent misassembly breakpoints are not excessive, even in the case when these edges have a large bias. VerityMap thus limits the maximum possible penalty using the *MisassemblyPenalty* threshold (default value is 5) and defines $penalty(a_M, b_M) = \min(MisassemblyPenalty, bias(a_M, b_M))$. Supplementary Note “Finding a longest path in the compatibility graph” describes how VerityMap finds a longest path in the compatibility graph.

The challenge of finding positions of all solid k -mers in a large genome. TandemMapper uses a naive strategy of *indexing* solid k -mers in a genome *Genome* (finding a list of positions for each solid k -mer) by traversing *Genome* twice. During the first traversal, it counts all k -mers by storing them in a hash-table. During the second traversal, it checks if a k -mer is solid using the constructed hash-table and records positions of all solid k -mers (total time and memory is $O(|Genome|*k)$). This strategy becomes prohibitively time-consuming unless mapping reads against relatively short regions (e.g., several Mbp) using short k -mers (e.g., with $k < 30$). VerityMap uses a different strategy that enables mapping reads against the entire human genome using long k -mers (e.g., with $k > 300$).

The complexity of pattern matching can be reduced by indexing *rolling hashes* of k -mers rather than the k -mers themselves (Karp and Rabin 1987). Given a rolling hash function and a hash value for a k -mer $a_i \dots a_{i+k-1}$ in a genome $\dots a_i \dots a_{i+k-1} a_{i+k} \dots$, one can compute the hash value of the next k -mer $a_{i+1} \dots a_{i+k-1} a_{i+k}$ in $O(1)$ time, thus, improving the run-time complexity from $O(|Genome|*k)$ to $O(|Genome|)$. Note, that a hash collision will lead to an overestimated frequency for k -mers with collided hashes. This might lead to misclassification of some solid k -mers as non-solid, but never leads to recruiting non-solid k -mers. Unfortunately, the existing k -mer counting tools do not allow querying a count of a rolling hash instead of a k -mer, and thus do not immediately lead to a fast indexing algorithm with $O(|Genome|)$ running time. Moreover, since most of the k -mers in the non-repetitive parts of the genome are solid, we aim to select only their small subset for constructing the (small) compatibility graph (see subsection “Filtering solid k -mers in non-repetitive regions”). Unfortunately, the existing k -mer counting tools do not produce a downsampled database of k -mer counts, which leads to an unnecessary computational burden. To overcome the limitations of existing approaches, VerityMap uses a probabilistic procedure for indexing solid k -mers inspired by the Jellyfish algorithm (Marçais and Kingsford 2011). Below we briefly describe the Bloom filter and the count-min sketch data structures that VerityMap uses for indexing solid k -mers.

Bloom filter and count-min sketch. The Bloom filter (Bloom 1970) is a space-efficient probabilistic data structure that is used to test whether an element is a member of a set and that may report false positives (with a small *false positive probability*) but never false negatives. Given a genome *Genome* (or its assembly), its Bloom filter is formed by *BloomNumber* independent *hash functions*, each mapping a k -mer from *Genome* into a bit array of size *BloomSize*. Storing all k -mers in a Bloom filter allows one to quickly query whether an arbitrary k -mer occurs in the genome. Given hash functions $h_1, h_2, \dots, h_{BloomNumber}$ and a k -mer a , one can quickly check whether all bits $h_1(a), h_2(a), \dots, h_{BloomNumber}(a)$ of the Bloom filter are equal to 1, an indication that the k -mer a may occur in the genome. Thus, a Bloom filter allows one to

efficiently query whether an arbitrary k -mer occurs in the genome (with a small false positive probability).

Bloom filter allows one to check whether an element belongs to a multiset but does not evaluate how frequent it is in this multiset. The count–min sketch (CMS) is a space-efficient probabilistic data structure that provides an upper bound on the count of an element in a multiset (Cormode and Muthukrishnan 2005). CMS is a modification of a single-bit Bloom filter array that uses a *counter* array of size $CMSNumber * CMSSize$ with cells containing several bits. This property allows it to count repeatedly inserted elements: when an element is inserted, the values of all $CMSNumber$ hash functions are computed and the corresponding cells in the counter arrays are incremented by 1 (in case an increment leads to an overflow, the corresponding cell gets frozen). To provide an upper bound for the count of a given element, CMS computes the values of all $CMSNumber$ hash functions and reports the *minimum* of the corresponding cells in the counter arrays (a frozen cell votes for the *infinity* count). Thus, CMS sometimes overestimates the count of an element but never underestimates it. VerityMap uses CMS to distinguish unique, rare, and frequent k -mers and with a controlled probability misclassifies a rare (unique) k -mer as frequent (rare).

Identification of rare k -mers using CMS. For each *Contig* in genome *Genome*, VerityMap counts all k -mers in this contig by storing them in the count-min-sketch $CMS(Contig) = CMS(Contig, k, Bits, CMSNumber, CMSSize)$ formed by $CMSNumber$ hash functions, each mapping a k -mer into a counter array of size $CMSSize$ where each counter takes $Bits$ bits. The memory footprint of CMS is $CMSNumber * Bits * CMSSize$ bits. The total run-time for inserting all k -mers from *Contig* into this CMS is $O(|Contig| * CMSNumber * Bits * k)$. Storing rolling hashes instead of k -mers further reduces the complexity to $O(|Genome| * CMSNumber * Bits)$. Supplementary Note “VerityMap parameters” describes how VerityMap sets the parameters of the CMS and the Bloom filters.

Identification of solid k -mers using the Bloom filter. To verify whether a rare k -mer is solid, VerityMap checks that (i) its reverse-complement does not appear in the genome and (ii) it appears only in a single contig. Storing all k -mers from the genome $Genome$ in a Bloom filter allows VerityMap to efficiently query whether an arbitrary k -mer is present in $Genome$. In order to check conditions (i) and (ii) for identifying solid k -mers, VerityMap constructs two Bloom filters — $BanBloomFilter(Genome)$ and $OnceBloomFilter(Genome)$ of an *optimal size* (Bloom 1970) for the expected number of k -mers ($|Genome|$ is an upper bound) and for the target false positive probability $BloomFPP$ (default value $BloomFPP = 0.00001$).

For every rare k -mer a , we ensure that its reverse complement a^* does not appear in the $Genome$ by inserting a^* in the $BanBloomFilter$ and, in the end, checking that a is not present in $BanBloomFilter$. To check if a k -mer a appears in a single contig, VerityMap scans all contigs and synchronously inserts a k -mer a from a contig $Contig$ in both $OnceBloomFilter(Genome)$ and $CMS(Contig)$. This synchronous construction of both $OnceBloomFilter(Genome)$ and $CMS(Contig)$ allows one to check whether the k -mer a occurs in a single contig, albeit with a small probability of error. Indeed, if it is already present in $OnceBloomFilter(Genome)$ but missing in $CMS(Contig)$, it likely appeared in a previously scanned contig. In this case, VerityMap filters out the k -mer a by inserting it into $BanBloomFilter(Genome)$. However, if it is present in both $OnceBloomFilter(Genome)$ and $CMS(Contig)$, the k -mer a is retained since it may occur in a single contig.

To reduce the complexity, instead of inserting k -mers in the Bloom filter, we store their rolling hashes. The total running time of finding all solid k -mers and determining their genomic positions is $O(|Genome| * (BloomHashNumber + CMSNumber))$. The total memory footprint is $O(|Genome| + (CMSNumber * CMSSize * Bits + BloomHashNumber * BloomSize))$.

Since the Bloom filter is a probabilistic data structure, some solid k -mers might be falsely removed as breaking either condition (i) or (ii). In practice, very few solid k -mers are filtered out (as controlled by the parameter $BloomFPP$) and that does not lead to deterioration of read mapping. Importantly, no true non-solid k -mer might be misclassified as solid. For example, a naive implementation

of an “exact” strategy (that uses deterministic albeit more memory consuming data structures) extracts 149,548,373 solid k -mers from ChrX in CHM13 assembly, while the “approximate” strategy described above (based on Bloom filters and CMSs) extracts 147,105,002 solid k -mers from the same dataset (more than 98% of truly solid k -mers). For the sake of easy comparison, we did not downsample solid k -mers in this example as described in Supplementary Note “Downsampling solid k -mers in non-repetitive regions”.

Distance Concordance test. VerityMap detects approximate locations of misassembly breakpoints without requiring a time-consuming nucleotide-level alignment. Indeed, if a read R spans a misassembly breakpoint, a longest path in the compatibility graph $G(R,A,S)$ often provides a hint for finding this breakpoint. For example, a deletion of length L in an assembly typically triggers an edge (a_M, b_M) with a large difference $\text{diff}(a_M, b_M) \cong L$, where a_M and b_M represent solid k -mers flanking the deletion breakpoint in such a way that a_M (b_M) precedes (follows) the deletion breakpoint as illustrated in Figure 1.

However, a single read may have non-zero values of $\text{diff}(a_M, b_M)$ even in the case of error-free assembly, especially when $d(a_R, b_R)$ is high. However, this difference is usually low for accurate HiFi reads and error-free assemblies, suggesting that large differences can be used for detecting misassemblies. To demonstrate that it is indeed the case, we considered 18 millions of unique k -mers ($k=301$) in **Cen9** and selected 100,000 pairs of unique k -mers such that k -mers in each pair are 10,000 bp apart in the assembly. The mean distance between the same pairs of k -mers in reads was $10,002 \pm 3.4$ bp. Thus, $\text{diff}(a_M, b_M)$ usually does not exceed $3 * 3.4 = 10.2$ bp if $d(a_A, b_A) = 10,000$ bp. Given our upper limit for a distance between two consecutive k -mers equal to 100,000 bp, we set a lower threshold for an indel size *MinIndelSize* to 100 bp. Thus, if nearly all longest paths containing k -mers a and b show a *systematic bias* in values of $\text{diff}(a_M, b_M)$ and $\text{diff}(a_M, b_M) > \text{MinIndelSize}$, then an indel of approximate size $\text{diff}(a_M, b_M)$ is likely contained between a_A and b_A . We classify each read that shows systematic bias in values of $\text{diff}(a_M, b_M)$ as a *discordant* read and report the fraction of discordant reads connecting each pair of solid k -mers in the assembly.

Formally, for each read R and for each pair a_M and b_M of consecutive *matches* in a longest path in compatibility graph we add the number $diff(a_M, b_M)$ to a set $Diff_s(a_A, b_A)$. We test the hypothesis that the mean of $Diff_s(a_A, b_A)$ is 0 and refer to the pair a_A and b_A as *significant* if this hypothesis is rejected (two-sided one-sample t -test). We use Bonferroni family-wise error rate correction and set up significance level α_{DC} (the default value $\alpha_{DC} = 0.05$). For each significant pair a_A and b_A of consecutive k -mers in a contig A , we compute the proportion $Concordance(a_A, b_A)$ of the most frequent difference $TopDiff(a_A, b_A)$ in $Diff_s(a_A, b_A)$. Figure 5 (left and middle) shows the $Concordance(a_A, b_A)$ across the **Cen9**_{del10} region when aligning simulated reads from the dataset **Cen9Sim** (**Cen9Sim-Het**) and reveals the artificial deletion (Het site) of length 10 kbp at position 20 Mbp. Figure 5 (right) shows the $Concordance(a_A, b_A)$ across cen19 when aligning real HiFi reads and reveals the likely assembly error and several Het sites (see Results).




Figure 5. Distance concordance test applied to simulated and real read-sets. (Left) Mapping reads from the dataset **Cen9Sim** reveals an artificial deletion at location 20 Mbp. $Concordance(a_A, b_A) = 1$ and $TopDiff(a_A, b_A) = 10,003$ for solid k -mer a_A at position 19,999,691 and b_A at position 20,000,001 revealing the approximate location of the deletion of size 10 kbp. (Middle) Mapping reads from the dataset **Cen9Sim-Het** reveals an artificial Het site at position 20 Mbp with $Concordance(a_A, b_A) = 0.33$ for the same solid k -mers a_A and b_A and $TopDiff(a_A, b_A) = 10,005$ pointing at insertion in **Cen9**_{del10} of size 10 kbp existing in one of the artificial haplotypes. (Right) Mapping HiFi reads reveals several Het variants and a likely assembly error (with $Concordance = 1$ at 28.5 Mbp) in the centromere of Chromosome 19. (see Supplementary Table 1, Supplementary Note “Extended benchmarking of long-read mapping tools” for description of all datasets).

Software Availability. The codebase of VerityMap version 2.1.2 generated in this study has been submitted to GitHub at <https://github.com/ablab/VerityMap>, to Zenodo <https://doi.org/10.5281/zenodo.6469695>, and to Supplemental Code.

Competing interests. The authors declare that they have no competing interests.

Acknowledgements. This work was supported by St. Petersburg State University, St. Petersburg, Russia [ID PURE 73023672]. We are indebted to Daniel Baker, Daniel Lemire, and Thomas Mueller for their suggestion to use count-min-sketch. We are grateful to Alexander Bzikadze and Cynthia Wu for helpful discussions and suggestions.

Authors' contributions. A.V.B. and A.M. conducted the experiments and wrote the code for VerityMap. P.A.P. supervised the research. All authors worked on the development of the VerityMap algorithm, wrote and edited the manuscript.

References

- Bankevich, Anton, Andrey V. Bzikadze, Mikhail Kolmogorov, Dmitry Antipov, and Pavel A. Pevzner. 2022. "Multiplex de Bruijn Graphs Enable Genome Assembly from Long, High-Fidelity Reads." *Nature Biotechnology*, February. <https://doi.org/10.1038/s41587-022-01220-6>.
- Bickhart, Derek M., Mikhail Kolmogorov, Elizabeth Tseng, Daniel M. Portik, Anton Korobeynikov, Ivan Tolstogonov, Gherman Uritskiy, et al. 2022. "Generating Lineage-Resolved, Complete Metagenome-Assembled Genomes from Complex Microbial Communities." *Nature Biotechnology*, January. <https://doi.org/10.1038/s41587-021-01130-z>.
- Bloom, Burton H. 1970. "Space/time Trade-Offs in Hash Coding with Allowable Errors." *Communications of the ACM* 13 (7): 422–26.
- Bushmanova, Elena, Dmitry Antipov, Alla Lapidus, Vladimir Suvorov, and Andrey D. Prjibelski. 2016. "naQUAST: A Quality Assessment Tool for de Novo Transcriptome Assemblies." *Bioinformatics* 32 (14): 2210–12.
- Bzikadze, Andrey V., Alla Mikheenko, and Pavel A. Pevzner. 2022. "VerityMap: Fast and Accurate Mapping of Long Reads to Complete Genome Assemblies." Zenodo. <https://doi.org/10.5281/ZENODO.6469229>.
- Bzikadze, Andrey V., and Pavel A. Pevzner. 2020. "Automated Assembly of Centromeres from Ultra-Long Error-Prone Reads." *Nature Biotechnology* 38 (11): 1309–16.
- Cheng, Haoyu, Gregory T. Concepcion, Xiaowen Feng, Haowen Zhang, and Heng Li. 2021. "Haplotype-Resolved de Novo Assembly Using Phased Assembly Graphs with Hifiasm." *Nature Methods* 18 (2): 170–75.
- Chin, Chen-Shan, Paul Peluso, Fritz J. Sedlazeck, Maria Nattestad, Gregory T. Concepcion, Alicia Clum, Christopher Dunn, et al. 2016. "Phased Diploid Genome Assembly with Single-Molecule Real-Time Sequencing." *Nature Methods* 13 (12): 1050–54.
- Cormode, Graham, and S. Muthukrishnan. 2005. "An Improved Data Stream Summary: The Count-Min Sketch and Its Applications." *Journal of Algorithms & Computational Technology* 55 (1): 58–75.
- Gurevich, Alexey, Vladislav Saveliev, Nikolay Vyahhi, and Glenn Tesler. 2013. "QUAST: Quality Assessment Tool for Genome Assemblies." *Bioinformatics* 29 (8): 1072–75.
- Gusfield, D. 1997. "Algorithms on Strings, Trees, and Sequences Cambridge University Press." *Cambridge, England*.
- Jain, Chirag, Arang Rhie, Nancy F. Hansen, Sergey Koren, and Adam M. Phillippy. 2022. "Long-Read Mapping to Repetitive Reference Sequences Using Winnowmap2." *Nature Methods*, April. <https://doi.org/10.1038/s41592-022-01457-8>.
- Jain, Chirag, Arang Rhie, Haowen Zhang, Claudia Chu, Brian P. Walenz, Sergey Koren, and Adam M. Phillippy. 2020. "Weighted Minimizer Sampling Improves Long Read Mapping." *Bioinformatics* 36 (Suppl_1): i111–18.
- Jain, Miten, Sergey Koren, Karen H. Miga, Josh Quick, Arthur C. Rand, Thomas A. Sasani, John R. Tyson, et al. 2018. "Nanopore Sequencing and Assembly of a Human Genome with Ultra-Long Reads." *Nature Biotechnology* 36 (4): 338–45.
- Karp, Richard M., and Michael O. Rabin. 1987. "Efficient Randomized Pattern-Matching Algorithms." *IBM Journal of Research and Development* 31 (2): 249–60.
- Kokot, Marek, Maciej Długosz, and Sebastian Deorowicz. 2017. "KMC 3: Counting and Manipulating K-Mer Statistics." *Bioinformatics* 33 (17): 2759–61.
- Kolmogorov, Mikhail, Derek M. Bickhart, Bahar Behsaz, Alexey Gurevich, Mikhail Rayko, Sung Bong Shin, Kristen Kuhn, et al. 2020. "metaFlye: Scalable Long-Read Metagenome Assembly Using Repeat Graphs." *Nature Methods* 17 (11): 1103–10.
- Kolmogorov, Mikhail, Jeffrey Yuan, Yu Lin, and Pavel A. Pevzner. 2019. "Assembly of Long, Error-Prone Reads Using Repeat

- Graphs.” *Nature Biotechnology* 37 (5): 540–46.
- Koren, Sergey, Brian P. Walenz, Konstantin Berlin, Jason R. Miller, Nicholas H. Bergman, and Adam M. Phillippy. 2017. “Canu: Scalable and Accurate Long-Read Assembly via Adaptive K-Mer Weighting and Repeat Separation.” *Genome Research* 27 (5): 722–36.
- International Human Genome Sequencing Consortium. 2001. Initial sequencing and analysis of the human genome. *Nature* 409: 860–921.
- Langmead, Ben, and Steven L. Salzberg. 2012. “Fast Gapped-Read Alignment with Bowtie 2.” *Nature Methods* 9 (4): 357–59.
- Li, Heng. 2016. “Minimap and Miniasm: Fast Mapping and de Novo Assembly for Noisy Long Sequences.” *Bioinformatics* 32 (14): 2103–10.
- Li, Heng. 2018. “minimap2: Pairwise Alignment for Nucleotide Sequences.” *Bioinformatics* 34 (18): 3094–3100.
- Li, Heng. 2021. “New Strategies to Improve minimap2 Alignment Accuracy.” *Bioinformatics*, October. <https://doi.org/10.1093/bioinformatics/btab705>.
- Li, Heng, and Richard Durbin. 2009. “Fast and Accurate Short Read Alignment with Burrows-Wheeler Transform.” *Bioinformatics* 25 (14): 1754–60.
- Lin, Yu, Jeffrey Yuan, Mikhail Kolmogorov, Max W. Shen, Mark Chaisson, and Pavel A. Pevzner. 2016. “Assembly of Long Error-Prone Reads Using de Bruijn Graphs.” *Proceedings of the National Academy of Sciences of the United States of America* 113 (52): E8396–8405.
- Logsdon, Glennis A., Mitchell R. Vollger, Pingshun Hsieh, Yafei Mao, Mikhail A. Liskovych, Sergey Koren, Sergey Nurk, et al. 2021. “The Structure, Function and Evolution of a Complete Human Chromosome 8.” *Nature* 593 (7857): 101–7.
- Mapleson, Daniel, Gonzalo Garcia Accinelli, George Kettleborough, Jonathan Wright, and Bernardo J. Clavijo. 2017. “KAT: A K-Mer Analysis Toolkit to Quality Control NGS Datasets and Genome Assemblies.” *Bioinformatics* 33 (4): 574–76.
- Marçais, Guillaume, and Carl Kingsford. 2011. “A Fast, Lock-Free Approach for Efficient Parallel Counting of Occurrences of K-Mers.” *Bioinformatics* 27 (6): 764–70.
- Mc Cartney, Ann M., Kishwar Shafin, Michael Alonge, Andrey V. Bzikadze, Giulio Formenti, Arkarachai Functamman, Kerstin Howe, et al. 2022. “Chasing Perfection: Validation and Polishing Strategies for Telomere-to-Telomere Genome Assemblies.” *Nature Methods*, March. <https://doi.org/10.1038/s41592-022-01440-3>.
- Miga, Karen H. 2019. “Centromeric Satellite DNAs: Hidden Sequence Variation in the Human Population.” *Genes* 10 (5). <https://doi.org/10.3390/genes10050352>.
- Miga, Karen H. 2020. “Centromere Studies in the Era of ‘Telomere-to-Telomere’ Genomics.” *Experimental Cell Research* 394 (2): 112127.
- Miga, Karen H., Sergey Koren, Arang Rhie, Mitchell R. Vollger, Ariel Gershman, Andrey Bzikadze, Shelise Brooks, et al. 2020. “Telomere-to-Telomere Assembly of a Complete Human X Chromosome.” *Nature* 585 (7823): 79–84.
- Miga, Karen H., Yulia Newton, Miten Jain, Nicolas Altemose, Huntington F. Willard, and W. James Kent. 2014. “Centromere Reference Models for Human Chromosomes X and Y Satellite Arrays.” *Genome Research* 24 (4): 697–707.
- Mikheenko, Alla, Andrey V. Bzikadze, Alexey Gurevich, Karen H. Miga, and Pavel A. Pevzner. 2020. “TandemTools: Mapping Long Reads and Assessing/improving Assembly Quality in Extra-Long Tandem Repeats.” *Bioinformatics* 36 (Suppl_1): i75–83.
- Mikheenko, Alla, Andrey Pribelski, Vladislav Saveliev, Dmitry Antipov, and Alexey Gurevich. 2018. “Versatile Genome Assembly Evaluation with QUAST-LG.” *Bioinformatics* 34 (13): i142–50.
- Mikheenko, Alla, Vladislav Saveliev, and Alexey Gurevich. 2016. “MetaQUAST: Evaluation of Metagenome Assemblies.” *Bioinformatics* 32 (7): 1088–90.
- Nisa, Israt, Prashant Pandey, Marquita Ellis, Leonid Olikier, Aydın Buluç, and Katherine Yelick. 2021. “Distributed-Memory K-Mer Counting on GPUs.” In *2021 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 527–36.
- Nurk, Sergey, Sergey Koren, Arang Rhie, Mikko Rautiainen, Andrey V. Bzikadze, Alla Mikheenko, Mitchell R. Vollger, et al. 2022. “The Complete Sequence of a Human Genome.” *Science* 376 (6588): 44–53.
- Nurk, Sergey, Brian P. Walenz, Arang Rhie, Mitchell R. Vollger, Glennis A. Logsdon, Robert Grothe, Karen H. Miga, Evan E. Eichler, Adam M. Phillippy, and Sergey Koren. 2020. “HiCanu: Accurate Assembly of Segmental Duplications, Satellites, and Allelic Variants from High-Fidelity Long Reads.” *Genome Research* 30 (9): 1291–1305.
- Ono, Yukiteru, Kiyoshi Asai, and Michiaki Hamada. 2021. “PBSIM2: A Simulator for Long-Read Sequencers with a Novel Generative Model of Quality Scores.” *Bioinformatics* 37 (5): 589–95.
- Rhie, Arang, Brian P. Walenz, Sergey Koren, and Adam M. Phillippy. 2020. “Mercury: Reference-Free Quality, Completeness, and Phasing Assessment for Genome Assemblies.” *Genome Biology* 21 (1): 245.
- Robinson, James T., Helga Thorvaldsdóttir, Wendy Winckler, Mitchell Guttman, Eric S. Lander, Gad Getz, and Jill P. Mesirov. 2011. “Integrative Genomics Viewer.” *Nature Biotechnology* 29 (1): 24–26.
- Ruan, Jue, and Heng Li. 2020. “Fast and Accurate Long-Read Assembly with wtdbg2.” *Nature Methods* 17 (2): 155–58.
- Sereika, Mantas, Rasmus Hansen Kirkegaard, Søren Michael Karst, Thomas Yssing Michaelsen, Emil Aarre Sørensen, Rasmus Dam Wollenberg, and Mads Albertsen. n.d. “Oxford Nanopore R10.4 Long-Read Sequencing Enables near-Perfect Bacterial Genomes from Pure Cultures and Metagenomes without Short-Read or Reference Polishing.” <https://doi.org/10.1101/2021.10.27.466057>.
- Shafin, Kishwar, Trevor Pesout, Ryan Lorig-Roach, Marina Haukness, Hugh E. Olsen, Colleen Bosworth, Joel Armstrong, et al. 2020. “Nanopore Sequencing and the Shasta Toolkit Enable Efficient de Novo Assembly of Eleven Human Genomes.” *Nature Biotechnology* 38 (9): 1044–53.

- Suzuki, Hajime, and Masahiro Kasahara. 2018. "Introducing Difference Recurrence Relations for Faster Semi-Global Alignment of Long Sequences." *BMC Bioinformatics* 19 (Suppl 1): 45.
- Venter, J. C., M. D. Adams, E. W. Myers, P. W. Li, R. J. Mural, G. G. Sutton, H. O. Smith, et al. 2001. "The Sequence of the Human Genome." *Science* 291 (5507): 1304–51.
- Wenger, Aaron M., Paul Peluso, William J. Rowell, Pi-Chuan Chang, Richard J. Hall, Gregory T. Concepcion, Jana Ebler, et al. 2019. "Accurate Circular Consensus Long-Read Sequencing Improves Variant Detection and Assembly of a Human Genome." *Nature Biotechnology* 37 (10): 1155–62.
- Wick, Ryan R., Louise M. Judd, Claire L. Gorrie, and Kathryn E. Holt. 2017. "Unicycler: Resolving Bacterial Genome Assemblies from Short and Long Sequencing Reads." *PLoS Computational Biology* 13 (6): e1005595.