



## Highly accurate assembly polishing with DeepPolisher

Mira Mastoras, Mobin Asri, Lucas Brambrink, et al.

*Genome Res.* 2025 35: 1595-1608 originally published online May 19, 2025

Access the most recent version at doi:[10.1101/gr.280149.124](https://doi.org/10.1101/gr.280149.124)

---

**References** This article cites 51 articles, 7 of which can be accessed free at:  
<http://genome.cshlp.org/content/35/7/1595.full.html#ref-list-1>

**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](#).



---

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

## Method

## Highly accurate assembly polishing with DeepPolisher

Mira Mastoras,<sup>1</sup> Mobin Asri,<sup>1</sup> Lucas Brambrink,<sup>2</sup> Prajna Hebbar,<sup>1</sup> Alexey Kolesnikov,<sup>2</sup> Daniel E. Cook,<sup>2</sup> Maria Nattestad,<sup>2</sup> Julian Lucas,<sup>1</sup> Taylor S. Won,<sup>1</sup> Pi-Chuan Chang,<sup>2</sup> Andrew Carroll,<sup>2</sup> Benedict Paten,<sup>1</sup> Kishwar Shafin,<sup>2</sup> and the Human Pangenome Reference Consortium<sup>3</sup>

<sup>1</sup>UC Santa Cruz Genomics Institute, University of California, Santa Cruz, California 95060, USA; <sup>2</sup>Google Incorporated, Mountain View, California 94043, USA

Accurate genome assemblies are essential for biological research, but even the highest-quality assemblies retain errors caused by the technologies used to construct them. Base-level errors are typically fixed with an additional polishing step that uses reads aligned to the draft assembly to identify necessary edits. However, current methods struggle to find a balance between over- and underpolishing. Here, we present an encoder-only transformer model for assembly polishing called DeepPolisher, which predicts corrections to the underlying sequence using Pacific Biosciences (PacBio) HiFi read alignments to a diploid assembly. Our pipeline introduces a method, PHAsing Reads in Areas Of Homozygosity (PHARAOH), which uses ultralong Oxford Nanopore Technologies (ONT) data to ensure alignments are accurately phased and to correctly introduce heterozygous edits in falsely homozygous regions. We demonstrate that the DeepPolisher pipeline can reduce assembly errors by approximately half, mostly driven by reductions in indel errors. We have applied our DeepPolisher-based pipeline to 180 assemblies from the next Human Pangenome Reference Consortium (HPRC) data release, producing an average predicted quality value (QV) improvement of 3.4 (54% error reduction) for the majority of the genome.

[Supplemental material is available for this article.]

Faithful reconstruction of the genome is increasingly essential for scientists to understand the underlying biology of any organism. Genome assembly, the process of digitally reconstructing a genome, is accomplished by piecing together segments of DNA (called sequencing “reads”) and obtaining a consensus. The past several years have seen a huge improvement in the quality of genome assembly, driven by long-read sequencing technology, which produces reads that are orders of magnitude longer than the short-read technologies that preceded them (Li and Durbin 2024; Taylor et al. 2024). Assembly algorithms have continued to advance alongside the improvements in sequencing, and it is now becoming routine to produce automated, fully phased, highly contiguous, and near-complete genome assemblies (Liao et al. 2023; Rautiainen et al. 2023; Cheng et al. 2024; Li and Durbin 2024; Porubsky et al. 2024; Taylor et al. 2024).

The quality of a genome assembly can be measured by an array of different metrics that assess different aspects of the reconstruction. Contiguity metrics describe the cumulative length and number of contigs of the assembly, whereas completeness metrics evaluate the extent of the expected sequence present in the assembly (Li and Durbin 2024). For organisms with multiple copies of their chromosomes (ploidy), fully resolving the haplotypes is a key component to accurate genome assembly, and chromosome-phasing metrics assess this (Rhie et al. 2020; Cheng et al. 2021). Another critical set of metrics evaluates the base-level accuracy of the assembled sequence (Rhie et al. 2020).

All sequencing technologies contain errors that are not completely random but are influenced by biases toward different

genomic features. Even using a consensus approach, recurrent errors in the reads can propagate into the assembly. Low complexity regions like homopolymers and tandem repeats are particularly challenging for Oxford Nanopore Technologies (ONT) (Kolmogorov et al. 2023) and Pacific Biosciences (PacBio) long reads (Wenger et al. 2019), and certain motifs such as those high in GC content can cause dropout, a phenomenon particularly prevalent in short reads like Illumina (Benjamini and Speed 2012; Chen et al. 2013). Small errors can also be introduced by biases in the assembler, and different assembly algorithms differ in their limitations and error characteristics (Jarvis et al. 2022; Rautiainen et al. 2023; Cheng et al. 2024).

Base-level errors interfere with scientists’ ability to make accurate conclusions from their data. For example, they can disrupt functional elements, as in protein-coding genes, leading to frameshifts, missense, and nonsense mutations (Aganezov et al. 2022; Miller et al. 2022; Wagner et al. 2022b). For projects like the Human Pangenome Reference Consortium (HPRC) (Liao et al. 2023) and Vertebrate Genome Project (VGP) (Rhie et al. 2021), whose goal is to produce haplotype-resolved population-scale reference genomes for entire groups of species, fixing errors is especially critical to avoid propagating mistakes to all the downstream analyses that make use of them, for example, by proposing sequence variations that are, in fact, errors (Church et al. 2015; Schneider et al. 2017; Aganezov et al. 2022).

The process of removing these small-scale errors in assemblies is known as “polishing.” Most methods for assembly polishing involve aligning reads back to the draft assembly to identify

<sup>3</sup>A complete list of the HPRC authors appears at the end of this paper. Corresponding authors: [awcarroll@google.com](mailto:awcarroll@google.com), [bpaten@ucsc.edu](mailto:bpaten@ucsc.edu), [shafin@google.com](mailto:shafin@google.com)

Article published online before print. Article, supplemental material, and publication date are at <https://www.genome.org/cgi/doi/10.1101/gr.280149.124>.

© 2025 Mastoras et al. 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/>.

potential sequence changes suggested by the reads. Because of the complex nature of sequencing and assembly error profiles, most polishing tools employ heuristic algorithms and models that need to be specific to a given sequencing platform in order to achieve high accuracy. Popular polishers like Racon (Vaser et al. 2017), Medaka (<https://github.com/nanoporetech/medaka>), PolishCLR (Rhie et al. 2021), and CONSENT (Morisse et al. 2021) were developed to correct assemblies using noisy long reads, whereas tools like Pilon (Walker et al. 2014), POLYPOLISH (Wick and Holt 2022), and POLCA (Zimin and Salzberg 2020) operate solely on short reads. Recently, new approaches have begun to emerge for correcting the state-of-the-art telomere-to-telomere (T2T) assemblies generated with a combination of PacBio high fidelity (HiFi) and ONT long reads. To polish the first fully complete human genome (CHM13), the T2T consortium leveraged DeepVariant (Poplin et al. 2018) calls from multiple sequencing technologies and a careful filtering strategy to avoid false-positive edits (Formenti et al. 2022; Mc Cartney et al. 2022). Algorithms like NextPolish2 (Hu et al. 2024) have also been developed to specifically target T2T genome polishing using HiFi reads. However, an ongoing major challenge for all polishing approaches is striking a balance between overcorrecting and undercorrecting assemblies.

Seq2seq transformer family models have been used to great success in many applications including language processing, machine translation tasks, and conversational AI (Brandes et al. 2022; Vaswani et al. 2023; Zhang et al. 2023). Previously, this model architecture was used in DeepConsensus to improve the quality of PacBio HiFi sequencing reads (Baid et al. 2022). In this work, we demonstrate a new adaptation of the transformer model for human and primate genome polishing called DeepPolisher, which takes HiFi sequencing reads aligned to a draft assembly as input. Along with DeepPolisher, we have developed a pipeline called PHARAOH for improving the phasing accuracy of HiFi read alignments to diploid assemblies in falsely homozygous regions using ONT ultralong (UL) reads. In this work, we show that DeepPolisher with PHARAOH outperforms current polishing approaches, reducing assembly errors by half, which is driven mostly by reductions in assembly insertion–deletion (indel) errors. We applied our pipeline to 180 haplotype-resolved assemblies from the next HPRC data release and show substantial improvements in accuracy.

## Results

### Overview of the DeepPolisher pipeline

An overview of the DeepPolisher pipeline is shown in Figure 1. First, PacBio HiFi reads are aligned to the diploid assembly with minimap2 (Li 2018) or winnowmap (Jain et al. 2020). For most of the genome, this alignment step is sufficient to assign the HiFi reads to the correct haplotype. However, a common error mode of modern assemblers is to mistakenly represent heterozygous sequence as homozygous. When these regions of false homozygosity are longer than the HiFi read (up to ~25 kb), the aligner cannot phase it and will randomly assign it to a haplotype, preventing any necessary polishing edits from being made there. To address this, we developed a pipeline called PHASing Reads in Areas Of Homozygosity (PHARAOH) that uses phasing information from ONT reads >100 kb to infer the correct haplotype for reads in long homozygous regions (Methods). For the rest of the genome, the tool Secphase (Liao et al. 2023) is used within the PHARAOH pipeline to ensure correct read phasing by revisiting secondary

alignments and calculating a marker consistency score to determine whether the read should be relocated to the other haplotype.

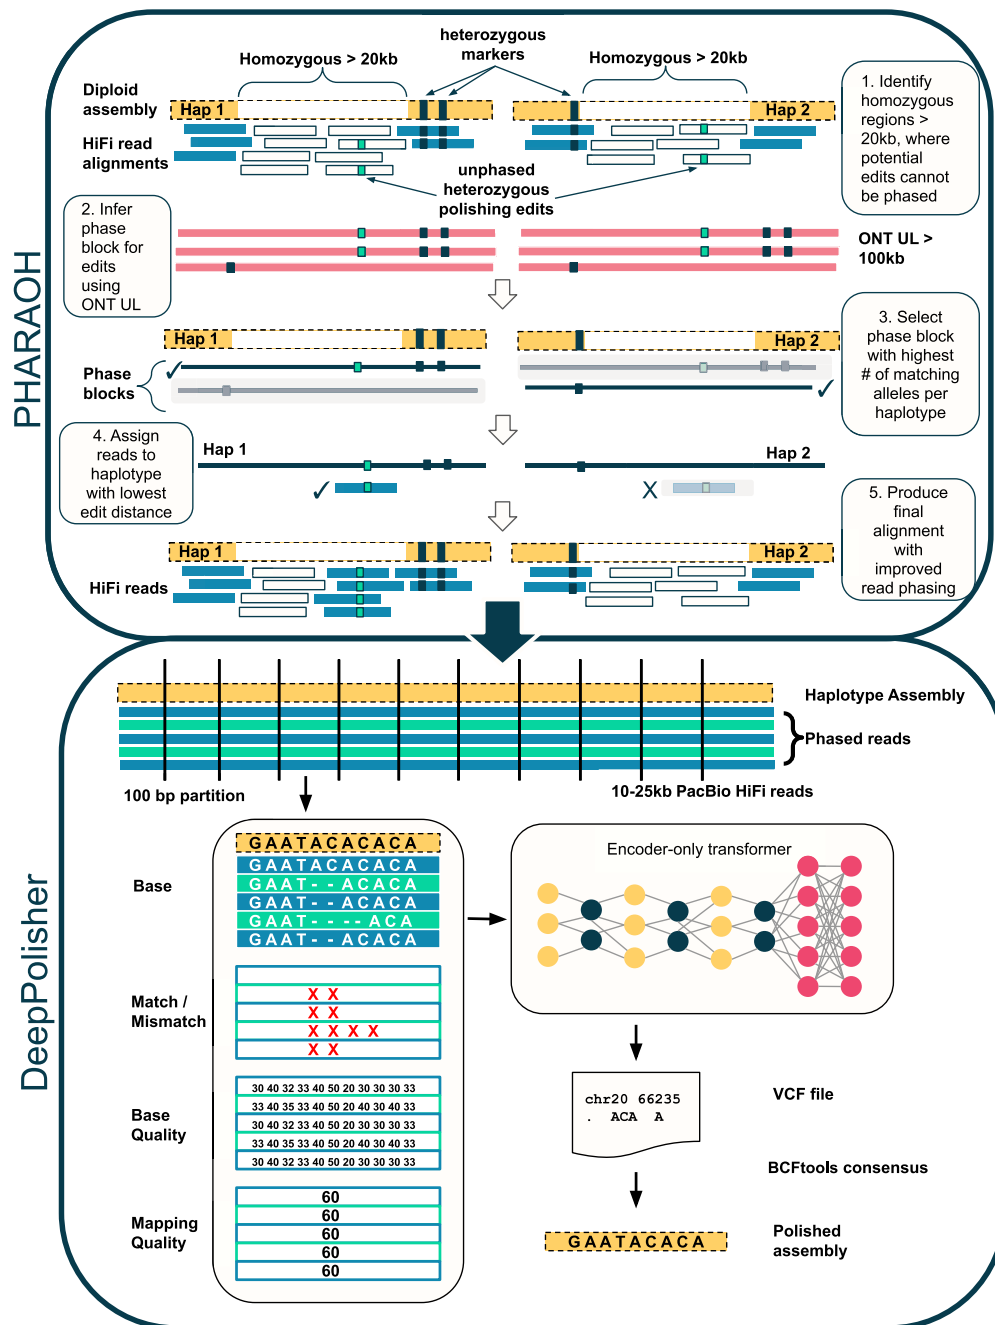
PHARAOH and Secphase allow us to provide improved read alignments as input to the DeepPolisher model. DeepPolisher first partitions these alignments into 100 bp segments and then represents them as a tensor object with auxiliary input features consisting of the base, whether the base is a match or mismatch to the assembly, the base quality, and the mapping quality. The tensor is provided to the encoder-only transformer neural network, which produces proposed assembly corrections in VCF format. These can be applied to the unpolished assembly with BCFtools (Danecek et al. 2021) consensus (Fig. 1; Supplemental Fig. 3). To train DeepPolisher we used a high-accuracy assembly of HG002 produced by the Q100 consortium (<https://github.com/marbl/HG002>) (v0.9), training across Chromosomes 1–19 (Methods). DeepPolisher can be run without PHARAOH if ONT UL reads are not available; however, assembly corrections in long stretches of false homozygosity will not be made.

### Alignment-based comparison of DeepPolisher and alternate polishing approaches

We compared the performance of DeepPolisher with existing polishing approaches for HiFi reads: the T2T consortium pipeline (Mc Cartney et al. 2022), DeepVariant using HiFi alignments, and NextPolish2. To test them, we ran the UL version of hifiasm (which uses both HiFi and ONT UL reads) (Cheng et al. 2024) to generate an assembly for a sample that was not included in our training set (HG005) and polished it with each method. A comparison of the computational resources used by each polishing approach on the HG005 hifiasm assembly can be found in Supplemental Table 1. After polishing with each tool, we called variants relative to GRCh38 using dipcall (Li et al. 2018), which enabled us to assess concordance with the Genome in a Bottle (GIAB) v4.2.1 benchmark call set (Methods) (Wagner et al. 2022a). This gives an alignment-based measure of polishing accuracy that spans ~80% of the HG005 genome, as defined by the GIAB high-confidence regions.

With DeepPolisher, we reduce the number of variant-calling errors in the unpolished assembly from 20,274 to 11,750, indicating a reduction in assembly error rate by 42% (8.14 errors per megabase to 4.72 errors per megabase) (Fig. 2A). This performance is reproduced on the held-out chromosome (20) from HG002 (Supplemental Table 2A,B). We also trained a DeepPolisher model on a Verkko (Rautiainen et al. 2023) assembly using the same data. According to the GIAB variant-calling results, the Verkko assembly for HG005 starts out at a higher quality than hifiasm (6.86 errors/megabase), and after DeepPolisher, it approaches the same quality as the polished hifiasm assembly at 4.87 errors/megabase (Supplemental Table 2E).

In comparison to DeepPolisher, the T2T polishing pipeline produces more conservative improvements, with only 3591 polishing edits passing their filters, leading to a removal of 372 errors (7.99 errors/megabase). Using DeepVariant to polish with HiFi alignments reduces assembly errors to 13,724 (5.51 errors/megabase). We note that much of the difference in GIAB variant-calling performance between DeepVariant and DeepPolisher on HG005 is owing to PHARAOH: Polishing the HG005 assembly with DeepVariant run on PHARAOH alignments produces a reduction to 11,508 assembly errors (4.62 errors/megabase). However, an important caveat is that DeepVariant was trained on several GIAB samples including HG005, with only Chr 20–22 held out, so this performance cannot be separated from training circularity.

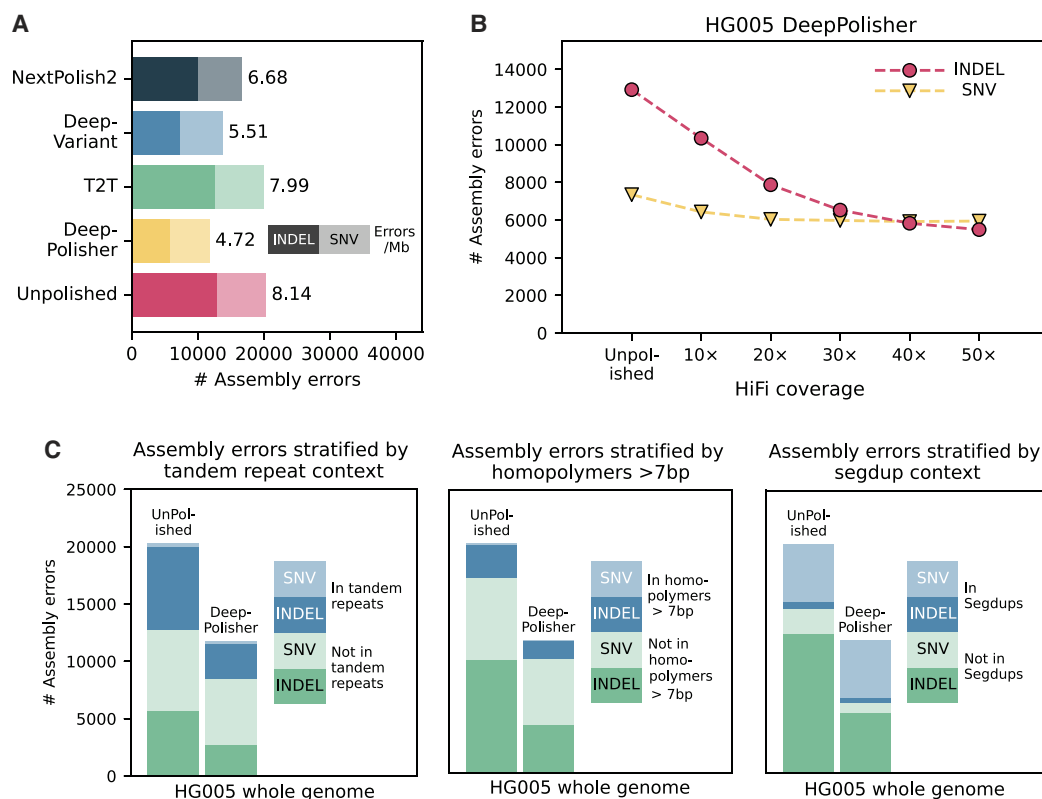


**Figure 1.** DeepPolisher pipeline overview. The PHARAOH pipeline leverages phase block information from ONT UL reads to correct the haplotype assignment of PacBio HiFi reads. The corrected alignment is passed to DeepPolisher, which is an encoder-only transformer model that predicts the underlying assembly sequence and proposes corrections in VCF format.

Finally, NextPolish2 removed 3260 errors from the unpolished assembly (6.68 errors/megabase) (Fig. 2A; Supplemental Table 2C).

We tested DeepPolisher across a variety of coverages and found that although 40× coverage produces optimal reduction in variant-calling errors, some assembly improvement can still be obtained with as low as 10× coverage (Fig. 2B). The DeepPolisher pipeline performance is also reproducible across HiFi read versions other than what it was trained on (DeepConsensus v1.2). This includes Sequel II, DeepConsensus v0.2, and Revio with DeepConsensus v1.1 (Supplemental Figs. 1, 2; Supplemental Tables 3, 8).

We stratified the GIAB variant-calling errors by genomic region and found that the majority of indel errors removed by DeepPolisher were in tandem repeats and homopolymers. The majority of SNV errors that were not fixed by DeepPolisher were located within segmental duplications. Some of these may be true errors, but others may be mapping artifacts or mistakes in the truth set (Fig. 2C; Supplemental Table 4). To understand DeepPolisher's performance in repetitive sequences that are excluded from the GIAB high-confidence regions, we assessed the concordance of the HG002 hifiasm assembly with the Q100 HG002 benchmark



**Figure 2.** Comparison of DeepPolisher and alternate polishing methods against GIAB v4.2.1 benchmark for HG005. (A) For each polishing method, GIAB v4.2.1 variant-calling (assembly) errors are separated by indels (darker shade) and single-nucleotide variants (SNVs; lighter shade), with the number of errors per megabase to the right of each bar. (B) Total GIAB variant-calling (assembly) errors for different HiFi read coverages, with indel errors represented in pink circles and SNV errors in yellow triangles. (C) Total GIAB variant-calling (assembly) errors stratified by presence in tandem repeats (*left*), homopolymers >7 bp (*middle*), and segmental duplications (segdups; *right*), with SNV errors in lighter shades and indel errors in darker shades.

genome within centromere annotations (Methods). Although it is expected that the centromere is enriched for errors within both assemblies, DeepPolisher improves the overall concordance with the Q100 genome within centromeres, removing 85 SNVs (out of 21,073 total) and 443 INDELS (out of 8516 total). (Supplemental Table 5).

To test how DeepPolisher impacts the coding regions, we extracted genes containing frameshift and premature stop mutations in the HG005 and HG002 assemblies before and after polishing. (Supplemental Table 6A). In the unpolished assemblies, we identified a mean of 106 genes with frameshifts and 46 with stop codons per haplotype. A mean of 77.87% of genes with frameshifts and 79.31% of genes with stop codons was supported by the GIAB v4.2.1 variant calls, indicating that the substantial majority are likely true mutations. These estimates are in line with those previously reported by the HPRC (Liao et al. 2023). Across the two samples, DeepPolisher reduced false frameshifts by a net seven and stop codons by three (Supplemental Table 6A). Out of 129 total frameshift and stop codons across both samples that are unvalidated by GIAB and that DeepPolisher is not fixing, 100 are outside of GIAB high-confidence regions, and 15 are coming from known polymorphic gene families (such as OR, KRT, FAM). We manually investigated 14 of these putative errors and noted that in all cases the HiFi reads appear to support the assembly.

To directly assess the impact of polishing on every gene, we extracted all transcripts that were altered after polishing, and

calculated their DNA identity and protein identity to the transcript in GRCh38. Of these, we excluded transcripts with only synonymous or nonsynonymous mutations and examined those for which the transcripts were impacted significantly by frameshifts or in-frame insertions and deletions (Supplemental Table 6B). Polishing improved the protein and DNA identity to the original transcript in GRCh38 for seven out of eight of these transcripts. In the one transcript in which polishing decreased the protein identity, the read alignments looked reliable and supported the polishing edit.

To assess the performance benefits of the PHARAOH pipeline, we ran DeepPolisher directly on minimap2 alignments. Without PHARAOH, DeepPolisher removes 8005 variant-calling errors from the unpolished HG002 assembly (32% error reduction) and 6917 errors from the unpolished HG005 assembly (39.5% error reduction). Using PHARAOH alignments as input to DeepPolisher removes an additional 1915 errors for HG002 (9% error reduction) and 519 for HG005 (2.5% error reduction). (Supplemental Table 2B,G). The impact of PHARAOH is expected to differ between assemblies based on the amount of false homozygosity that is present.

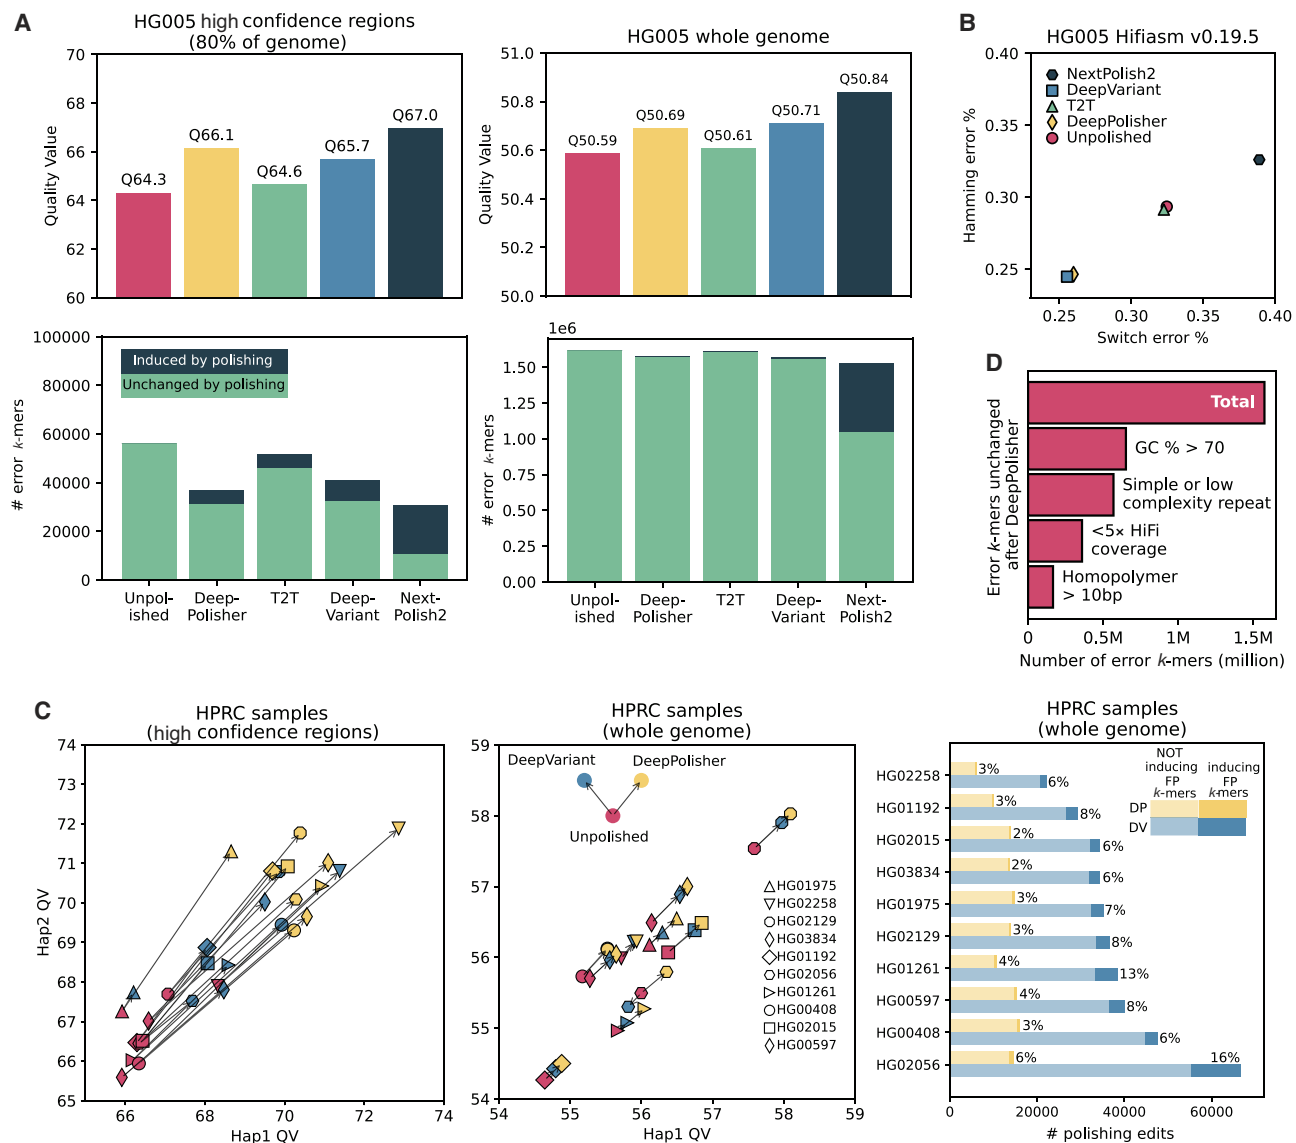
#### *k*-mer-based comparison of DeepPolisher and alternate polishing methods

*k*-mer-based methods assess the accuracy of an assembly using an alignment-free approach (Rhie et al. 2020). To compare the *k*-mer and alignment-based assessments, we first restricted the *k*-

mer assessment to just the assembled sequence contained within the GIAB high-confidence regions, which span ~80% of the HG005 assembly (Methods). *k*-mer methods generally report a quality value (QV), a log base 10-scaled measure of assembly error, with higher values indicating more accurate assemblies. We find that DeepPolisher improves the *k*-mer-based assembly QV by 1.8, equivalent to a 34% reduction in errors, which is approximately consistent with the quantity of improvement suggested by the GIAB variant-calling performance (Fig. 3A; Supplemental Table 7). Out of all the polishing methods compared, DeepPolisher induces the least amount of error *k*-mers to the assembly (Fig. 3A). NextPolish2, which uses Illumina *k*-mers in its algorithm for as-

sembly polishing, produces the highest QV improvement, yet with the trade-off of inducing the highest number of new error *k*-mers to the assembly. The introduction of new error *k*-mers likely lies behind the increase in switch and hamming error rate produced by NextPolish2 (0.03% and 0.06%, respectively). In contrast, the other polishing methods manage to reduce switch and hamming error rates (Fig. 3B).

Because the DeepVariant model was trained on the HG005 GIAB variants, it is critical to also compare its performance against DeepPolisher on samples that were completely held out from training in both methods. We applied both tools to the same PHARAOH alignments for eight of the HPRC samples and assessed



**Figure 3.** *k*-mer-based comparison of DeepPolisher and alternate polishing approaches for HG005. (A) Top panels display QV scores for each polishing method. Bottom panels depict total error *k*-mers, divided by error *k*-mers induced by polishing (dark blue) and error *k*-mers unchanged after polishing (green). Left panels show results for the GIAB high-confidence regions; right panels, whole genome. (B) Switch (x-axis) and hamming (y-axis) error rates for each polishing method. (C) Comparison of DeepVariant and DeepPolisher for eight HPRC samples. Left and middle panels show Hap1 (x-axis) and Hap2 (y-axis) QV for eight HPRC samples, with an arrow connecting the unpolished QV (pink) to the QV after polishing with DeepVariant (blue) and DeepPolisher (yellow). The left panel is within the GIAB high-confidence regions; middle panel, whole genome. The right panel shows the number of polishing edits from DeepPolisher (yellow) and DeepVariant (blue). Lighter shades indicate edits not inducing error (FP) *k*-mers; darker shades show edits that induce error *k*-mers. (D) Number of error *k*-mers unchanged by polishing with DeepPolisher falling into sequence annotation categories.

their QV. For all samples, DeepPolisher produces a higher QV improvement (average 3.95 QV, 60% error reduction) within the high-confidence regions compared with DeepVariant (average 2.26 QV, 41% error reduction) (Fig. 3C). In addition, the percentage of edits inducing error *k*-mers is more than twice as high for DeepVariant (average 9%) than DeepPolisher (average 4%) (Fig. 3C; Supplemental Table 9).

### *k*-mer-based assessments underestimate assembly error rates

Relative to alignment-based metrics like those used in the GIAB assessment above, *k*-mer-based QVs are unbiased by alignment and easy to assess. However, *k*-mer-based methods lack an underlying ground truth, instead relying on consistency between assembly and sequencing technology. This results in a potential limitation in which results could be consistent but not correct owing to technology-specific sequencing bias. Because *k*-mer-based quality metrics are agnostic to the sequence context outside of the *k*-mer window, they also are limited in their ability to capture errors in repetitive sequence or in regions with a loss of heterozygosity.

For the unpolished HG005 assembly within the high-confidence regions, we see a QV estimate of 66.97 using Illumina *k*-mers (*k*=31), suggesting 0.811 errors/megabase, whereas the GIAB variant analysis suggests the same assembly has about 10× more errors per megabase (8.14; equivalent to a QV of 50.9) (Supplemental Table 10). Using shorter *k*-mer sizes or a hybrid *k*-mer set built with both Illumina and HiFi reads as proposed by Mc Cartney et al. (2022) led to substantially more inflated estimates of accuracy (0.258 and 0.201 errors/megabase, respectively; Methods) (Supplemental Table 10). To explore the source of this difference in error rate, we performed a manual analysis investigating the concordance between the errors reported by the QV and the GIAB variants for HG005 after polishing with DeepPolisher (Methods) (Supplemental Table 11). We found that error *k*-mers reported by QV were almost always validated by errors reported using GIAB, indicating that most of the QV errors within the high-confidence regions are likely true. However, for 35% of the cases examined, the QV did not flag the pre- nor postpolishing sequence as an error. Because it is impossible for both versions of the assembly sequence to be correct, we attribute this to the fact that *k*-mer-based quality estimates often miss errors in repetitive sequences, owing to their presence in the “truth” *k*-mer set coming from other regions of the genome. Therefore, we conclude that the non-unique *k*-mer sequence is likely the cause of the underestimated error rates in QV metrics.

### *k*-mer-based methods identify residual errors in unpolished sequence

Looking at the *k*-mer QV whole genome, we found little difference between the unpolished assembly and the four different polishing methods, with all QV scores in the range of 50.56–50.84. This is because the vast majority of the error *k*-mers leading to this QV score is unchanged after polishing (97.3% for DeepPolisher) (Fig. 3A). Investigating these unchanged error *k*-mers, we found that 41% have >70% GC content, and when we plotted the distribution of their GC content, we saw a clear bias relative to randomly permuted *k*-mers (Methods) (Supplemental Fig. 4). This indicates that the observed GC bias (Benjamini and Speed 2012; Chen et al. 2013) in short reads is likely inducing a substantial number of falsely reported errors. However, this still leaves a majority of these error *k*-mers that may be real, unaddressed errors. Indeed,

we find that 22.8% of predicted error *k*-mers are located in regions with <5× coverage of HiFi data. These regions were likely bridged by ONT sequence in the assembly stage and therefore have a higher error rate but are inaccessible to DeepPolisher owing to the lack of HiFi read coverage. Additionally, we found that 36% of the unchanged error *k*-mers overlapped simple and low-complexity repeat annotations, with 10% overlapping homopolymers >10 bp. These regions tend to have higher error rates in HiFi reads (Wenger et al. 2019; Stoler and Nekrutenko 2021), which makes it likely these are true errors that DeepPolisher was unable to fix. Together, these observations indicate that although the residual predicted error *k*-mers contain many false positives, the *k*-mer-based methods are identifying significant sequence outside of GIAB high-confidence regions that is erroneous and unchanged by our current polishing methods (Fig. 3D; Supplemental Table 12).

### Polishing the HPRC release 2 assemblies

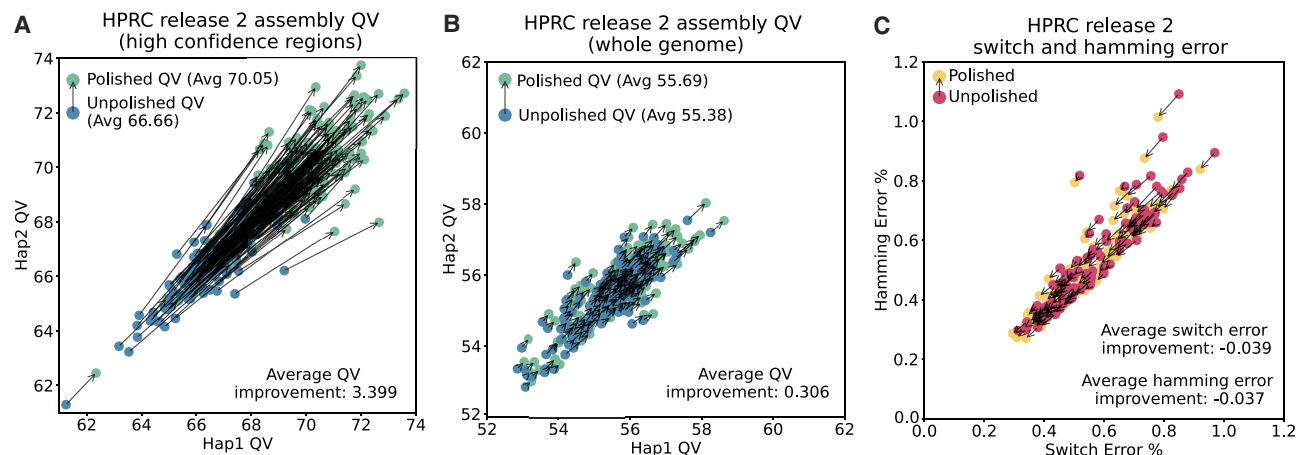
We applied our polishing pipeline to 180 HPRC assemblies made available as part of the second HPRC data release. Within the high-confidence regions (81.4% of the genome on average), and consistent with our analysis of HG005, we noted an average QV improvement of 3.4 (54% reduction in errors), with the average unpolished QV of 66.66 moving to an average of 70.05 after polishing (Fig. 4A). Whole genome, we improved the QV by 0.306 points on average (7% reduction in errors) (Fig. 4B), owing to the presence of the lower-quality ONT-only regions that cannot be polished with the current model and regions subject to sequencing bias as previously discussed. Polishing improved phasing accuracy for all 111 samples that contained trio data: On average, it reduced switch error by 0.039 and hamming errors by 0.037 (Fig. 4C; Supplemental Table 13).

### Application to non-human primate assemblies

We generated hifiasm assemblies for two marmoset individuals, Baguette and Maisie, using HiFi and ONT long-read and Illumina short-read data derived from fibroblast cell lines, in order to test the DeepPolisher pipeline’s performance on non-human species. Whole genome, DeepPolisher improved the QV for Baguette from 57.1495 to 57.5501 (9% error reduction) and for Maisie from 57.0333 to 57.5419 (11% error reduction), which is comparable to the whole-genome error reduction observed for the HPRC assemblies (Supplemental Table 14). To assess the performance in regions accessible to the current DeepPolisher model, we excluded from the QV calculation the regions in which HiFi coverage of mapping quality greater than one was lower than 5×. For Baguette, this excluded 5.8% of the genome and resulted in a QV improvement from 59.0761 to 59.7212 (14% error reduction). For Maisie, this excluded 3.3% of the genome and resulted in a QV improvement from 59.3859 to 60.2831 (18.7% error reduction).

### Removing additional errors with element data

Element Biosciences AVITI (Element) short-read data have gained popularity recently for its high accuracy in tandem repeats and homopolymers. We tested whether further polishing with Element data after DeepPolisher could fix residual errors. Because of the challenges associated with phasing short-read data, we restricted our initial experiment to polishing just homozygous locations. To obtain potential homozygous edits, we aligned 50× Element Cloudbreak reads (Carroll et al. 2023) to each haplotype of the



**Figure 4.** Polishing results for 180 HPRC assemblies. (A) Hap1 QV (*x*-axis) and Hap2 QV (*y*-axis) in the high-confidence regions for 180 HPRC samples from the second release. For each sample, unpolished QV is in blue with an arrow pointing to the polished QV. (B) The same as A but for whole-genome QV. (C) Switch (*x*-axis) and hamming (*y*-axis) error rate for the 107 samples with trio data. Unpolished in pink with an arrow pointing to polished in yellow.

HG002 and HG005 polished assemblies. We ran DeepVariant on each alignment and selected just the homozygous-alternate calls with a GQ greater than seven. This approach removed 9% (1182) of the remaining variant-calling errors after polishing with DeepPolisher for HG002 and 5% (600) for HG005, a promising number of errors removed given the relatively low number of polishing edits applied (5586 for HG002 and 8795 for HG005) (Supplemental Table 15).

## Discussion

Rapid advancements in long-read sequencing and assembly algorithms in recent years are ushering in a new era of population-scale, T2T genome assembly projects (Ebert et al. 2021; Rhie et al. 2021; Liao et al. 2023). Many of these assemblies are being used as reference genomes, making it important that the sequences be highly accurate. To remove the remaining errors in these sequences, deep learning approaches are a promising way to model sequencing bias in regions like homopolymers and low-complexity repeats. In this work, we developed a new polishing tool called DeepPolisher, an encoder-only transformer model, that takes in HiFi reads aligned to a diploid assembly and predicts the underlying sequence. Accompanying this, we introduced a pipeline called PHAsing Reads in Areas Of Homozygosity (PHARAOH) that ensures read alignment inputs to DeepPolisher are assigned to the correct haplotype and phases potential polishing edits in long homozygous assembly regions.

Using an alignment-based assessment of assembly accuracy measured against the GIAB v4.2.1 benchmarking variant calls, we demonstrate that DeepPolisher reduces assembly errors by approximately half, largely driven by reductions in indel errors, across the large majority of sequence that is mappable by HiFi sequencing. Although alignment-based assessments of high-quality benchmark sets are close to a gold standard, current benchmarks do not cover the full genome and are dependent on accurate alignment. To get around these issues, we also assessed assembly quality with *k*-mer methods. Comparing *k*-mer methods to the alignment-based assessment within the same genomic regions predicted similar percentage reductions in errors, which is consistent and reassuring, but also revealed that the *k*-mer methods miss around  $\geq 90\%$  of residual errors. Although some of these may be false pos-

itives in the alignment benchmark, our manual analysis suggests that most are simply missed by the *k*-mer methods owing to their limitations regarding modeling repeat *k*-mers that are likely often enriched in assembly errors. We would therefore urge users to recognize this and note that such QV estimates are likely overly optimistic because they have a high false negative rate.

However, *k*-mer methods also reveal large numbers of predicted errors in regions that remain untouched with the current iteration of DeepPolisher and are located outside of the regions assessed by the alignment-based benchmarks. Although nearly half of these predicted errors may be false positives caused by GC content bias, it is probable that many are real. We attribute a quarter of these errors to regions of HiFi coverage dropout in which ONT was primarily used during the assembly stage. We plan to expand the DeepPolisher model to work on ONT reads in the future to address this fraction of errors. Around one-third of the residual errors lie in homopolymers and tandem repeats, sequences that are challenging even for highly accurate HiFi reads. We experimented with using Element Biosciences data to fix these errors and demonstrate that we can remove a minority of errors using the approach. However, Element reads are currently relatively short, preventing accurate mapping across repetitive sequence and limiting phasing. We predict the remaining class of errors will be challenging to fix.

To resolve the limitations of current alignment and *k*-mer-based approaches to assessing base-level assembly accuracy, the Q100 project (<https://github.com/marbl/HG002>) aims to create a “perfect” genome assembly to be used as a benchmark for assembly methods. Although achieving a completely perfect assembly will be challenging, the direct comparison of a draft assembly to an assembly benchmark will help expand the alignment-based assessment of assembly quality to more difficult regions of the genome than are covered by the GIAB high-confidence regions. In the future, we hope that genome benchmarks for more samples will be produced to expand the training and testing sets to multiple genomes for models like DeepPolisher.

Despite the difficulties inherent in assessing the performance benefits of the DeepPolisher model, we demonstrate that it provides the best balance between over- and underpolishing for high-quality assembly projects like the HPRC. We have applied the DeepPolisher and PHARAOH pipeline to the next release of HPRC assemblies and demonstrate consistent improvements to QV across all samples,

which will be critical for downstream applications of this new reference data set. We have demonstrated that the current iteration of DeepPolisher, although trained on human data, is applicable to non-human primate species, which may be useful for projects like the VGP (Rhie et al. 2021) and the T2T consortium's T2T-Primates project (Yoo et al. 2025). We anticipate that for expansion into more distantly related clades, new models will need to be trained.

## Methods

### Assembly

To create HG002 and HG005 hifiasm assemblies for training and testing DeepPolisher, we ran hifiasm (UL) v0.19.5 in trio mode with the default parameters. To create HG002 and HG005 Verkko assemblies for training and testing DeepPolisher, we ran Verkko v2.0 in trio mode with the `--correct-overlap-batches 64` and `--screen human` and otherwise default parameters. As input we used 40× HiFi DeepConsensus v1.2 and 40× ONT UL >100 kb reads from the HPRC. The ONT reads for the HG005 and HG002 samples were sequenced using R9.4.1 flow cells and basecalled using Guppy software (HG005 with version 5.0.7 and HG002 with version 6.0.6). For phasing assembly haplotypes in the trio mode, both assemblers needed *k*-mer databases that were created from parental short reads, and Verkko required *k*-mer databases from the child's short reads. For hifiasm we used yak and for Verkko we used meryl to create the related *k*-mer databases. For HG002, we used the 300× parental and child Illumina data from the HPRC to create input yak and meryl databases (using default settings for both) and used the 100× Illumina data for HG005 parental and child databases. All data used to generate and polish the assemblies were downloaded from the HPRC ([https://s3-us-west-2.amazonaws.com/human-pangenomics/index.html?prefix=working/HPRC\\_PLUS/](https://s3-us-west-2.amazonaws.com/human-pangenomics/index.html?prefix=working/HPRC_PLUS/)). The HPRC release 2 assemblies and their associated data were downloaded ([https://github.com/human-pangenomics/hprc\\_intermediate\\_assembly/tree/main/data\\_tables](https://github.com/human-pangenomics/hprc_intermediate_assembly/tree/main/data_tables)).

### DeepPolisher development and training

DeepPolisher uses an encoder-only transformer model to predict potential errors present in a haplotype-resolved genome assembly. The framework of DeepPolisher is adopted from our previous work DeepConsensus (Baid et al. 2022). The input to DeepPolisher is an alignment file in which the reads are aligned to the correct haplotype assembly. Then, it performs the following steps to identify potential errors:

- 1. Make images.** In the `make_images` step
  - DeepPolisher takes a 25 kb window of the assembly and tries to find potential positions at which there could be an error. This is done by comparing the reads to the assembly. If at any position, there are more than two reads containing a sequence that does not match the assembly (mismatched base, insertion or deletion), we pick that position to be a potential error site.
  - The positions that are within 50 bp or closer are grouped together to create a window that does not exceed 100 bp sequence.
  - For each window, we take alignment features (base, base quality, mapping quality, match, or mismatch) and create a tensor-like representation of the window that potentially contains assembly errors.
  - This tensor-like representation of the alignment is then fed to the transformer-model for prediction.

- 2. Inference.** In the inference step, we take the tensor-like representation of the 100 bp window with potential sequencing errors and use that as an input to the encoder-only transformer model. The model predicts a sequence based on the read alignments. We then compare the predicted sequence to the assembly, and any difference between the predicted sequence and the assembly is reported in a VCF format file, which would indicate potential errors in the assembly. These VCF edits can be applied to the assembly after DeepPolisher in a separate step using BCFtools (Danecek et al. 2021) consensus -H2.

We trained DeepPolisher on the HG002 hifiasm assembly taking the HG002-T2T-v0.9 (2024) as the truth. We took HiFi reads aligned to the hifiasm assembly using the PHARAOH pipeline. Then, we aligned the hifiasm diploid assembly to the HG002-T2T-v0.9 assembly to create an assembly-to-truth alignment, which we use as the truth sequence for training the model. During training, the predicted sequence is compared against the truth sequence, and the loss is calculated via the alignment loss function we introduced in our previous work DeepConsensus (Baid et al. 2022). We use a similar approach to train a model for Verkko assembly, in which we use the Verkko assembly for training.

### PHARAOH pipeline

The first step in the PHARAOH workflow involves aligning all HiFi reads to the diploid assembly using minimap2 with parameters `"-L --eqx --cs -c -k19 -x map-hifi."` Next, we identify stretches of homozygosity in the assembly by aligning the two haplotype FASTA files to each other with minimap2 using the parameters `"-L --eqx --cs -c -x asm5."` The resulting PAF file is passed into a Python script ([https://github.com/mobinasri/secphase/blob/main/programs/src/find\\_homozygous\\_regions.py](https://github.com/mobinasri/secphase/blob/main/programs/src/find_homozygous_regions.py)) that parses the cigar string to return stretches of 100% identical sequence >20,000 bp (the average length of HiFi reads). Next, the reads within the homozygous regions are extracted and aligned separately to the maternal and paternal haplotypes using minimap2 with parameters `"-L --eqx --cs -c -k19 -map-hifi."` All reads with gap-compressed mismatch ratio exceeding 0.02 are removed to avoid calling spurious variants in reads too diverged from the assembly. DeepVariant is then used to detect heterozygous variants in these alignments, and all variants with a genotype quality less than 10 are filtered out. We align ONT UL reads >100,000 bp separately to each haplotype assembly using minimap2 with parameters `"--cs --eqx -L -Y -map-ont."` These UL alignments are passed to WhatsHap (Martin et al. 2016), which is used to phase the heterozygous variants called by DeepVariant in homozygous regions.

To reassign reads to the correct haplotype assembly using the phased variants from WhatsHap, we added a new mode to the tool Secphase. From WhatsHap, we have phased variants grouped into phase blocks. Each phase block contains two haplotypes, and Secphase picks the haplotype with a higher number of reference alleles (i.e., the one that is more similar to the assembly sequence) covered by that phase block. Secphase then retains the variants with alternative alleles in the selected haplotype of the phase block. Next, a variant block is created around each selected variant. The minimum block size can be set as a parameter (default = 100), and the actual size is adjusted to be twice the variant size if the minimum block size is smaller than twice the variant size. Overlapping variant blocks are merged, so a single block may encompass multiple variants. Each variant block, defined in assembly coordinates, is then projected onto the read coordinates. A single read may have multiple alignments, typically one per assembly haplotype, with each alignment having its own projected blocks. To create nonoverlapping blocks crucial for distinguishing

assembly haplotypes, all projected variant blocks are merged on the read coordinates. These merged blocks serve as proxies to compare haplotypes. The variants are then applied to the assembly haplotypes, effectively “pseudopolishing” the assembly only in the created blocks. The pseudopolished sequence of each haplotype is then compared against the read sequence within each variant block, with edit distances calculated using the Edlib (Šošić and Šikić 2017) library. The edit distances across all merged variant blocks are summed for each alignment. The haplotype with the lower total edit distance is selected as the correct haplotype for the related read.

In the PHARAOH pipeline, this new variant mode of Secphase is used for the homozygous regions >20 kb containing a potential heterozygous polishing edit, and the original marker mode is used for the rest of the genome. After reassigning reads to the new location determined by Secphase, we remove alignments with a gap-compressed mismatch ratio exceeding 0.002, producing the final alignment for input to DeepPolisher. PHARAOH is implemented as a wdl workflow and is publicly available at GitHub (<https://github.com/miramastoras/PHARAOH>). The tool Secphase may be found at GitHub (<https://github.com/mobinasri/Secphase>). The entire pipeline for running PHARAOH and DeepPolisher may be found at GitHub ([https://github.com/human-pangenomics/hpp\\_production\\_workflows/tree/master/polishing](https://github.com/human-pangenomics/hpp_production_workflows/tree/master/polishing)).

### DeepPolisher GQ filter optimization

We selected four HPRC samples (HG01975, HG04115, HG02129, HG01993) with varying QV improvements after polishing (Supplemental Table 16) to optimize a set of GQ filters for DeepPolisher. First, we annotated the error  $k$ -mers in the assemblies by whether they were induced, fixed or did not change with polishing. To accomplish this we first aligned each haplotype of the unpolished assembly to the corresponding haplotype of the polished assembly with minimap2 and parameters `-x asm5 -L --eqx --cs -c`. We took the polished assembly `*_only.bed` file produced by Merqury (which contains the assembly coordinates for the error  $k$ -mers), merged the  $k$ -mers with BEDTools (Quinlan and Hall 2010) `merge -c 1`, and projected them to the unpolished assembly using the script [https://github.com/mobinasri/flagger/blob/main/programs/src/project\\_blocks\\_multi\\_thread.py](https://github.com/mobinasri/flagger/blob/main/programs/src/project_blocks_multi_thread.py). This allowed us to have the polished assembly error  $k$ -mers in unpolished assembly coordinates. We then subtracted the polished projected error  $k$ -mer BED file from the raw error  $k$ -mer BED file with BEDTools `subtract -A` to extract the error  $k$ -mers fixed by polishing. We performed the reverse to obtain the error  $k$ -mers induced by polishing. Finally, we used BEDTools `intersect` to produce the error  $k$ -mers unchanged by polishing, which were common to both BED files. To annotate the polishing edits by whether they induced, fixed, or did not change error  $k$ -mers, we intersected the polishing VCF with the error  $k$ -mer BED files labeled in the previous step.

We found that the majority of DeepPolisher edits inducing error  $k$ -mers were 1 bp insertions or deletions (Supplemental Fig. 4). Based on this observation, we tested three filtering scenarios: (1) a single GQ filter for all variants, (2) a GQ filter for 1 bp insertions and a separate one for the remaining edits, and (3) a separate GQ filter for 1 bp insertions, 1 bp deletions, and all other edits. We combined the annotated polishing edits from all four samples, and for each of these scenarios, we swept the GQ filter cutoff from zero to 25 and calculated the estimated QV improvement using the formula  $-10 \times \log_{10}(\text{error\_k-mer\_after/error\_k-mer\_before})$ . We found the optimal GQ cutoff to be GQ 20 for 1 bp insertions, GQ 12 for 1 bp deletions, and GQ 5 for all other variants (Supplemental Table 16). All

DeepPolisher results shown in the paper have this filter applied unless otherwise stated.

### Implementing alternate polishing methods

We adapted the T2T polishing pipeline described by Mc Cartney et al. (2022) for diploid assemblies to compare against DeepPolisher. We aligned 40× HiFi DeepConsensus v1.2 reads from the HPRC for both cell lines to the diploid hifiasm assemblies with winnowmap v2.03 and parameters `--cs --eqx -L -Y -I8g -map-pb`. We aligned ONT R9 40× reads from the HPRC to the diploid hifiasm assembly using winnowmap v2.03 and parameters `--cs --eqx -I8g -Y -L -p0.5 -map-hifi`. We also aligned 30× Illumina data to the diploid hifiasm assembly with BWA-MEM v.0.7.17 and default parameters. HiFi and Illumina BAM files were merged with SAMtools, and DeepVariant v1.6.1 with parameter `--model-type=HYBRID_PACBIO_ILLUMINA` was used to call variants. PEPPER-Margin-DeepVariant was used to call variants for the ONT alignments. All GQ and VAF filters described by Mc Cartney et al. (2022) were applied. Merfin mode `-polish` was also used to further filter polishing edits. Edits were applied to the assembly with BCFtools `consensus -H1`.

We used the same 40× HiFi DeepConsensus v1.2 winnowmap read alignments described above as input to NextPolish2 and DeepVariant. For NextPolish2, 30× Illumina data from the HPRC were first preprocessed with the tool `fastp` (Chen et al. 2018) with parameters `-f 5 --cut_front --cut_tail` as suggested by the authors (<https://github.com/Nextomics/NextPolish2>). The preprocessed data were then used to create a yak  $k=21$  and  $k=31$  database as input to NextPolish2, using default settings. NextPolish2 was run with the default settings. DeepVariant v1.6.1 was run with the default settings and flag `--model-type=PACBIO`. For DeepVariant, BCFtools `consensus -H1` was used to apply edits to the original assembly.

### Measuring GIAB variant-calling performance

To assess the accuracy of the assemblies against the GIAB v4.2.1 benchmark set, we first ran `dipcall` (Li et al. 2018) with the default settings against GRCh38 obtained from <https://ftp-trace.ncbi.nlm.nih.gov/ReferenceSamples/giab/release/references/GRCh38/>. We next ran the tool `hap.py` (<https://github.com/Illumina/hap.py>) to calculate performance metrics of the `dipcall` VCF file against the GIAB v4.2.1 benchmark VCF.

For HG002, we restricted the `hap.py` analysis to regions within the GIAB high-confidence regions subset that were also validated by the T2T Q100 v1.0.1 (<https://github.com/marbl/HG002>) assembly. To create this subset, we ran `dipcall` (Li et al. 2018) on the T2T v1.0.1 assembly against GRCh38 using minimap parameters `-z200000,10000` (to align across larger SVs and more divergent regions like the MHC). We next intersected the HG002 high-confidence regions BED file provided by GIAB with the BED file produced by `dipcall` in the previous step to produce a BED file of high-confidence regions that are alignable between GRCh38 and T2T v1.0.1. We ran `hap.py` on the output VCF of `dipcall` VCF containing variants between T2T v1.0.1 and GRCh38 against the GIAB v4.2.1 VCF, restricting the analysis to the high-confidence regions that are alignable between GRCh38 and T2T v1.0.1 by passing that BED file in with parameter `-f`. Using BEDTools `subtract`, we removed 50 bp surrounding any FP or FN designated variants produced in that `hap.py` run from the GIAB high-confidence regions BED file in order to remove regions that were not concordant with the T2T v1.0.1. Eight hundred thirty total variants were subtracted, and 63.7% of them were located within the GRCh38 segmental duplication track.

To create the final BED file to use with `hap.py -f` in benchmarking the HG002 assemblies, we intersected the BED file containing regions concordant between the Q100 T2T v1.0.1 assembly and GIAB v4.2.1 (described in the previous paragraph) with the BED file produced from the `dipcall` output of the unpolished assembly to GRCh38. For HG005, we repeat this step but with the original HG005 high-confidence regions BED file provided by GIAB. For all `hap.py` runs, we used the options `--pass-only --no-roc --no-json --engine=vcfEval` and the GIAB v3.3 stratifications to obtain performance statistics per genomic region.

To investigate the performance of DeepPolisher in the centromere, we ran `dipcall` of our HG002 assembly (pre- and postpolishing) against Q100 HG002 v1.0.1 with the default parameters. Using `samtools -q 60`, we subset the output BAM file from `dipcall` to regions where mapping quality between the assemblies was 60. We used `BEDTools` to convert this BAM file to BED and then intersected it with the provided `cenSat` annotations BED file (<https://github.com/marbl/HG002>) to define centromere regions alignable between the Q100 assembly and our HG002 hifiasm assembly. Using `BEDTools` again to intersect the `dipcall` VCF file with these regions, we counted the number of SNVs and INDELS relative to the Q100 genome before and after polishing with DeepPolisher.

### Stratifying QV by whole genome and within GIAB high-confidence regions

For QV calculations, we ran Merqury and Yak with the default settings, using  $\sim 30\times$  coverage Illumina data for all samples. For trio samples, we used Yak trioEval along with parental Illumina data to calculate switch and hamming rates. To assess QV just within the high-confidence regions, we first took an intersection of the GIAB high-confidence regions for HG002 and HG005 using `BEDTools` to define a subset of high-confidence regions that is not genome specific. We aligned each haplotype assembly to GRCh38 with `minimap2` and parameters `-x asm5 -L --eqx --cs -c`. We then used the script [https://github.com/mobinasri/flagger/blob/main/programs/src/project\\_blocks\\_multi\\_thread.py](https://github.com/mobinasri/flagger/blob/main/programs/src/project_blocks_multi_thread.py) to project the GIAB high-confidence regions bedfile from GRCh38 to the assembly. Next, we used `BEDTools getfasta` with the default settings to extract only the sequence found within this BED file and passed it to Merqury and Yak for QV calculation. To count the number of error  $k$ -mers induced, fixed, or unchanged after polishing and to obtain the polishing edits that were the cause of inducing, fixing, or failing to change the error  $k$ -mers, we performed the same annotation procedure described under the section “DeepPolisher GQ filter optimization.” To transform QV into an estimate of error per megabase, we used the formula  $(10^{-QV/10}) \times 10^6$ .

### Comparing different methods for estimating QV

We compared several different methods for calculating QV on our HG005 hifiasm assembly and found them to be highly variable (Supplemental Table 9). Merqury (Rhie et al. 2020) and Yak (Cheng et al. 2021) typically report values two to three QV points apart (approximately a  $2\times$  difference in error rate). We found that using  $k$ -mers of size 31 led to more conservative scores than  $k$  of 21 (between two and four points lower, equivalent to an  $\sim 50\%$  difference in error rate). These longer  $k$ -mers can capture differences in longer repeats, which likely accounts for the higher predicted error rates. We also tested Merqury with a hybrid of HiFi and Illumina  $k$ -mers following the approach suggested by Mc Cartney et al. (2022). We found that the hybrid  $k$ -mer database led to QV scores around 10 points higher (a predicted order of magnitude lower error rate) than when we used just Illumina. Although using hybrid  $k$ -mers may remove some Illumina-specific bias, specifically

known issues with GC bias (Benjamini and Speed 2012; Chen et al. 2013), we observed that it introduces too many HiFi-specific sequencing errors to the “truth”  $k$ -mer space and so reduces the apparent error rate by masking true errors. We also experimented with using Element cloudbreak (Carroll et al. 2023) short-read  $k$ -mers to calculate QV and found it predicted double the error rate relative to Illumina (Supplemental Table 9). However, we found the same impact of GC bias on the error  $k$ -mers calculated by Merqury QV as was described for Illumina (Supplemental Fig. 4). Given that Element data were not available for the HPRC samples, and to be conservative, we report Merqury QV with Illumina  $k$ -mers of size 31 for this paper unless otherwise stated.

### Manual investigation of concordance between GIAB variant calls and Illumina QV

We performed a manual investigation of the concordance between the GIAB variant call errors and the Illumina QV reported errors for HG005 (polished by DeepPolisher) within the high-confidence regions using IGV (Robinson et al. 2011). To obtain BED tracks for viewing the GIAB variant call errors in respect to the raw assembly coordinates, we used `BCFtools` (Danecek et al. 2021) `isec` to obtain the intersections, unions, and complements between the raw and polished VCFs output by `hap.py` (these VCFs contain variants relative to GRCh38 labeled by true positive, false positive, or false negative). We used the tool `vcf2bed` from `BedOps` (Neph et al. 2012) to convert the VCF files to BED format, expanded the intervals by 10 bp on each side to enable projecting of the indels over to the raw assembly coordinates with the script [https://github.com/mobinasri/flagger/blob/main/programs/src/project\\_blocks\\_multi\\_thread.py](https://github.com/mobinasri/flagger/blob/main/programs/src/project_blocks_multi_thread.py).

For viewing the QV errors in IGV, we took the `*_error.bed` file (containing the locations in the assembly of the error  $k$ -mers) produced by Merqury using Illumina  $k$ -mers of size 31 and used the procedure described in the section “DeepPolisher GQ filter optimization” to annotate the polishing VCF by whether the edits induced, fixed, or failed to change the error  $k$ -mers, as well as to label the error  $k$ -mers by whether they were fixed, induced, or unchanged by polishing.

We used the files described above to randomly select 10 GIAB errors induced by DeepPolisher, 10 GIAB errors fixed by DeepPolisher, 10 polishing edits inducing error  $k$ -mers, and 10 polishing edits removing error  $k$ -mers. We carefully evaluated all 40 examples in IGV alongside the PHARAOH read alignments. For the GIAB errors, we looked at both haplotypes to understand if one or both alleles were considered an error, and in two cases, there was a polishing edit made to both haplotypes for a given GIAB error, so edits on both haplotypes were assessed. For all 42 polishing edits, we recorded the error status in both GIAB and QV. All details for the 42 edits may be found in Supplemental Table 10.

In summary, out of 42 polishing edits manually inspected, there were three cases in which GIAB indicated the edit induced an error, whereas the QV indicated that it fixed an error. For 24/42 polishing edits, the QV and GIAB metrics both agreed on whether the edit fixed or induced an error. For 15/42 polishing edits, the QV reported no error  $k$ -mer in either the pre- or postpolishing sequence, indicating that both were correct according to the QV metric. Because it is impossible for both versions of the sequence at a given location to be correct, this 33% of inspected edits likely falls in nonunique sequence, in which either the pre- or postpolishing sequence is found elsewhere in the genome, so the QV failed to flag it as false. This proportion may also be caused by noise from sequencing errors in the Illumina reads. In addition, all but two of the 42 edits examined were found within long tandem repeats or homopolymers, which are more likely to be

nonunique in the genome. We propose that this category accounts for the much lower error rate reported by the QV than by the GIAB analysis within the high-confidence regions. Because all of the QV reported errors we examined were corroborated as false by GIAB (and there were only three cases in which QV indicated an error was fixed whereas GIAB indicated it was induced), we conclude that the majority of the QV errors within the high-confidence regions are most likely true errors.

### Characterizing the context of unchanged error *k*-mers

After obtaining the error *k*-mers unchanged by polishing as described in the section “DeepPolisher GQ filter optimization,” we sought to establish how many were located within coverage drop-outs. We ran the tool `mosdepth` with parameter `--quantize 0:5:10:150` on the PHARAOH alignment to obtain locations with less than 5× coverage and then `BEDTools intersect` to obtain the error *k*-mers unchanged by polishing within those regions. To obtain the numbers that were within simple or low complexity repeats, we ran `dipcall` with the default parameters to align the HG005 raw assembly against CHM13 and then the script [https://github.com/mobinasri/flagger/blob/main/programs/src/project\\_blocks\\_multi\\_thread.py](https://github.com/mobinasri/flagger/blob/main/programs/src/project_blocks_multi_thread.py) to project the repeat annotations produced for T2T-CHM13 at <https://github.com/marbl/CHM13> to HG005 coordinates. We then used `BEDTools intersect` with the unchanged error *k*-mers and the projected repeat annotations. To calculate the GC content of all error *k*-mers, we used `BEDTools nuc`. To obtain regions in the HG005 raw assembly with homopolymers >10 bp, we used a custom Python script that may be found at GitHub ([https://github.com/miramastoras/DeepPolisher\\_manuscript/blob/main/paper\\_analysis/quantify\\_error\\_kmers\\_HG005.md](https://github.com/miramastoras/DeepPolisher_manuscript/blob/main/paper_analysis/quantify_error_kmers_HG005.md)) and `BEDTools intersect`.

### Gene impact analysis

We annotated the HG002 and HG005 assemblies before and after polishing with the Comparative Annotation Toolkit (CAT2.0) (Fiddes et al. 2018) using the `transMap` and `liftoff` modules based on the GENCODE v46 gene annotations for hg38 as the reference. For these annotation sets, we identified the locations of frameshifts by iterating over the coding sequence of every transcript and looking for gaps in the alignment. If the gap length was not a multiple of three or if the length was >30 bp, the gap was determined to be a frameshift. To identify the number of nonsense mutations that would cause early stop codons in the annotation sets, we iterated through each codon in the coding sequence and looked for an early stop codon before the canonical stop codon at the end of the transcript. We then sought to validate how many of the mutations were likely true mutations, rather than errors, by checking for their presence in an orthogonal set of variant calls (GIAB v4.2.1). The mutations were lifted over to the GRCh38 reference using `halLiftover` (Hickey et al. 2013) on a GRCh38-based Cactus alignment and then intersected with the GIAB set using `BEDTools intersect` (Quinlan and Hall 2010). To study the impact of the polishing edits on the coding regions of the genome, we intersected the total set of edits made to the assembly with the CDS regions of the assembly. The DNA and protein sequences of transcripts that were affected by these edits were extracted from the raw and polished assemblies. These were aligned against the DNA and protein sequences extracted from the GRCh38 transcripts using `parasail` (Daily 2016) to look at changes in DNA and protein sequence identity for these transcripts caused by polishing.

### Marmoset assembly polishing

To create the hifiasm assemblies for the two marmosets (Baguette and Maisie), we ran `hifiasm` (UL) v0.24.0 in HiC mode with

the default parameters. As input, we used 60× PacBio HiFi DeepConsensus reads and 30× ONT UL >100 kb reads. The ONT reads were sequenced using R10.4.1 flow cells and basecalled using Dorado v4.3.0. For phasing this assembly, we used HiC data at 50× coverage. The assemblies were polished using the HiFi and ONT data as input to the PHARAOH + DeepPolisher pipeline. To exclude regions of low HiFi coverage from the QV calculation, we ran `mosdepth` (Pedersen and Quinlan 2018) with parameters `-Q1 --quantize 0:5:10:150` on the PHARAOH BAM files for Baguette and Maisie. We extracted regions from zero to 5× coverage from the resulting BED file and used `bedtool subtract` to obtain the inverse of these regions. We then used `BEDTools getfasta` with the default settings to extract only the sequence with greater than 5× coverage of mapping quality greater than one. We ran `Merqury` with `illumina k-mers` of size 31 for the QV calculation.

### Element polishing

We obtained 50× element data for HG002 and HG005 from Carroll et al. (2023) and aligned all reads to each polished haplotype assembly separately. We used `DeepVariant` v1.6.1 with `--model_type=WGS` to call variants. We selected only homozygous alt PASS variants with genotype quality greater than seven and applied them to the assembly with `BCFtools consensus -H 2`.

### Software availability

DeepPolisher code is available publicly on GitHub (<https://github.com/google/deeppolisher>). PHARAOH is implemented in a pipeline available on GitHub (<https://github.com/miramastoras/PHARAOH>). Scripts for analysis and benchmarking are available on GitHub ([https://github.com/miramastoras/DeepPolisher\\_manuscript](https://github.com/miramastoras/DeepPolisher_manuscript)). Code used to generate the HPRC assembly data release 2 may be found at GitHub ([https://github.com/human-pangenomics/hprc\\_intermediate\\_assembly](https://github.com/human-pangenomics/hprc_intermediate_assembly)). Source code for all of these repositories is also available as [Supplemental Code](#).

### Data access

The HG002 and HG005 assemblies and DeepPolisher VCFs used for benchmarking are available at [https://s3-us-west-2.amazonaws.com/human-pangenomics/index.html?prefix=publications/DeepPolisher\\_2024/](https://s3-us-west-2.amazonaws.com/human-pangenomics/index.html?prefix=publications/DeepPolisher_2024/).

## Human Pangenome Reference Consortium Release 2

Ahmad Abou Tayoun,<sup>4,5</sup> Derek Albracht,<sup>6</sup> Jamie Allen,<sup>7</sup> Alawi A. Alsheikh-Ali,<sup>8</sup> Casey Andrews,<sup>9</sup> Dmitry Antipov,<sup>10</sup> Lucinda Antonacci-Fulton,<sup>6</sup> Mobin Asri,<sup>11</sup> Marcelo Ayllon,<sup>12</sup>

<sup>4</sup>Center for Genomic Discovery, Mohammed Bin Rashid University, Dubai Health, P.O. Box 505055, UAE

<sup>5</sup>Dubai Health Genomic Medicine Center, Dubai Health, P.O. Box 505055, UAE

<sup>6</sup>McDonnell Genome Institute, Washington University School of Medicine, St. Louis, MO 63108, USA

<sup>7</sup>European Molecular Biology Laboratory, European Bioinformatics Institute (EMBL-EBI), Wellcome Genome Campus, Hinxton, Cambridge CB10 1SD, UK

<sup>8</sup>Center for Applied and Translational Genomics (CATG), Mohammed Bin Rashid University of Medicine and Health Sciences, P.O. Box 505055, Dubai, United Arab Emirates

<sup>9</sup>Department of Genetics, Washington University School of Medicine, St. Louis, MO 63110, USA

<sup>10</sup>Genome Informatics Section, Center for Genomics and Data Science Research, National Human Genome Research Institute, National Institutes of Health, Bethesda, MD 20892, USA

<sup>11</sup>UC Santa Cruz Genomics Institute, University of California, Santa Cruz, Santa Cruz, CA 95060, USA

Jennifer R. Balacco,<sup>13</sup> Edward A. Belter Jr.,<sup>6</sup> Halle D. Bender,<sup>11</sup> Andrew P. Blair,<sup>11</sup> Silvia Buonaiuto,<sup>14</sup> Davide Bolognini,<sup>15</sup> Katherine E. Bonini,<sup>16</sup> Christina Boucher,<sup>17</sup> Guillaume Bourque,<sup>18,19,20</sup> Shuo Cao,<sup>14</sup> Andrew Carroll,<sup>21</sup> Ann M. Mc Cartney,<sup>11</sup> Monika Cechova,<sup>11</sup> Pi-Chuan Chang,<sup>21</sup> Xian Chang,<sup>11</sup> Jitender Cheema,<sup>7</sup> Haoyu Cheng,<sup>22</sup> Claudio Ciofi,<sup>23</sup> Sarah Cody,<sup>6</sup> Vincenza Colonna,<sup>14</sup> Holland C. Conwell,<sup>24</sup> Robert Cook-Deegan,<sup>25</sup> Mark Diekhans,<sup>11</sup> Maria Angela Diroma,<sup>23</sup> Daniel Doerr,<sup>26,27,28</sup> Zheng Dong,<sup>9</sup> Richard Durbin,<sup>29,30</sup> Jana Ebler,<sup>26,31</sup> Evan E Eichler,<sup>12,32</sup> Jordan M. Eizenga,<sup>11</sup> Parsa Eskandar,<sup>11</sup> Eddie Ferro,<sup>17</sup> Anna-Sophie Fiston-Lavier,<sup>33,34</sup> Sarah M. Ford,<sup>24</sup> Willard W. Ford,<sup>35</sup> Giulio Formenti,<sup>13</sup> Adam Frankish,<sup>7</sup> Mallory A. Freeberg,<sup>7</sup> Qichen Fu,<sup>9</sup> Stephanie M. Fullerton,<sup>36</sup> Robert S. Fulton,<sup>6</sup> Yan Gao,<sup>37</sup> Gage H. Garcia,<sup>12</sup> Obed A. Garcia,<sup>38</sup> Joshua M.V. Gardner,<sup>11</sup> Shilpa Garg,<sup>39</sup> Erik Garrison,<sup>14</sup> Nanibaa' A. Garrison,<sup>40,41,42</sup> John Garza,<sup>6</sup> Mohammadmehdi Ghorbani,<sup>43</sup>

Tina Graves-Lindsay,<sup>6</sup> Richard E. Green,<sup>24</sup> Cristian Groza,<sup>44</sup> Andrea Guarracino,<sup>14</sup> Melissa Gymrek,<sup>35</sup> Leanne Haggerty,<sup>7</sup> Ira M. Hall,<sup>45,46</sup> Nancy F. Hansen,<sup>10</sup> Mohammad Amiruddin Hashmi,<sup>8</sup> Maximilian Haeussler,<sup>11</sup> David Haussler,<sup>11</sup> Prajna Hebbar,<sup>11</sup> Peter Heringer,<sup>26,27,28</sup> Glenn Hickey,<sup>11</sup> Todd L. Hillaker,<sup>11</sup> S. Nakib Hossain,<sup>7</sup> Neng Huang,<sup>37,47</sup> Sarah E. Hunt,<sup>7</sup> Toby Hunt,<sup>7</sup> Nafiseh Jafarzadeh,<sup>11</sup> Nivesh Jain,<sup>13</sup> Erich D. Jarvis,<sup>13,32</sup> Juan Jiang,<sup>9</sup> Jonathan LoTempio Jr.,<sup>48</sup> Eimear E. Kenny,<sup>16</sup> Juhyun Kim,<sup>10</sup> Bonhwang Koo,<sup>13</sup> Sergey Koren,<sup>10</sup> Milinn Kremitzki,<sup>6,9</sup> Ben Langmead,<sup>49</sup> Xiaoyu Zhuo,<sup>9</sup> Heather A. Lawson,<sup>9</sup> Daofeng Li,<sup>9</sup> Heng Li,<sup>37,47</sup> Wen-Wei Liao,<sup>45,46</sup> Jiadong Lin,<sup>12</sup> Tianjie Liu,<sup>9</sup> Glennis A. Logsdon,<sup>48</sup> Ryan Lorig-Roach,<sup>11</sup> Hailey Loucks,<sup>11</sup> Jane E. Loveland,<sup>7</sup> Jianguo Lu,<sup>50</sup> Shuangjia Lu,<sup>45,46</sup> Julian K. Lucas,<sup>11</sup> Juan F. Macias-Velasco,<sup>6,9,51</sup> Maximilian G. Marin,<sup>37</sup> Franco L. Marsico,<sup>14</sup> Kateryna D. Makova,<sup>52</sup> Christopher Markovic,<sup>9</sup> Tobias Marschall,<sup>26,31</sup> Fergal J. Martin,<sup>7</sup> Mira Mastoras,<sup>11</sup> Capucine Mayoud,<sup>33</sup> Brandy McNulty,<sup>11</sup> Jack A. Medico,<sup>13</sup> Julian M. Menendez,<sup>11</sup> Karen H. Miga,<sup>11</sup> Anna Minkina,<sup>53</sup> Matthew W. Mitchell,<sup>54</sup> Saswat K. Mohanty,<sup>55</sup> Younes Mokrab,<sup>43,56,57</sup> Jean Monlong,<sup>58</sup> Shabir Moosa,<sup>43</sup> Avelina Moreno-Ochando,<sup>59,60</sup> Shinichi Morishita,<sup>61</sup> Jonathan M. Mudge,<sup>7</sup> Katherine M. Munson,<sup>12</sup> Njagi Mwaniki,<sup>62</sup> Nasna Nassir,<sup>8</sup> Chiara Natali,<sup>23</sup> Shloka Negi,<sup>11</sup> Lingbin Ni,<sup>12</sup> Adam M. Novak,<sup>11</sup> Chie Owa,<sup>61</sup> Sadye Paez,<sup>13</sup> Benedict Paten,<sup>11</sup> Hiram Clawson,<sup>11</sup> Clelia Peano,<sup>15,63</sup> Adam M. Phillippy,<sup>10</sup> Brandon D. Pickett,<sup>10</sup> Laura Pignata,<sup>14</sup> Nadia Pisanti,<sup>62</sup> David Porubsky,<sup>12</sup> Pjotr Prins,<sup>14</sup> Anandi Radhakrishnan,<sup>11</sup> Brian J. Raney,<sup>11</sup> Mikko Rautiainen,<sup>64</sup>

<sup>12</sup>Department of Genome Sciences, University of Washington School of Medicine, Seattle, WA 98195, USA

<sup>13</sup>The Vertebrate Genome Laboratory, The Rockefeller University, New York, NY 10065, USA

<sup>14</sup>Department of Genetics, Genomics and Informatics, University of Tennessee Health Science Center, Memphis, TN 38163, USA

<sup>15</sup>Human Technopole, 20157 Milan, Italy

<sup>16</sup>Institute for Genomic Health, Icahn School of Medicine at Mount Sinai, New York, NY 10029, USA

<sup>17</sup>Department of Computer and Information Science and Engineering, University of Florida, Gainesville, FL 32611, USA

<sup>18</sup>Canadian Center for Computational Genomics, McGill University, Montréal, QC H3A 0G1, Canada

<sup>19</sup>Department of Human Genetics, McGill University, Montréal, QC H3A 0G1, Canada

<sup>20</sup>Victor Phillip Dahdaleh Institute of Genomic Medicine, Montréal, QC H3A 0G1, Canada

<sup>21</sup>Google LLC, Mountain View, CA 94043, USA

<sup>22</sup>Department of Biomedical Informatics and Data Science, Yale School of Medicine, New Haven, CT 06510, USA

<sup>23</sup>University of Florence, Department of Biology, Sesto Fiorentino (FI) 50019, Italy

<sup>24</sup>Department of Ecology and Evolutionary Biology, University of California, Santa Cruz, Santa Cruz, CA 95064, USA

<sup>25</sup>Arizona State University, Consortium for Science, Policy & Outcomes, Washington, DC 20006, USA

<sup>26</sup>Center for Digital Medicine, Heinrich Heine University Düsseldorf, 40225 Düsseldorf, Germany

<sup>27</sup>Department for Endocrinology and Diabetology, Medical Faculty and University Hospital Düsseldorf, Heinrich Heine University Düsseldorf, 40225 Düsseldorf, Germany

<sup>28</sup>German Diabetes Center (DDZ), Leibniz Institute for Diabetes Research, 40225 Düsseldorf, Germany

<sup>29</sup>University of Cambridge, Cambridge Biomedical Campus, Cambridge CB2 0QQ, UK

<sup>30</sup>Wellcome Sanger Institute, Genome Campus, Hinxton, CB10 1HH, UK

<sup>31</sup>Institute for Medical Biometry and Bioinformatics, Medical Faculty and University Hospital Düsseldorf, Heinrich Heine University, 40225 Düsseldorf, Germany

<sup>32</sup>Howard Hughes Medical Institute, Chevy Chase, MD 20815, USA

<sup>33</sup>ISEM, Univ Montpellier, CNRS, IRD, 34095 Montpellier, France

<sup>34</sup>Institut Universitaire de France, 75005 Paris, France

<sup>35</sup>Department of Computer Science and Engineering, University of California San Diego, La Jolla, CA 92093, USA

<sup>36</sup>Department of Bioethics & Humanities, University of Washington School of Medicine, Seattle, WA 98195, USA

<sup>37</sup>Department of Data Science, Dana-Farber Cancer Institute, Boston, MA 02215, USA

<sup>38</sup>Department of Anthropology, University of Kansas, Lawrence, KS 66045, USA

<sup>39</sup>University of Manchester, Manchester M13 9PL, UK

<sup>40</sup>Traditional, ancestral and unceded territory of the Gabrielino/Tongva peoples, Institute for Society & Genetics, University of California, Los Angeles, Los Angeles, CA 90095, USA

<sup>41</sup>Traditional, ancestral and unceded territory of the Gabrielino/Tongva peoples, Institute for Precision Health, David Geffen School of Medicine, University of California, Los Angeles, Los Angeles, CA 90095, USA

<sup>42</sup>Traditional, ancestral and unceded territory of the Gabrielino/Tongva peoples, Division of General Internal Medicine & Health Services Research, David Geffen School of Medicine, University of California, Los Angeles, Los Angeles, CA 90095, USA

<sup>43</sup>Medical and Population Genomics Lab, Sidra Medicine, P.O. Box 26999, Doha, Qatar

<sup>44</sup>Montreal Heart Institute, Montreal, Quebec H1T 1C8, Canada

<sup>45</sup>Center for Genomic Health, Yale University School of Medicine, New Haven, CT 06510, USA

<sup>46</sup>Department of Genetics, Yale University School of Medicine, New Haven, CT 06510, USA

<sup>47</sup>Department of Biomedical Informatics, Harvard Medical School, Boston, MA 02115, USA

<sup>48</sup>Department of Genetics, Epigenetics Institute, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104, USA

<sup>49</sup>Department of Computer Science, Johns Hopkins University, Baltimore, MD 21218, USA

<sup>50</sup>Sun Yat-sen University, Guangzhou 510275, China

<sup>51</sup>Edison Family Center for Genome Sciences & Systems Biology, Washington University School of Medicine, St. Louis, MO 63110, USA

<sup>52</sup>Center for Medical Genomics, Penn State University, University Park, PA 16802, USA

<sup>53</sup>Division of Medical Genetics, Department of Medicine, University of Washington School of Medicine, Seattle, WA 98195, USA

<sup>54</sup>Coriell Institute for Medical Research, Camden, NJ 08103, USA

<sup>55</sup>Department of Biology, Penn State University, University Park, PA 16802, USA

<sup>56</sup>Department of Biomedical Science, College of Health Sciences, Qatar University, P.O. Box 2713, Doha, Qatar

<sup>57</sup>Department of Genetic Medicine, Weill Cornell Medicine-Qatar, P.O. Box 24144, Doha, Qatar

<sup>58</sup>IRSD - Digestive Health Research Institute, University of Toulouse, INSERM, INRAE, ENVT, UPS, 31024 Toulouse, France

<sup>59</sup>MATCH biosystems, S.L., 03202 Elche, Spain

<sup>60</sup>Universidad Miguel Hernández de Elche, 03202 Elche, Spain

<sup>61</sup>Department of Computational Biology and Medical Sciences, The University of Tokyo, Kashiwa, Chiba 277-8561, Japan

<sup>62</sup>University of Pisa, 56126 Pisa, Italy

<sup>63</sup>Institute of Genetics and Biomedical Research, UoS of Milan, National Research Council, 20133 Milan, Italy

<sup>64</sup>Institute for Molecular Medicine Finland, Helsinki Institute of Life Science, University of Helsinki, 00290 Helsinki, Finland

Alessandro Raveane,<sup>15</sup> Luyao Ren,<sup>12,32</sup> Arang Rhie,<sup>10</sup> Farnaz Salehi,<sup>14</sup> Samuel Sacco,<sup>24</sup> Michael C. Schatz,<sup>49,65</sup> Laura B. Scheinfeldt,<sup>54</sup> Aarushi Sehgal,<sup>35</sup> William E. Seligmann,<sup>24</sup> Mahsa Shabani,<sup>66</sup> Kishwar Shafin,<sup>21</sup> Shadi Shahatit,<sup>33</sup> Ruhollah Shemirani,<sup>16</sup> Vikram S. Shivakumar,<sup>49</sup> Swati Sinha,<sup>7</sup> Jouni Sirén,<sup>11</sup> Linnéa Smeds,<sup>55</sup> Steven J. Solar,<sup>10</sup> Marco Sollitto,<sup>13,23</sup> Nicole Soranzo,<sup>15,29,30</sup> Andrew B. Stergachis,<sup>12,53</sup> Marie-Marthe Suner,<sup>7</sup> Yoshihiko Suzuki,<sup>61</sup> Arda Söylev,<sup>26,31</sup> Jack A.S. Tierney,<sup>7</sup> Chad Tomlinson,<sup>6</sup> Francesca Floriana Tricomi,<sup>7</sup> Mohammed Uddin,<sup>8,67</sup> Matteo Tommaso Ungaro,<sup>24,68</sup> Rahul Varki,<sup>17</sup> Flavia Villani,<sup>14</sup> Mitchell R. Vollger,<sup>53</sup> Brian P. Walenz,<sup>10</sup> Charles Wang,<sup>69</sup> Lisa E. Wang,<sup>16</sup> Ting Wang,<sup>6,9,51</sup> Aaron M. Wenger,<sup>70</sup> Conor V. Whelan,<sup>13</sup> Zilan Xin,<sup>9</sup> Zheng Xu,<sup>9</sup> Kai Ye,<sup>71</sup> DongAhn Yoo,<sup>12</sup> Wenjin Zhang,<sup>9</sup> Ying Zhou,<sup>37</sup> Ivo Violich,<sup>11</sup> Giulia Zunino<sup>15</sup>

## Competing interest statement

A.C., D.E.C., P.-C.C., A.K., L.B., M.N., and K.S. are employees of Google and own Alphabet stock as part of the standard compensation package.

## Acknowledgments

B.P. and UCSC personnel were in part supported by National Institutes of Health (NIH) grants R01HG010485, U01HG010961, U24HG010262, OT2OD026682, and U24HG011853. M.M. was in part supported by NIH grant T32HG012344. The HPRC acknowledges the National Human Genome Research Institute (NHGRI) for funding the following grants supporting the creation of the human pangenome reference: U41HG010972, U01HG010971, U01HG013760, U01HG013755, U01HG013748, U01HG013744, R01HG011274, and the Human Pangenome Reference Consortium (BioProject ID: PRJNA730823).

**Author contributions:** K.S. and B.P. devised the study. M.M., B.P., and K.S. drafted the manuscript. A.C., D.E.C., P.-C.C., A.K., L.B., M.N., and K.S. developed DeepPolisher. M.A. and M.M. developed PHARAOH. M.M., M.A., and T.S.W. performed data analysis. P.H. performed gene impact analysis. M.A. and J.L. generated assemblies.

## References

- Aganezov S, Yan SM, Soto DC, Kirsche M, Zarate S, Avdeyev P, Taylor DJ, Shafin K, Shumate A, Xiao C, et al. 2022. A complete reference genome improves analysis of human genetic variation. *Science* **376**: eabl3533. doi:10.1126/science.abl3533
- Baid G, Cook DE, Shafin K, Yun T, Llinares-López F, Berthet Q, Belyaeva A, Töpfer A, Wenger AM, Rowell WJ, et al. 2022. DeepConsensus improves the accuracy of sequences with a gap-aware sequence transformer. *Nat Biotechnol* **41**: 232–238. doi:10.1038/s41587-022-01435-7
- Benjamini Y, Speed TP. 2012. Summarizing and correcting the GC content bias in high-throughput sequencing. *Nucleic Acids Res* **40**: e72. doi:10.1093/nar/gks001
- <sup>65</sup>Department of Biology, Johns Hopkins University, Baltimore, MD 21218, USA
- <sup>66</sup>University of Amsterdam, 1012 WX Amsterdam, Netherlands
- <sup>67</sup>GenomeArc Inc., Mississauga, ON L4W 5M1, Canada
- <sup>68</sup>Department of Biology and Biotechnologies “Charles Darwin”, University of Rome “La Sapienza”, 00185 Rome, Italy
- <sup>69</sup>Center for Genomics, Loma Linda University School of Medicine, Loma Linda, CA 92350, USA
- <sup>70</sup>PacBio, Menlo Park, CA 94025, USA
- <sup>71</sup>The first affiliated hospital of Xi’an Jiaotong University, Xi’an Jiaotong University, Xi’an, Shaanxi, 710049, China

- Brandes N, Ofer D, Peleg Y, Rappoport N, Linial M. 2022. ProteinBERT: a universal deep-learning model of protein sequence and function. *Bioinformatics* **38**: 2102–2110. doi:10.1093/bioinformatics/btac020
- Carroll A, Kolesnikov A, Cook DE, Brambrink L, Wiseman KN, Billings SM, Kruglyak S, Lajoie BR, Zhao J, Levy SE, et al. 2023. Accurate human genome analysis with element avidity sequencing. bioRxiv doi:10.1101/2023.08.11.553043
- Chen Y-C, Liu T, Yu C-H, Chiang T-Y, Hwang C-C. 2013. Effects of GC bias in next-generation-sequencing data on de novo genome assembly. *PLoS One* **8**: e62856. doi:10.1371/journal.pone.0062856
- Chen S, Zhou Y, Chen Y, Gu J. 2018. fastp: an ultra-fast all-in-one FASTQ preprocessor. *Bioinformatics* **34**: i884–i890. doi:10.1093/bioinformatics/bty560
- Cheng H, Concepcion GT, Feng X, Zhang H, Li H. 2021. Haplotype-resolved de novo assembly using phased assembly graphs with hifiasm. *Nat Methods* **18**: 170–175. doi:10.1038/s41592-020-01056-5
- Cheng H, Asri M, Lucas J, Koren S, Li H. 2024. Scalable telomere-to-telomere assembly for diploid and polyploid genomes with double graph. *Nat Methods* **21**: 967–970. doi:10.1038/s41592-024-02269-8
- Church DM, Schneider VA, Steinberg KM, Schatz MC, Quinlan AR, Chin C-S, Kitts PA, Aken B, Marth GT, Hoffman MM, et al. 2015. Extending reference assembly models. *Genome Biol* **16**: 13. doi:10.1186/s13059-015-0587-3
- Daily J. 2016. Parasail: SIMD C library for global, semi-global, and local pairwise sequence alignments. *BMC Bioinformatics* **17**: 81. doi:10.1186/s12859-016-0930-z
- Danecek P, Bonfield JK, Liddle J, Marshall J, Ohan V, Pollard MO, Whitwham A, Keane T, McCarthy SA, Davies RM, et al. 2021. Twelve years of SAMtools and BCFtools. *GigaScience* **10**: giab008. doi:10.1093/gigascience/giab008
- Ebert P, Audano PA, Zhu Q, Rodriguez-Martin B, Porubsky D, Bonder MJ, Sulovari A, Ebler J, Zhou W, Serra Mari R, et al. 2021. Haplotype-resolved diverse human genomes and integrated analysis of structural variation. *Science* **372**: eabf7117. doi:10.1126/science.abf7117
- Fiddes IT, Armstrong J, Diekhans M, Nachtweide S, Kronenberg ZN, Underwood JG, Gordon D, Earl D, Keane T, Eichler EE, et al. 2018. Comparative annotation Toolkit (CAT)—simultaneous clade and personal genome annotation. *Genome Res* **28**: 1029–1038. doi:10.1101/gr.233460.117
- Formenti G, Rhie A, Walenz BP, Thibaud-Nissen F, Shafin K, Koren S, Myers EW, Jarvis ED, Phillippy AM. 2022. Merfin: improved variant filtering, assembly evaluation and polishing via k-mer validation. *Nat Methods* **19**: 696–704. doi:10.1038/s41592-022-01445-y
- Hickey G, Paten B, Earl D, Zerbino D, Haussler D. 2013. HAL: a hierarchical format for storing and analyzing multiple genome alignments. *Bioinformatics* **29**: 1341–1342. doi:10.1093/bioinformatics/btt128
- Hu J, Wang Z, Liang F, Liu S-L, Ye K, Wang D-P. 2024. NextPolish2: a repeat-aware polishing tool for genomes assembled using HiFi long reads. *Genomics Proteomics Bioinformatics* **22**: qzad009. doi:10.1093/gpbjnl/qzad009
- Jain C, Rhie A, Zhang H, Chu C, Walenz BP, Koren S, Phillippy AM. 2020. Weighted minimizer sampling improves long read mapping. *Bioinformatics* **36**: i111–i118. doi:10.1093/bioinformatics/btaa435
- Jarvis ED, Formenti G, Rhie A, Guarracino A, Yang C, Wood J, Tracey A, Thibaud-Nissen F, Vollger MR, Porubsky D, et al. 2022. Semi-automated assembly of high-quality diploid human reference genomes. *Nature* **611**: 519–531. doi:10.1038/s41586-022-05325-5
- Kolmogorov JM, Billingsley KJ, Mastoras M, Meredith M, Monlong J, Lorig-Roach R, Asri M, Alvarez Jerez P, Malik L, Dewan R, et al. 2023. Scalable nanopore sequencing of human genomes provides a comprehensive view of haplotype-resolved variation and methylation. *Nat Methods* **20**: 1483–1492. doi:10.1038/s41592-023-01993-x
- Li H. 2018. Minimap2: pairwise alignment for nucleotide sequences. *Bioinformatics* **34**: 3094–3100. doi:10.1093/bioinformatics/bty191
- Li H, Durbin R. 2024. Genome assembly in the telomere-to-telomere era. *Nat Rev Genet* **25**: 658–670. doi:10.1038/s41576-024-00718-w
- Li H, Bloom JM, Farjoun Y, Fleharty M, Gauthier L, Neale B, MacArthur D. 2018. A synthetic-diploid benchmark for accurate variant calling evaluation. *Nat Methods* **15**: 595–597. doi:10.1038/s41592-018-0054-7
- Liao W-W, Asri M, Ebler J, Doerr D, Haukness M, Hickey G, Lu S, Lucas JK, Monlong J, Abel HJ, et al. 2023. A draft human pangenome reference. *Nature* **617**: 312–324. doi:10.1038/s41586-023-05896-x
- Martin M, Patterson M, Garg S, Fischer SO, Pisanti N, Klau GW, Schöenhuth A, Marschall T. 2016. WhatsHap: fast and accurate read-based phasing. bioRxiv doi:10.1101/085050
- Mc Cartney AM, Shafin K, Alonge M, Bzikadze AV, Formenti G, Functammanan A, Howe K, Jain C, Koren S, Logsdon GA, et al. 2022. Chasing perfection: validation and polishing strategies for telomere-to-telomere genome assemblies. *Nat Methods* **19**: 687–695. doi:10.1038/s41592-022-01440-3

- Miller CA, Walker JR, Jensen TL, Hooper WF, Fulton RS, Painter JS, Sekeres MA, Ley TJ, Spencer DH, Goll JB, et al. 2022. Failure to detect mutations in UZAF1 due to changes in the GRCh38 reference sequence. *J Mol Diagn* **24**: 219–223. doi:10.1016/j.jmoldx.2021.10.013
- Morisse P, Marchet C, Limasset A, Lecroq T, Lefebvre A. 2021. Scalable long read self-correction and assembly polishing with multiple sequence alignment. *Sci Rep* **11**: 761. doi:10.1038/s41598-020-80757-5
- Neph S, Kuehn MS, Reynolds AP, Haugen E, Thurman RE, Johnson AK, Rynes E, Maurano MT, Vierstra J, Thomas S, et al. 2012. BEDOPS: high-performance genomic feature operations. *Bioinformatics* **28**: 1919–1920. doi:10.1093/bioinformatics/bts277
- Pedersen BS, Quinlan AR. 2018. Mosdepth: quick coverage calculation for genomes and exomes. *Bioinformatics* **34**: 867–868. doi:10.1093/bioinformatics/btx699
- Poplin R, Chang P-C, Alexander D, Schwartz S, Colthurst T, Ku A, Newburger D, Dijamco J, Nguyen N, Afshar PT, et al. 2018. A universal SNP and small-indel variant caller using deep neural networks. *Nat Biotechnol* **36**: 983–987. doi:10.1038/nbt.4235
- Porubsky D, Dashnow H, Sasani TA, Logsdon GA, Hallast P, Noyes MD, Kronenberg ZN, Mokveld T, Koundinya N, Nolan C, et al. 2024. A familial, telomere-to-telomere reference for human *de novo* mutation and recombination from a four-generation pedigree. bioRxiv doi:10.1101/2024.08.05.606142
- Quinlan AR, Hall IM. 2010. BEDTools: a flexible suite of utilities for comparing genomic features. *Bioinformatics* **26**: 841–842. doi:10.1093/bioinformatics/btq033
- Rautiainen M, Nurk S, Walenz BP, Logsdon GA, Porubsky D, Rhie A, Eichler EE, Phillippy AM, Koren S. 2023. Telomere-to-telomere assembly of diploid chromosomes with Verkko. *Nat Biotechnol* **41**: 1474–1482. doi:10.1038/s41587-023-01662-6
- Rhie A, Walenz BP, Koren S, Phillippy AM. 2020. Merqury: reference-free quality, completeness, and phasing assessment for genome assemblies. *Genome Biol* **21**: 245. doi:10.1186/s13059-020-02134-9
- Rhie A, McCarthy SA, Fedrigo O, Damas J, Formenti G, Koren S, Uliano-Silva M, Chow W, Fungtammasan A, Kim J, et al. 2021. Towards complete and error-free genome assemblies of all vertebrate species. *Nature* **592**: 737–746. doi:10.1038/s41586-021-03451-0
- Robinson JT, Thorvaldsdóttir H, Winckler W, Guttman M, Lander ES, Getz G, Mesirov JP. 2011. Integrative genomics viewer. *Nat Biotechnol* **29**: 24–26. doi:10.1038/nbt.1754
- Schneider VA, Graves-Lindsay T, Howe K, Bouk N, Chen H-C, Kitts PA, Murphy TD, Pruitt KD, Thibaud-Nissen F, Albracht D, et al. 2017. Evaluation of GRCh38 and *de novo* haploid genome assemblies demonstrates the enduring quality of the reference assembly. *Genome Res* **27**: 849–864. doi:10.1101/gr.213611.116
- Šošić M, Šikić M. 2017. Edlib: a C/C++ library for fast, exact sequence alignment using edit distance. *Bioinformatics* **33**: 1394–1395. doi:10.1093/bioinformatics/btw753
- Stoler N, Nekrutenko A. 2021. Sequencing error profiles of Illumina sequencing instruments. *NAR Genom Bioinform* **3**: lqab019. doi:10.1093/nargab/lqab019
- Taylor DJ, Eizenga JM, Li Q, Das A, Jenike KM, Kenny EE, Miga KH, Monlong J, McCoy RC, Paten B, et al. 2024. Beyond the human genome project: the age of complete human genome sequences and Pangenome references. *Annu Rev Genomics Hum Genet* **25**: 77–104. doi:10.1146/annurev-genom-021623-081639
- Vaser R, Sović I, Nagarajan N, Šikić M. 2017. Fast and accurate *de novo* genome assembly from long uncorrected reads. *Genome Res* **27**: 737–746. doi:10.1101/gr.214270.116
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser L, Polosukhin I. 2023. Attention is all you need. *arXiv:1706.03762* [cs.CL]. doi:10.48550/arXiv.1706.03762
- Wagner J, Olson ND, Harris L, Khan Z, Farek J, Mahmoud M, Stankovic A, Kovacevic V, Yoo B, Miller N, et al. 2022a. Benchmarking challenging small variants with linked and long reads. *Cell Genomics* **2**: 100128. doi:10.1016/j.xgen.2022.100128
- Wagner J, Olson ND, Harris L, McDaniel J, Cheng H, Fungtammasan A, Hwang Y-C, Gupta R, Wenger AM, Rowell WJ, et al. 2022b. Curated variation benchmarks for challenging medically relevant autosomal genes. *Nat Biotechnol* **40**: 672–680. doi:10.1038/s41587-021-01158-1
- Walker BJ, Abeel T, Shea T, Priest M, Abouelliel A, Sakthikumar S, Cuomo CA, Zeng Q, Wortman J, Young SK, et al. 2014. Pilon: an integrated tool for comprehensive microbial variant detection and genome assembly improvement. *PLoS One* **9**: e112963. doi:10.1371/journal.pone.0112963
- Wenger AM, Peluso P, Rowell WJ, Chang P-C, Hall RJ, Concepcion GT, Ebler J, Fungtammasan A, Kolesnikov A, Olson ND, et al. 2019. Accurate circular consensus long-read sequencing improves variant detection and assembly of a human genome. *Nat Biotechnol* **37**: 1155–1162. doi:10.1038/s41587-019-0217-9
- Wick RR, Holt KE. 2022. Polypolish: short-read polishing of long-read bacterial genome assemblies. *PLoS Comput Biol* **18**: e1009802. doi:10.1371/journal.pcbi.1009802
- Yoo D, Rhie A, Hebbar P, Antonacci F, Logsdon GA, Solar SJ, Antipov D, Pickett BD, Safonova Y, Montinaro F, et al. 2025. Complete sequencing of ape genomes. *Nature* **641**: 401–418. doi:10.1038/s41586-025-08816-3
- Zhang S, Fan R, Liu Y, Chen S, Liu Q, Zeng W. 2023. Applications of transformer-based language models in bioinformatics: a survey. *Bioinform Adv* **3**: vbad001. doi:10.1093/bioadv/vbad001
- Zimin AV, Salzberg SL. 2020. The genome polishing tool POLCA makes fast and accurate corrections in genome assemblies. *PLoS Comput Biol* **16**: e1007981. doi:10.1371/journal.pcbi.1007981

Received October 22, 2024; accepted in revised form April 30, 2025.