

## Automated assembly of high-quality diploid human reference genomes

Erich D. Jarvis\*<sup>§1,2</sup>, Giulio Formenti\*<sup>§1</sup>, Arang Rhie<sup>3</sup>, Andrea Guarracino<sup>4</sup>, Chentao Yang<sup>5</sup>, Jonathan Wood<sup>6</sup>, Alan Tracey<sup>6</sup>, Francoise Thibaud-Nissen<sup>7</sup>, Mitchell R. Vollger<sup>8</sup>, Haoyu Cheng<sup>9,10</sup>, Mobin Asri<sup>11</sup>, Mark J.P. Chaisson<sup>12</sup>, Chen-Shan Chin<sup>13</sup>, Sarah Cody<sup>14</sup>, Joanna Collins<sup>6</sup>, Peter Ebert<sup>15</sup>, Merly Escalona<sup>16</sup>, Olivier Fedrigo<sup>1</sup>, Robert S. Fulton<sup>14</sup>, Lucinda L. Fulton<sup>14</sup>, Shilpa Garg<sup>17</sup>, Jay Ghurye<sup>18</sup>, Edward Green<sup>18</sup>, Ira Hall<sup>19</sup>, Patrick Hasenfeld<sup>20</sup>, Alex Hastie<sup>21</sup>, Marina Haukness<sup>11</sup>, Miten Jain<sup>11</sup>, Melanie Kirsche<sup>22</sup>, Mikhail Kolmogorov<sup>23</sup>, Jan O. Korbel<sup>20</sup>, Sergey Koren<sup>3</sup>, Jonas Korlach<sup>24</sup>, Joyce Lee<sup>21</sup>, Daofeng Li<sup>25</sup>, Tina Lindsay<sup>14</sup>, Glennis A. Logsdon<sup>8</sup>, Julian Lucas<sup>11</sup>, Feng Luo<sup>26</sup>, Tobias Marschall<sup>15</sup>, Jennifer McDaniel<sup>27</sup>, Fan Nie<sup>28</sup>, Hugh E. Olsen<sup>11</sup>, Nathan Olson<sup>27</sup>, Carnevali Paolo<sup>29</sup>, Trevor Pesout<sup>11</sup>, David Porubsky<sup>8</sup>, Daniela Puiu<sup>30</sup>, Allison Regier<sup>31</sup>, Jue Ruan<sup>32</sup>, Steven L. Salzberg<sup>30</sup>, Ashley D. Sanders<sup>33</sup>, Michael C. Schatz<sup>22</sup>, Anthony Schmitt<sup>34</sup>, Valerie A. Schneider<sup>7</sup>, Siddarth Selvaraj<sup>34</sup>, Kishwar Shafin<sup>11</sup>, Alaina Shumate<sup>30</sup>, Catherine Stober<sup>20</sup>, James Torrance<sup>6</sup>, Justin Wagner<sup>27</sup>, Jianxin Wang<sup>28</sup>, Aaron Wenger<sup>24</sup>, Chuanle Xiao<sup>36</sup>, Aleksey V. Zimin<sup>30</sup>, Guojie Zhang<sup>36</sup>, Ting Wang<sup>8,13,36</sup>, Heng Li<sup>9</sup>, Erik Garrison<sup>37</sup>, David Haussler<sup>38</sup>, Justin M. Zook<sup>27</sup>, Evan E. Eichler<sup>8</sup>, Adam M. Phillippy<sup>3</sup>, Benedict Paten<sup>11</sup>, Kerstin Howe<sup>§6</sup>, Karen H. Miga<sup>§11</sup>, Human Pangenome Reference Consortium

\*co-first authors

§co-corresponding authors

### Corresponding Author Contact Information:

### Affiliations for Primary Authors:

1. Vertebrate Genome Laboratory, The Rockefeller University, New York, NY, USA, 10065
2. Investigator, Howard Hughes Medical Institute, Chevy Chase, MD, 20815
3. National Human Genome Research Institute, National Institutes of Health, 9000 Rockville Pike, Bethesda, MD 20892
4. Genomics Research Centre, Human Technopole, Viale Rita Levi-Montalcini, 1, 20157 Milano MI, Italy
5. BGI-Shenzhen, Shenzhen 518083, China
6. Tree of Life, Wellcome Sanger Institute, Cambridge, CB10 1SA
7. NCBI, National Library of Medicine, National Institutes of Health, Bethesda, MD 20894
8. Department of Genome Sciences, University of Washington School of Medicine, Seattle, WA 98195
9. Department of Data Sciences, Dana-Farber Cancer Institute, Boston, MA
10. Department of Biomedical Informatics, Harvard Medical School, Boston, MA
11. UC Santa Cruz Genomics Institute, University of California, Santa Cruz, CA
12. Quantitative and Computational Biology, University of Southern California, 1050 Childs Way RRI 408H, Los Angeles, CA 90089
13. Foundation for Biological Data Science, 3410 Plateau Dr., Belmont, CA
14. McDonnell Genome Institute, Washington University School of Medicine, St. Louis, MO
15. Institute for Medical Biometry and Bioinformatics, Heinrich Heine University, Düsseldorf, Germany
16. Department of Biomolecular Engineering, University of California Santa Cruz, Santa Cruz, CA, 95064
17. Department of Biology, University of Copenhagen, Copenhagen Denmark
18. Dovetail Genomics, Scotts Valley, CA 95066
19. Yale School of Medicine, New Haven, CT
20. Genome Biology, European Molecular Biology Laboratory, Meyerhofstr. 1 69117 Heidelberg Germany
21. Bionano Genomics, San Diego, CA

22. Department of Computer Science, Johns Hopkins University, Baltimore, MD
23. Department of Computer Science and Engineering, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093
24. Pacific Biosciences, Menlo Park, CA, USA
25. Department of Genetics, Washington University School of Medicine, St. Louis, MO
26. School of Computing, Clemson University, Clemson, SC 29634-0974, USA.
27. Material Measurement Laboratory, National Institute of Standards and Technology, 100 Bureau Dr., Gaithersburg, MD, 20899, USA
28. School of Information Science and Engineering, Central South University, Changsha 410083, People's Republic of China
29. Chan Zuckerberg Initiative, Redwood City, CA
30. Department of Biomedical Engineering, Johns Hopkins University, 3400 N. Charles St., Baltimore, MD 21218
31. DNAnexus, Mountain View, CA, USA
32. Agricultural Genomics Institute, Chinese Academy of Agricultural Sciences, Shenzhen, China.
33. Berlin Institute for Medical Systems Biology, Max Delbrück Center for Molecular Medicine in the Helmholtz Association (MDC), Hannoversche Str. 28 10115 Berlin, Germany
34. Arima Genomics, San Diego, CA, USA.
35. State Key Laboratory of Ophthalmology, Zhongshan Ophthalmic Center, Sun Yat-sen University, #7 Jinsui Road, Tianhe District, Guangzhou, P.R. China
36. Edison Family Center for Genome Sciences and Systems Biology, Washington University School of Medicine, St. Louis, MO
37. Department of Genetics, Genomics and Informatics, University of Tennessee Health Science Center, Memphis, TN
38. Department of Ecology and Evolutionary Biology, University of California Santa Cruz, Santa Cruz, CA, USA

## Abstract

The current human reference genome, GRCh38, represents over 20 years of effort to generate a high-quality assembly, which has greatly benefited society<sup>1,2</sup>. However, it still has many gaps and errors, and does not represent a biological human genome since it is a blend of multiple individuals<sup>3,4</sup>. Recently, a high-quality telomere-to-telomere reference genome, CHM13, was generated with the latest long-read technologies, but it was derived from a hydatidiform mole cell line containing a duplicated female haplotype, and is thus completely homozygous<sup>5</sup>. To address these limitations, the Human Pangenome Reference Consortium (HPRC) recently formed with the goal of creating a collection of high-quality, cost-effective, diploid genome assemblies for a pangenome reference that represents human genetic diversity<sup>6</sup>. Here, in our first scientific report, we determined which combination of current genome sequencing and automated assembly approaches yields the most complete, accurate, and cost-effective diploid genome assemblies with minimal manual curation. Approaches that used high-accuracy long reads and parent-child data to sort haplotypes during assembly outperformed those that did not. Developing a combination of all the top methods, we generated our first high-quality diploid reference assembly, containing only ~4 gaps (range 0-12) per chromosome, most within  $\pm 1\%$  of CHM13's length. Nearly 1/4th of protein coding genes have synonymous amino acid changes between haplotypes, and centromeric regions had the highest density of variation. Our findings serve as a starting point for assembling near-complete human genomes at the scale required for constructing a human pangenome reference that captures all genetic variation from single nucleotides to large structural rearrangements.

## Keywords

Long reads, linked reads, optical maps, Strand-seq, genome assembly, human diversity, haplotype phasing

## Main

The current human reference genome assembly, GRCh38, reflects several decades of effort, and was the primary output of the Human Genome Project (HGP). It is a clone-based assembly created from the physical mapping of thousands of individually sequenced 40-2,000 kb BAC, YAC, and fosmid clones, supplemented with whole genome sequence data<sup>1,2</sup>. This original reference contains a combination of DNA sequences from 20 anonymous volunteers, with one individual representing ~70% of the assembly<sup>2</sup>. During the course of the HGP, the assembly went from having over 150,000 gaps to currently 349 gaps<sup>2,3</sup>, with the total cost exceeding \$2.7 billion. Despite being one of the most complete references available, GRCh38 represents an incomplete composite genome, as opposed to any one human genome and does not adequately capture the full spectrum of humankind genomic variation<sup>7</sup>.

In the years following the HGP, several technological limitations prevented the generation of new human reference genomes of similar or higher quality at scale. Large duplications, particularly those larger than the sequence read lengths, are particularly challenging to assemble. The sequencing enzymes used have difficulty reading through GC-rich regions with complex structures, such as promoters that regulate gene expression<sup>8,9</sup>. We also now know that mixing diverse haplotypes into a single assembly, even from the same individual, can introduce many errors with standard assembly tools<sup>8,10,11</sup>: including switch errors where variants from each haplotype are assembled into the same pseudo-haplotype; haplotype false duplications and associated gaps where the two haplotypes are sufficiently divergent such that they assemble separately as false paralogs; and consensus errors due to collapses between haplotypes. While resequencing efforts using less expensive short reads have contributed to revealing more single-nucleotide variation, structural variation is not fully captured<sup>9,12</sup>.

Major improvements have since been made in sequence read lengths<sup>4,13</sup>, sequence accuracy of long reads<sup>14</sup>, contig algorithms, scaffolding of contigs into chromosomal maps<sup>8,15-17</sup>, haplotype-resolving algorithms<sup>18-20</sup>, and technologies that allow sequencing at reduced cost. These include advances made for vertebrates by the Vertebrate Genomes Project (VGP)<sup>8</sup>, the Human Genome Structural Variation Consortium (HGSVC)<sup>12</sup>, and the Telomere-to-Telomere (T2T) consortium, which produced for the first time a complete and error-free human reference genome of CHM13<sup>5</sup>. The CHM13 cell line is from a hydatidiform mole, i.e. derived from an sperm cell containing an X chromosome with nearly two identical paternal haploid complements, which eliminated the need to separate haplotypes and thus the associated errors with assembling a diploid individual. The T2T CHM13 assembly required a significant amount of manual curation as. Thus, despite improvements in technology, errors still remain in the automated sequencing and assembly pipelines, particularly for diploid genomes which is critical to advancing technology to clinically relevant samples and to understanding human genetic variation.

To help overcome these limitations, in 2019 the National Human Genome Research Institute (NHGRI) invested in an international Human Pangenome Reference Consortium (HPRC), which aims to produce a high-quality, pangenome human reference, with an aspired goal of representing over 99% of human genetic diversity for minor alleles at least 1% or higher frequency in the human population<sup>6</sup>. At current estimates, one could start to approach this value with complete *de novo* assemblies of ~450 individuals (e.g. 900 haplotypes) from the world population (**Supplementary Note 1**). Starting in 2020, we tested the current best practices in sequencing technologies and automated assembly algorithms on one human sample, HG002, an openly-consented Ashkenazi individual from the Personal Genome Project<sup>21</sup>. We included parent samples (HG003-father, HG004-mother) for trio-based assemblies, where parental genomic sequence data are used to sort the parental haplotypes in the offspring sequence data<sup>19,8</sup>. Extensive evaluation of the resulting assemblies alongside GRCh38 (and CMH13 when it later became available) led to new approaches that yielded the best of over 60 metrics, and new biological discoveries. Evaluation also

identified areas of needed improvement to achieve automated complete and error-free diploid genome assemblies.

### Data types and algorithms

We chose HG002 because of the availability of prior extensive public data<sup>22</sup> and benchmarks<sup>23</sup>, generated by the Genome in a Bottle (GIAB) consortium. As a male sample, it enables the assembly and evaluation of both X and Y chromosomes. We obtained or generated additional state-of-the-art sequence data types, including long reads (>10 kb) for generating contigs and long-range linking information (e.g. 10X linked reads, Hi-C linked reads, optical maps, and Strand-seq) for scaffolding the contigs into chromosomes (**Supplementary Table 1**). These choices were made based on lessons learned for producing high-quality assemblies from other consortia (e.g. VGP<sup>8</sup>, T2T<sup>5</sup>, HGSVC<sup>12</sup>) or individual labs<sup>24–26</sup>. In particular, long-read based assemblies are more contiguous and more structurally accurate than short-read based assemblies, long-range linking information can place contigs into chromosome-level scaffolds, and haplotype phasing and high base accuracy helps prevent false duplications and other common assembly errors.

We made an open call to the international genome community to produce the most complete and highest-quality genome assembly possible of HG002 with the data provided, in an automated manner without manual curation (<https://humanpangenome.org/hg002/>). We received 23 assembly combinations (named asm1 to asm23, with suffix a/b for haplotypes), from 14 groups (including HPRC members) that used different data types and algorithms for contig assembly, scaffolding, and/or haplotype phasing when attempted (**Table 1**). Among these 23, 12 assembly algorithms were used: Canu and HiCanu<sup>14</sup>, CrossStitch, DipAsm<sup>26</sup>, FALCON Unzip<sup>18</sup>, Flye<sup>27</sup>, hifiasm<sup>28</sup>, MaSuRCA<sup>29</sup>, NECAT<sup>30</sup>, Peregrine<sup>31</sup>, Shasta<sup>25</sup>, and wtdbg<sup>32</sup> (**Table 1**). We classified the assemblies into four categories: 1) diploid scaffolded assemblies, where there was an attempt to assemble comparable contigs and scaffolds of both haplotypes or two pseudohaplotypes (mixed paternal and maternal derived sequences); 2) diploid contig-only assemblies, where there was an attempt to assemble comparable contigs only of both haplotypes/pseudohaplotypes or to assemble a more complete assembly representing one pseudohaplotype; 3) haploid scaffolded assemblies, where contigs and scaffolds were all merged into one pseudohaplotype assembly; and 4) haploid contig-only assemblies, where only contigs were generated and all merged into one pseudohaplotype assembly (**Table 1**). CrossStitch is reference-based (GRCh38) by design, and MaSuRCA used the GRCh38 reference for scaffolding, to order and orient assembled HG002 contigs into chromosome-level scaffolds, followed by filling in the gaps with GRCh38 sequence. Since these assemblies (asm1, asm15, and asm17) are not “pure” *de novo* assemblies, they might contain some bias towards the reference and certainly reference sequence for gap filling. They are included for metric comparisons, but considered cautiously for *de novo* methods going forward. Following the VGP model<sup>8</sup>, we assessed ~54 metrics under 14 categories, including continuity, structural accuracy, base accuracy, haplotype phasing, and completeness (**Supplemental Table 2**).



ID	Team	Pipeline	Technologies	Contigs	Scaffolders
<b>Diploid scaffold assemblies</b>					
asm23a,b	Rockefeller	Trio VGP	CLR, 10X , BN, Hi-C	Trio Canu	Trio based: Scaff10x, Solve, Salsa
asm10a,b	UCPH	Dipasm	HiFi, HiC	Peregrine	DipAsm, 3D-DNA, HapCUT2, Whatshap
asm2a,b	Dovetail	Dipasm HiRise	HiFi, HiC	Peregrine	HiRise, HapCUT2
asm22a,b	Dovetail	Dipasm Salsa	HiFi, HiC	Peregrine	Salsa, HapCUT2
asm14a,b	HHU/UW	PGAS	HiFi, Strand-seq	Peregrine	SaaRclust
asm17a,b	JHU	CrossStich	HiFi, ONT UL, HiC	Crossstitch	Ref-based to GRCh38, HapCUT2
<b>Diploid contig assemblies</b>					
asm6a,b	NHGRI	Trio Flye ONT std	ONT	Trio Flye	NA
asm7a,b	NHGRI	Trio Flye ONT UL	ONT UL >100Kb	Trio Flye	NA
asm19a,b	NHGRI	Trio HiCanu	HiFi	Trio HiCanu	NA
asm20a,b	NHGRI	Trio HiPeregrine	HiFi	Trio Peregrine	NA
asm9a,b	DFCI Harvard	Trio HiFiasm	HiFi	Trio Hifiasm	NA
asm11a,b	UCPH	Dipasm HiRise	HiFi, HiC	DipAsm, HiRise?	NA
asm3a,b	FBDS	Peregrine HiFi 25kb	HiFi long	Peregrine	NA
asm4a,b	FBDS	Peregrine HiFi 20kb	HiFi	Peregrine	NA
asm16a,b	PacBio	FALCON Unzip	HiFi	FALCON unzip	NA
asm8a,b	NHGRI	HiCanu	HiFi	HiCanu, Purge_dups	NA
<b>Haploid scaffold assemblies</b>					
asm5	UCSD	Flye ONT	ONT, HiFi	Flye	Flye
asm18	UCSC-CZI	Shasta ONT HiRise	ONT UL, Hi-C	Shasta	HiRise
asm21	UCSC-CZI	Shasta ONT Salsa	ONT UL, Hi-C	Shasta	Salsa2
asm15	JHU	MaSuRCA Flye ONT	ONT UL >120kb, HiFi	Flye	Ref-based to GRCh38, MaSuRCA
asm1	JHU	MaSuRCA Combo	Old ONT, III, HiFi	MaSuRCA	Ref-based to GRCh38, MaSuRCA
<b>Haploid contig assemblies</b>					
asm3a	FBDS	Peregrine HiFi 25K	HiFi long	Peregrine	NA
asm4a	FBDS	Peregrine HiFi	HiFi	Peregrine	NA
asm13	CAAS-AGIS	wtdbg2 HiFi	HiFi, III	wtdbg2	NA
asm12	Clemson	NECAT ONT	ONT (no UL)	NECAT	NA
<b>Final diploid</b>					
HPRC	HPRC	Trio HPRC v1.0	HiFi, ONT UL, BN, Hi-C	Trio Hifiasm	Trio based: Solve, Salsa, gap fill, curated

**Table 1. Summary of sequencing and assembly approaches tested.** Listed are the 23 assemblies generated, categorized into four broad types based on whether there were diploid (a,b) or merged haploid, and scaffolded or contigs only. They are color coded according to type of haplotype phasing: blue, trio-based; light blue, non-trio more complete for both haplotypes; turquoise, non-trio, more complete for one haplotype; green, a merging of haplotypes. The final HPRC reference v1.0 is colored magenta to distinguish it from the test assemblies, but is trio-based. Grey-text entries are also referenced-based assemblies. The sequencing technologies are the primary data types used (**Supplementary Table 1**), and the contig and scaffold assemblers are the algorithms used, although the pipelines involve more steps (see methods and **Supplementary Table 2a,b**).

### Non-human contamination and organelle genomes

First, we screened for non-human DNA and found all *de novo* assemblies had between 1 to 25 contigs or scaffolds with library adaptor sequence contamination, which were not successfully removed during read processing (**Extended Data Fig. 1a, Supplemental Table 2c**). The presence of adaptor sequences introduced gaps between some of the human-based contigs and were retained due to them being within a read, or two or more were concatenated together to make contigs of just adaptors (**Supplementary Note 2**). We also found instances of assembled bacterial (*Escherichia coli*) and yeast (*Saccharomyces cerevisiae*) genomes, either as stand-alone microbial contigs or scaffolds (3 assemblies), chimeric scaffolds with human genomic DNA (4 assemblies), or both (4 assemblies; **Extended Data Fig. 1b, Supplemental Table 2c**). There were typically less than 10 copies of these microbial contigs per assembly, except in the wtdbg2 assembly, with 35 *E. coli* and 46 *S. cerevisiae* contigs. There were also from 1 to ~40 assembled human mitochondrial (MT) contigs in ~74% (17/23) of the assemblies (**Extended Data Fig. 1c, Supplemental Table 2c**). In trio-based assemblies, they were all associated with the maternal haplotype, indicating that the MT reads were correctly sorted during haplotype phasing before generating contigs (in the VGP Trio


assembly, the MT was purposely included in both haplotypes to avoid NUMT overpolishing<sup>8</sup>). This finding further demonstrates<sup>33</sup> that with long reads, most new assembly algorithms can assemble a mitochondrial genome in one contig. Part of the reason for the differential presence of MT genomes in the assemblies is presumably due to differential read length thresholds used for initial contig assembly; the higher the size threshold, the less likely MT reads will be included<sup>33</sup>.

### Highly contiguous phased assemblies

Our assembly targets were an expected maternal genome size of ~3.06 Gb (22 autosomes + X) and an expected paternal size of ~2.96 Gb (22 autosomes + Y), as the X (155.3 Mb) is ~96 Mb bigger than the Y (~60 Mb)<sup>34</sup>. The goal was to minimize the number of gaps for a theoretical maximum gapless contig NG50 that equals a chromosome NG50 of ~150 Mb for human (the contig where at least half of the assembled contigs are this size and bigger)<sup>5</sup>. Almost all assemblies, including the diploid ones, were close to the expected sizes of a human genome (~3.0 Gb; range 2.8-3.1 Gb; **Extended Data Fig. 2a-c, Supplemental Table 2d-f**). Only the diploid pair asm19a and asm19b were bigger by ~3%. For the trio-based assemblies, the maternal (mat) haplotypes were all longer than the paternal (pat), consistent with sex chromosome differences. The non-trio assemblies tended to be skewed to the expected size of a maternal assembly (with X chromosome), in part because some had both X and Y in both haplotypes, or were missing a Y assembly altogether (**Supplemental Table 2d**, assessed for the diploid scaffolded assemblies). The two assemblies that came closest to the theoretical expected total size (98-100%) for both maternal and paternal haplotypes were the Trio VGP scaffolded assemblies (asm23a,b) and the Trio hifiasm assemblies (asm9a,b; **Extended Data Fig. 2a**). Although the assembly sizes were close to expected, the scaffolded assemblies still had quite a range of missing sequence (Ns), ~40 kb to 50 Mb, either in the unassembled gaps between contigs or trailing at the ends of contigs (**Extended Data Fig. 2c; Supplemental Table 2f**). In comparison, the GRCh38 reference has ~151 Mb of gaps. We note that with the exception of Bionano optical maps, most scaffolding tools place arbitrary gap sizes. Most assemblies also had from 0.3-2.3% false duplications (**Extended Data Fig. 2d, Supplemental Table 2g**), which could also explain why some were bigger than expected<sup>8</sup>. The GRCh38 reference is also known to still contain false duplications<sup>5,35</sup>, although the precise estimate might be difficult to measure due to the mixture of multiple haplotypes that went into the assembly.

Half of the assemblies had contig NG50 sizes in the 20-40 Mb range, including for both haplotypes in some of the diploid assemblies (**Fig. 1a; Supplemental Table 2e**). This is between ~13-27% of the theoretical maximum of 150 Mb for the maternal haplotype, indicating partial chromosomal length contigs. The exceptions well below 10 Mb were the alternative (alt) haplotypes of approaches like FALCON Unzip or HiCanu that generate a partial diploid assembly by design (asm16b and asm8b, respectively), with the primary pseudohaplotype being more contiguous (asm16a and asm8a), and both haplotypes of the Dovetail implementation of the DipAsm assembler (asm2 and asm22), where Hi-C data was used to phase the haplotypes. In contrast, the original implementation of DipAsm with Hi-C created two assemblies with contig NG50 greater than 20 Mb (a,b). Not surprisingly, the assembly (asm7a,b) that used the ONT ultralong reads (ONT UL > 100 kb) had the highest contig NG50s (48.6 Mb maternal; 39.8 Mb paternal). The trio-based ONT assemblies and the trio-based hifiasm (asm9a,b) HiFi assemblies had the fewest contigs (~600-900) compared to all the other diploid assemblies (**Extended Data Fig. 3a**). All scaffolded assemblies had scaffold NG50 values ranging from 80-140 Mb (**Fig. 1b**), 53-93% of the theoretical maximum of 150 Mb. All non-trio diploid scaffolded assemblies had 23-30 scaffolds in the final assembly, at or close to the expected 23 chromosomes per haplotype (**Supplemental Table 2f**). However, this particular metric comparison is made less informative since DipAsm inherently filters out scaffolds less than 10 kb, PGAS excludes contigs less than 500 kb as the Strand-seq signal is too sparse to scaffold small contigs, and CrossStitch only includes contigs/scaffolds that align to the GRCh38 reference. Such exclusions have the potential to result in higher numbers of gaps with missing sequence. This idea is supported by the diploid Trio VGP assembly (asm23a,b) that used multiple scaffolding data types, without exclusions, and had the highest number of scaffolds (over 2000 each), but the fewest gaps (673 maternal and 917 paternal,

relative to 900-4000; **Extended Data Fig. 3b,c**). The size of the largest scaffold (max) for most assemblies approached the size of chromosome 1 (248 Mb; range 149-242 Mb; **Supplementary Table 2f**), indicating nearly whole chromosome assemblies, albeit with gaps. These findings demonstrate an important shift in recent assembly tools to generate two separate assemblies per individual, representing the two haplotypes or pseudohaplotypes.



**Fig. 1 | Assembly continuity, phasing, and base call accuracy metrics. a, Contig NG50 values. b, Scaffold NG50**

values. **c**, Haplotype phase block NG50 values. **d**, QV base call accuracy; as an example, QV60 is about 1 error per Mb. Color code designated broad type of haplotype phasing performed. Grey shaded regions in (b) are not applicable for scaffolds, as these are contig-only assemblies; however, several contig only assemblies had gaps, because the Flye assembler inserts gaps in contigs where there is uncertainty of a repeat sequence; and the `purge_dups` function applied to the HiCanu contigs removes false duplications within contigs and creates a gap in the removal location. Grey shading in (c) indicates not applicable for phase blocks, because the GRCh38 has many haplotypes and CHM13 is from a haploid cell line. Numbers in brackets along the x-axis are the assembly numbers given.

Despite the high levels of contiguity among the assemblies, manual curation using gEVAL alignments<sup>36</sup>, Hi-C interaction plots (**Extended Data Fig. 4a**), and Bionano maps revealed a handful to several hundred contigs with missed joins that should have connected them, misjoins between unrelated sequences, other scaffolding errors (such as inversions, translocations, and false duplications), and contig errors (such as chimeric and collapsed sequences; **Extended Data Fig. 5a-d, Supplementary Table 2h**). There was no one approach, without using a highly curated reference (i.e. CrossStitch or MaSuRCA; asm1, 15, and 17), that was free of one or more scaffolding or contig errors in an automated process. As an independent measure of structural accuracy, we used Strand-seq data<sup>37,38</sup>, a method for selectively sequencing of the plus (Crick) and minus (Watson) strand of genomic DNA from cultured cells, to pinpoint large scale assembly errors such as misorientations (inversions or reverse complements) and chimeric or collapsed contigs or scaffolds (**Supplementary Fig. 1**). Nearly all assemblies had 1-25 misorientation errors, with the exception of the Strand-seq based asm14. Most had only 0 to 10 chimeric errors, except the Trio VGP maternal, Dipasm asm10, Peregrine assemblies asm20, and Flye ONT asm5 as outliers (**Supplementary Fig. 1**). We think this is due to CLR data used for the VGP Trio assembly, and that Peregrine and Flye assemblers are more greedy in connecting contigs. Despite these errors, each approach avoided at least one type of error that others did not.

### Base call accuracy

The base accuracy of the underlying assemblies is critical for subsequent annotation of protein-coding genes and non-coding regulatory DNA as well as the characterization of genetic variation especially with more complex regions that have traditionally been collapsed or excluded as a result of their repetitive content. We compared *k-mer* frequencies between unassembled Illumina sequencing reads and each assembly to estimate consensus accuracy. PGAS Strand-seq (asm14a,b) generated the highest base call accuracy (QV) among scaffolded diploid assemblies while Trio hifiasm (asm9a,b) and HiCanu (asm8a) generated the highest among the contig-only diploid assemblies (QV  $\geq$  50, or no more than 1 base call error per 100,000 bp; **Fig. 1d**). Among the merged haploid assemblies, Flye ONT.phap (asm5) performed the best with 2 rounds of base call polishing each with ONT and HiFi reads. Obtaining such a high degree of accuracy (QV > 50) in long reads has only been a recent advance, specifically with HiFi<sup>14</sup>. What these four assemblies share in common is the use of HiFi reads, either for high level read/contig filtering (asm14a,b, asm8a), polishing (asm5), and/or phasing of haplotypes (asm9a,b).

### Variant benchmarking

One of the primary goals of the HPRC is to detect and catalog genetic variation among haplotypes in the population. To determine how well each assembly detects haplotype variation, we developed a benchmarking pipeline that calculates assembly quality metrics based on variant calling. We first aligned each assembly to GRCh38, then used `dipcall`<sup>39</sup> to call variants and compared them to a manually validated ground truth, the v4.2.1 small variant HG002 benchmark from GIAB<sup>23</sup>, following the Global Alliance for Genomics and Health (GA4GH) benchmarking best practices<sup>40</sup>. We found that all of the diploid-based assemblies had high true positive (TP) rates (aka recall or sensitivity) above 90% for single nucleotide variants (SNVs), whereas the haploid assemblies were all around 40%, due to merging of haplotypes (**Extended Data Fig. 6a**). The Trio hifiasm graph-based haplotype phasing assembly (asm9) once again had the highest TP rate (99.47%; **Supplementary Table 3a**). As expected, the values were higher (65-74%) for the haploid assemblies when ignoring genotype (GT; when only one haplotype has to match the

benchmark), but still not as high as the diploid assemblies (**Supplementary Table 3a**). When examining variants in the segmental duplications alone, the diploid assemblers performed poorly (TP dropping by 9-32%), with the exception of Trio hifiasm (94.10%) and Trio HiCanu (91.38%) assemblies (which only dropped by 5 and 6% respectively, **Supplementary Table 3a**).

We next assessed the accuracy of small indels (insertions/deletions < 50 bp), which are particularly problematic and highly variable due to their association with short tandem repeats (STRs). All the HiFi-based diploid assemblies (~92-98%) outperformed the haploid assemblies (~38-59%), as well as the ONT diploid assemblies (~52-58%; **Supplementary Table 3b**), the latter due to the high indel error rate in ONT reads. Trio hifiasm (asm9a,b) had the highest combination of TP rates for both SNVs and small indels. The Trio hifiasm assembly was used to further improve the GIAB benchmark for SNVs, small indels, and larger SVs (indels, inversions, and translocations) in 273 challenging, medically-relevant genes that were not well-represented in the GIAB v4.2.1 benchmark or GIAB v0.6 SV benchmark. Extensive curation by GIAB found that the Trio hifiasm assembly produced more accurate variant calls across SNVs, small indels, and SVs in these challenging regions, and the primary error mode was inaccurate genotypes in highly homozygous regions, particularly for indels in long homopolymers<sup>41</sup>. These results demonstrate that diploid assemblies are not only highly concordant but exceed existing variant benchmarks in regions resolved by mapping-based methods. They, thus, show the greatest promise for resolving more challenging regions and variants not included in current benchmarks.

### Annotation

We performed annotation for each assembly by aligning to them the human NCBI RefSeq transcriptome dataset of 78,492 transcripts from 27,225 autosomal genes, and measured mapping statistics, using GRCh38 and CHM13 as controls. Surprisingly, most of the assemblies had 100-400 genes (over 1,600 for the haploid wtdbg2 asm13 assembly) where transcripts did not align at all. Exceptions were the Trio VGP assembly (asm23a,b), the Trio Canu (asm19a,b), Trio hifiasm (asm9a,b), and reference-based assemblies (asm1, 15, 17) with only ~60-70 unaligned genes for each haplotype, twice the missing number of 36 for GRCh38 but similar to 66 genes for CHM13 (**Extended Data Fig. 9a; Supplementary Table 2m**). Most of the contig-only assemblies had more genes (~100-500) split between contigs than those split in the scaffolded assemblies (**Extended Data Fig. 9a; Supplementary Table 1m**), consistent with scaffolding also bringing separate parts of more genes together. The Trio VGP (asm23a,b) scaffolded assembly and Trio hifiasm (asm9a,b) contig-only assembly had the fewest split genes (~30-40) among the de novo assemblies, the reference-based assemblies had even fewer (1-9), and less than the GRCh38 reference (10 genes). Most assemblies had 100-700 genes (over 4,000 in the alts of asm16b and asm8b) that were less than 95% complete, except for the Trio VGP, Trio hifiasm, and reference-based assemblies (range 32-89 genes; **Extended Data Fig. 9b**). For almost all assemblies, there were 200-600 genes apparently collapsed as assessed by overlapping transcript mapping (**Extended Data Fig. 9c**); those that used HiFi were on the lower end of this distribution. Similarly, the number of genes that required frameshift error corrections were ~1,000 for assemblies that used CLR reads (Trio VGP, asm 23a,b), ~1,500 that used the 25kb longer but less accurate HiFi reads (asm3, asm4), incredibly ~6,000-16,000 (more than half of the genes) that used unpolished ONT reads, but only ~100-200 genes with the shorter (15kb) but more accurate HiFi reads (**Extended Data Fig. 9a; Supplementary Table 1m**). These findings demonstrate that a critical combination of read length, base call accuracy, structural accuracy, and haplotype phasing are necessary to obtain the most complete and accurate annotation possible, and those not annotated in the highest quality assemblies appear to represent real biological differences between individuals.

### Trio approaches yielded higher phasing accuracy

The original Trio assembly approach of binning long reads into their respective maternal and paternal haplotypes before generating contigs was originally implemented with the Canu assembler, as TrioCanu<sup>19</sup>, but this approach had not yet been tested in a head-to-head comparison with different assemblers and data

types. Here we tested haplotype-binned reads with different contig assembly algorithms (Flye, HiCanu, HiFiasm, and Peregrine), different long read data types (HiFi, CLR, and ONT), and with trio-sorted scaffolding data types (linked reads, optical maps, and Hi-C). We found that all of the trio-based approaches yielded more extended phasing of the same haplotype than most non-trio counterparts. In particular, trios that used HiFi or CLR data had the largest NG50 haplotype phase blocks ( $\sim 10\text{-}30$  Mb vs  $< \sim 0.2\text{-}5.0$  Mb; **Fig. 1c**), lowest haplotype switch errors within contigs or scaffolds ( $\sim 0.01\text{-}0.02\%$  vs  $0.20\text{-}7.3\%$ ), the highest bp total that were phased, and the most complete separation ( $\sim 99\%$ ) of paternal and maternal haplotype k-mers when using HiFi reads (**Extended Data Fig. 7a-c**, **Supplementary Table 2i-k**). Several of the trios approaches (Trio HiCanu and Trio hifiasm) had the least collapsed sequence in the repetitive hard to assemble regions (**Extended Data Fig. 8a-c**; **Supplementary Table 2l**). The only non-trio method that approached the phasing accuracy for shared maternal and paternal alleles of a trio method used Strand-seq for phasing and scaffolding (asm14; **Fig. 1c**; **Extended Data Fig. 7a**), but it suffered from having the highest within scaffold errors (**Extended Data Fig. 5c**) and twice the number of collapsed bases in the hard to assemble regions (**Extended Data Fig. 8a**). The trio-based ONT contig assemblies had lower haplotype phase blocks (NG50s  $\sim 3\text{-}6$  Mb; **Fig. 1c**), and higher haplotype switch errors ( $\sim 0.3\text{-}0.5\%$ ; **Extended Data Fig. 7a**), presumably due to their higher sequence error rates. In contrast to prior findings<sup>8</sup>, not all trio-based assemblies had lower haplotype false duplication rates, as assessed by either k-mers or BUSCO duplicate gene copies (**Extended Data Fig. 2d**, **Supplementary Table 2g**). This appears to be due to improvements in the higher read accuracy of Pacbio HiFi vs CLR and non-trio based algorithms.

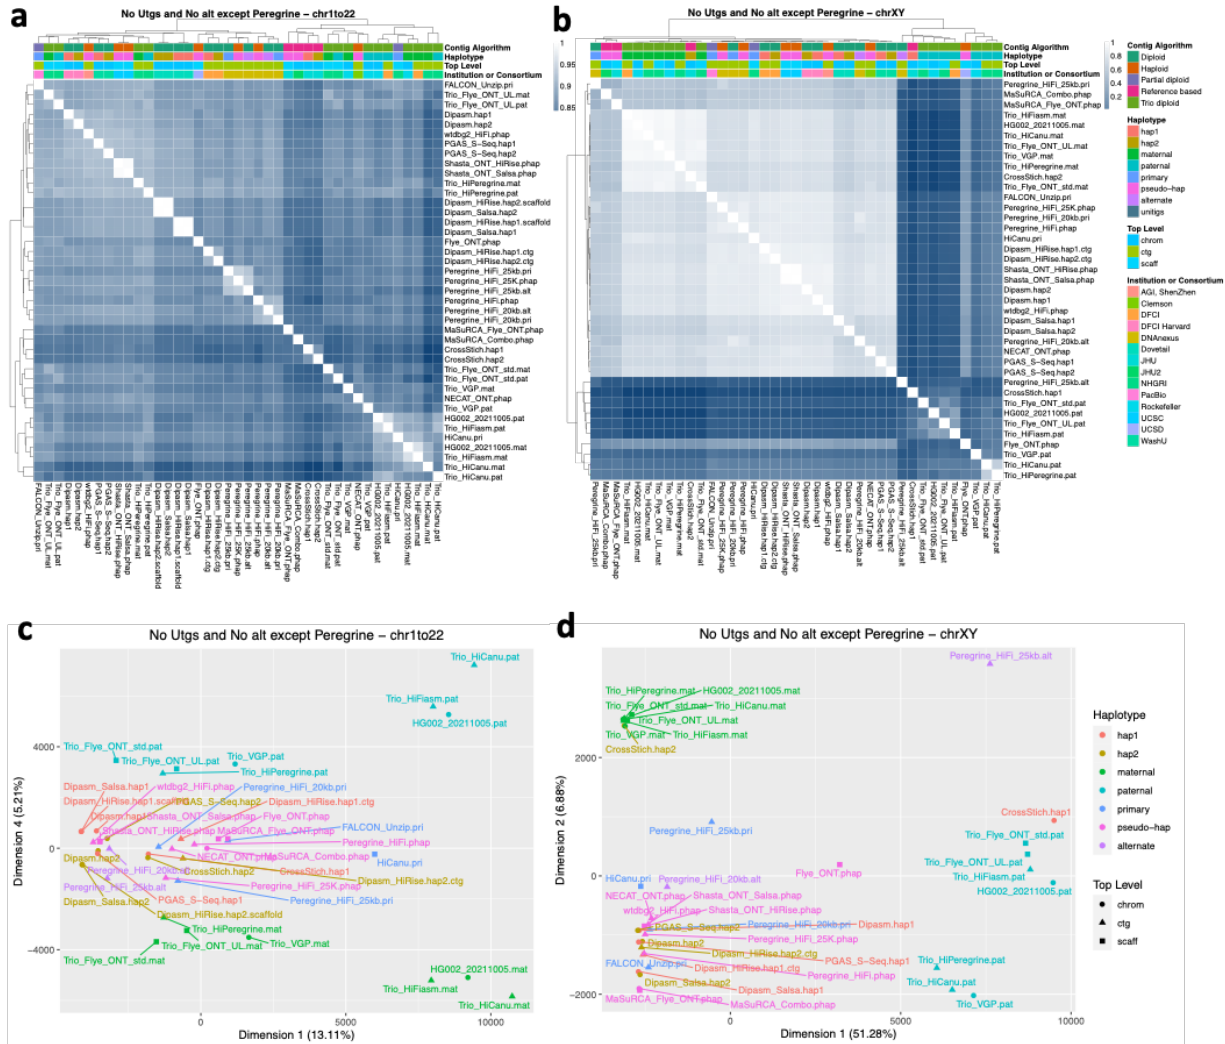
### Graph phasing assembly is most complete and accurate

The trio-based approaches fell into two principal categories: 1) those that use parental reads to bin the child's reads by haplotype before assembly (e.g. Trio VGP, Trio Flye, Trio HiCanu, Trio Peregrine); or 2) those that generate an assembly graph of the child's genome first and then label haplotypes in the graph using the parental reads (e.g. Trio hifiasm). As presented in a complementary study conducted simultaneously<sup>28</sup> and further advanced here, we find that the graph-based phasing approach generally outperforms the two-step binning trio approach when high-accuracy long reads are used to build the initial assembly graph. In particular, among the diploid assemblies, the Trio hifiasm maternal (asm9a) and paternal (asm9b) assemblies had the highest combination of high-quality metrics, including the highest QV accuracies (**Fig. 1d**), the third highest NG50 haplotype phase blocks (**Fig. 1c**; Trio VGP was the highest), among the least false duplications (**Extended Data Fig. 2d**), the fewest contigs (**Extended Data Fig. 3a**), among the lowest haplotype switch errors (**Extended Data Fig. 7a**), the highest genome completeness (**Supplementary Table 2k**), and the least collapsed repeats (**Extended Data Fig. 8a**). These findings indicate that a graph-based phasing assembly is more accurate and complete as the combination of the graph with haplotype information can correct errors made by either method alone. A prerequisite to highly accurate graph-based haplotype phasing is well-resolved diploid assembly graph as generated from high-accuracy long reads (e.g. HiFi).

### Shared properties of a pan-assembly alignment

To identify both shared and distinct features of the assemblies, we utilized a pangenomic approach, performing an all-vs-all alignment for 45 assemblies (both haplotypes each; **Extended Data Fig. 10a**), excluding the alternate contigs or unitigs of pseudohaplotype assemblies as they were highly fragmented. We annotated the alignment according to chromosomes in GRCh38 and CHM13. Pairwise Jaccard similarity analyses on the autosomes (Chr 1-22) clustered the Trio hifiasm and Trio HiCanu assemblies as more similar to each other and distinct from the other assemblies (**Fig. 2a**); at one branch higher, these trio assemblies clustered with the other trio assemblies (except Trio HiPeregrine) and with the MaSuRCA and CrossStich reference-based assemblies. The remaining super-cluster of assemblies subclustered mostly by assembly pipeline, indicating that approach drives the assembly similarities the most. More pronounced than the autosomes, the Jaccard similarity analyses on the XY sex chromosomes grouped all of the trio paternal assemblies in one cluster, with distinctions amongst themselves, relative to nearly all of the

remaining assemblies into a sister supercluster with the trio maternal assemblies (**Fig. 2b**), indicating that most non-trio based assemblers were more successful at putting together an X chromosome than a Y chromosome. Two exceptions were the haploid Flye ONT.phap assembly (asm5) and the reference-based CrossStich hap1 assembly (asm17a), which grouped with the trio paternal assemblies and had a more complete Y chromosome (asm17a due to using the GRCh38 Y as a reference). PCA multidimensional analyses supported these conclusions, where in the 4th dimension the trio-based autosomes 1 through 22 clustered by parental haplotype without the presence of the sex chromosomes (**Fig. 2c,d**). The Trio hifiasm and Trio HiCanu autosome assemblies were the most distinctly clustered by parental haplotype. Clustering on each chromosome alone, 15 of the 22 autosomes (68%) had similar patterns (**Supplemental Table 4**). These findings suggest that there is sufficient biological signal among the autosomes to separate them by parental haplotype source.



**Figure 2. Multidimensional relationship among assemblies.** a and b, Clustering of pairwise Jaccard similarities between pairs of assemblies, for the autosomes 1-22 and sex chromosomes XY, respectively. Blue heatmap, the lighter the blue (closer to 1) the more similarity between pairs of assemblies. At the top of each heatmap, assemblies are annotated with four different color-coded classifications. c and d, Principle component analysis on the multidimensional Euclidean distances amongst assemblies, for the autosomes



1-22 and sex chromosomes XY, respectively. PCA dimensions shown are those where the paternal and maternal haplotypes separated out the strongest.

### Combined approaches for an HPRC HG002 reference

Based on our findings, we developed a pipeline that combines the best practices of the different approaches and used it to generate a higher-quality diploid *de novo* assembly (**Extended Data Fig. 10b**). We first removed the remaining HiFi reads with unremoved vectors (adaptors) using HiFiAdapterFilter (**Supplementary Note 2**). We then generated HiFi maternal and paternal contigs with graph-based haplotype phasing of Trio Hifiasm v0.14.1. This latest version incorporates bug fixes we found after generating the initial HG002 assemblies, including: a) enhancing contig QV by constructing the contig golden path through high-quality portions of corrected reads; b) resolving more segmental duplications by selecting high-occurrence seeds at the overlapping stage; and c) improving contig N50 by rescuing contained reads that break contigs on one haplotype when the read actually comes from the other haplotype<sup>42</sup>. In addition, we titrated HiFi coverage with hifiasm and found a level (~160X) given the data that yielded an optimal contiguity and haplotype switch error (**Supplementary Fig. 2**). We then separately scaffolded the paternal and maternal HiFi-based contigs with paternal and maternal Bionano optical maps. Conflicts between the HiFi contigs and Bionano optical maps were manually evaluated, where we accepted 5 of 15 maternal and 3 of 13 paternal joins or breaks indicated by the Bionano maps (**Supplemental Table 5a**). The majority of these conflicts (25 of 28 total) were in segmental duplications, near centromeres, particularly of the acrocentric chromosomes (Chr 15, 21, 22), and included natural haplotype SV differences in HG002; the remaining three were in known tandemly repeated genes (*IgK*, *IgH*, and *TSP*), with the first as known sources of programmatic structural variation associated with B lymphocytes. We then further scaffolded the paternal and maternal assemblies with haplotype-filtered (Meryl) Hi-C (OmniC) data and the Salsa 2.3 algorithm. Scaffolding with Arima Hi-C v2 data yielded similar results. We performed a manual curation with Hi-C contact maps, which resulted in less than ~10 breaks and ~50 additional joins in each haplotype assembly (**Extended Data Fig. 4b**, **Supplemental Tables 2h, 4b**). Most of the breaks were at centromeres to allow their satellite placement.

Next, we performed gap filling with a conservative version of the pipeline used to fill the gaps in the initial draft of the T2T CHM13 assembly<sup>5</sup>. ONT-UL reads were base recalled with Guppy 4.2.2, binned with trio-Canu, and then assembled into haplotype specific contigs using Flye. Draft ONT-UL contigs were polished to increase consensus accuracy. Variant calls were generated using Medaka and ONT long reads, and filtered with Merfin<sup>43</sup> using kmers from Illumina short reads and then applying them to increase the quality of the consensus sequence. The polished contigs were aligned to their respective haplotypes of the curated HiFi-based scaffolds from the Hi-C step above, and used to fill gaps. This resulted in 10 and 5 gaps filled in the maternal and paternal assemblies respectively. Of these 15 gaps, 10 contained GA-rich repeats and 2 were long segmental duplications (**Supplementary Fig. 1**). A final manual curation fixed 37 items in the maternal and 60 in the paternal assemblies (**Supplementary Tables 2h and 4c**), much fewer than the hundreds of manual fixes that normally would be required (e.g. **Extended Data Fig. 4a**). A contamination screen removed multiple copies (41 maternal; 45 paternal) of endogenous human EBV viral genomes, present as single contigs, as well as a yeast contig in the paternal assembly; we did not find any non-human contamination within the human contigs and scaffolds. Approximately 98% of the remaining assembled sequence was assignable to the 22 autosomes and the X and Y sex chromosomes. These new assemblies were named HPRC-HG002.mat and HPRC-HG002.pat references.

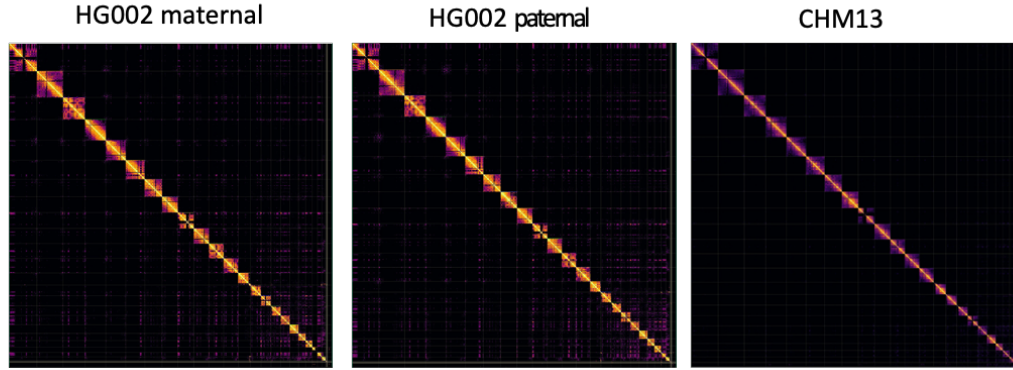
These two *de novo* assemblies had the highest quality of all metrics compared to the bakeoff assemblies, including the referenced-based assemblies and the GRCh38 reference itself: the largest contig (62.9 and 81.6 Mb) and comparable scaffold (154.4 and 146.7 Mb) NG50s, close to the theoretical maximum for the scaffold values (**Fig. 1a,b**); the fewest contigs and gaps in scaffolds (**Extended Data Fig.**




**3a-c**); the highest QVs (~60; **Fig. 1d**); the most complete haplotype phasing (99.7 maternal, 99.0% paternal; **Fig. 3a,b**) with NG50 phase blocks of 106.7 and 90.4 Mb, respectively (**Fig. 1c, Supplementary Table 2j**); the least collapsed repeats (18.5 and 17.6 Mb, respectively; **Extended Data Fig. 8a**); and among the highest annotation metrics (**Extended Data Fig. 9; Supplementary Table 2m**) and the highest SNV and small indel TP rate (**Supplementary Table 3a,b**). They clustered closest with the Trio hifiasm and Trio Canu assemblies (**Fig. 2**). Assessing against GIAB HG002 benchmarks further, this diploid assembly produced highly accurate SNV concordance (F1 score) of 99.7% and small indel concordance of 98.6% between HPRC-HG002.mat and HPRC-HG002.pat haplotypes. These concordance values were 0.2% and 0.8% lower than the best-performing mapping-based variant callers in a 2020 precision FDA Truth Challenge<sup>44</sup>. We found that 70% of the discordant SNVs fell in segmental duplications, most in the regions of complex SVs that could not be accurately benchmarked. In fact, many of these differences appeared to be more accurate in the new HPRC HG002 assembly than in the mapping-based benchmark or precisionFDA entries. The primary limitation of the assemblies were small indels in homopolymers and 51 to 200 bp tandem repeats (i.e structural variants), making up 80% of all discordant indels; curation revealed that the final HPRC assembly had infrequent errors due to collapsing haplotypes and/or to noise in starting HiFi reads. When benchmarking larger SVs in the new HG002 assemblies with respect to GRCh37 against the GIAB v0.6 SV benchmark, which excludes segmental duplications and centromeres<sup>45</sup>, TP was 98% (compared to 93% for asm9a,b) and precision was 89%, and most putative errors were actually just differences in representation of SVs in tandem repeats or errors in the benchmark. We found that some repetitive gene families that are known to be difficult to assemble were completely assembled, including the major histocompatibility complex (MHC) genes, where variants were > 99.99% concordant with the GIAB v4.2.1 benchmark (**Extended Data Fig. 10c**). Overall, this high concordance between the assembly-based variants and existing benchmarks for all variant types, including some regions where the assembly is more accurate than the benchmarks, demonstrates substantial promise for phased, whole genome assembly-based benchmarks that include challenging genomics regions.

All of these metrics, particularly for the HG002 maternal haplotype, were on par with the CHM13 T2T assembly (**Fig. 1; Extended Data Figs. 2-3 and 6-9**), including a scaffold NG50 (154 MB) similar to CHM13's contig NG50 (154Mb, essentially chr X). We aligned the two HG002 haplotype assemblies to CHM13 (+Y from GRCh38), and found very good length correlations to it (**Figs. 3c and 4a; Supplementary Table 7**). The HG002 chromosomes were nearly completely phased correctly for each haplotype (**Fig. 3b**), with only three small exceptions being small contigs in the centromeres of Chr 3, 6, and 17 (**Fig. 4a**); the two ends of Chr Y aligned to CHM13 chr X, because they are missing in the GRCh38 reference. Nearly all chromosomes were within 90-100% of CHM13's length, with most (16 maternal and 15 paternal) within  $\pm 1\%$ , suggesting that the missing alignments in parts of the centromeres are likely due to much greater divergence in centromeres from the CHM13 assembly. The biggest exceptions were the short arms of the acrocentric chromosomes 13, 14, 15, 21, 22, with Chr 21, 22 being the two outliers at ~85% of CHM13's length for the maternal and ~75% for the paternal haplotype (**Fig. 3d**); the acrocentric regions of these chromosomes are notoriously difficult to assemble owing to their highly repetitive structure comprising an rDNA array, numerous satellite arrays, and segmentally duplicated sequences<sup>5</sup>. Consistent with this, the several hundred remaining small unplaced scaffolds appeared to be largely centromeric, satellites, and parts of the Y chromosome. Nevertheless, most chromosomes have only 1-12 remaining gaps (**Fig. 3e**), and the maternal Chr 11 and 12 have no gaps (e.g. one contig each). We checked suggesting that a mostly automated phased T2T assembly is now possible for some human chromosomes.


a, Hi-C contact maps after curation




b, Phase plot of scaffolds




c, HG002 w/o Ns:CHM13 %




d, HG002 chromosome sizes relative to CHM13

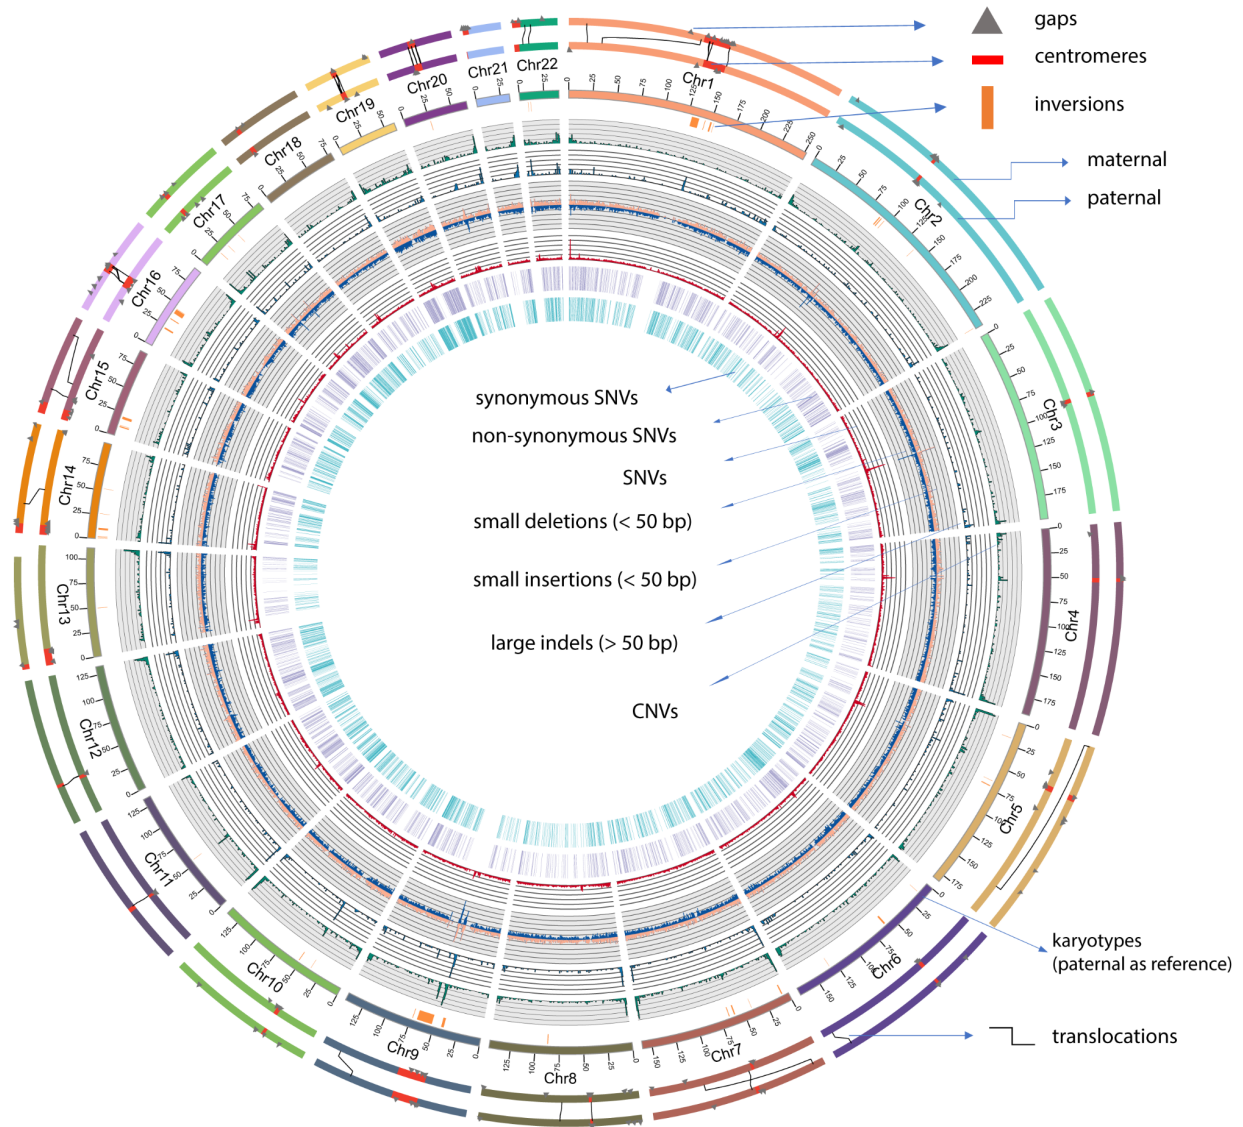


e, Remaining gaps in HG002 chromosomes



**Fig. 3 | Near complete haplotype separation of scaffolds.** a, Hi-C contact maps to the final curated HPRC-HG002 maternal and paternal assemblies in comparison to the CHM13 assembly. Numbers designated chromosome numbers. b, Blob plot using Illumina parental  $k$ -mers for the scaffolds of the HPRC HG002 haplotypes. Axes in sqrt scale to highlight the small phase blocks. c, % size of HG002 diploid assembled chromosome relative to CHM13 chromosomes. d, absolute size values of all three assemblies compared. e, Remaining gaps in each chromosome of each HG002 haplotype. Chr Y has the most gaps, maternal Chr 11 and 12 have no gaps, highlighted with \*.






**Figure 4. HPRC HG002 features.** **a**, Chromosome (scaffold) alignments between HPRC-HG002 maternal and paternal assemblies and CHM13. Haplotype separation is nearly complete, and thus colors are solid blue (paternal) and red (maternal). Black tick marks, gaps between contigs or unaligned divergent sequence with CHM13. Alignment gaps represent structural variation or missing sequence in the HG002 assembly compared to CHM13. Grey, regions enriched with centromeric satellites. The Y chromosome reference is from GRCh38, since CHM13 does not have a Y. **b**, Circos plot of heterozygosity landscape between the two haploid assemblies of HG002. Tracks from inside out: synonymous amino acid changes; non-synonymous changes; SNV density (window size 500kb, range 0~3.1%), small deletion and small insertion (< 50 bp) density (window size 1 Mb, count 0~850); large indel density ( $\geq 50$  bp, window size 1 Mb, count 0~20), CNV density (window size 1 Mb, count 0~40). The black line links in the outermost circles denote translocations between paternal (inner) and maternal (outer) assemblies. Orange bars, inversions ( $\geq 50$  bp). Red bars, centromeres, Grey triangles, gaps.

### Missing genes among individuals and haplotypes

The similar number of autosomal genes that did not align to both HG002 haplotypes (68 mat and 75 pat) and CHM13 (66; **Supplemental Table 2m**) led us to manually curate this annotated data set for assessing possible gene loss diversity between individuals. We found 20 genes absent from all four reference

assemblies (GRCh38, HPRC HG002.mat, HPRC HG002.pat, and CHM13), and greater overlap of 74% (55 of 68 genes) not present in CHM13 and one or both haplotypes of HG002 (**Fig. 5**). The inverse had lower overlap, with 62% (41/66) overlap for the HG002.mat and 58% (44/76) for the HG002.pat haplotype also absent in CHM13. Similarly, the two maternal and paternal haplotypes of HG002 shared 65% (43/66) and 56% (43/76) of gene loss with each other, respectively. Conversely, CHM13 and each haplotype of HG002 had uniquely 12-18 genes absent (**Fig. 5**). However, about 15 of the total HPRC HG002.pat unaligned genes were present in one or more of the Trio paternal bakeoff assemblies, indicating that either they were missed in the new HPRC HG002 reference assembly or were false haplotype duplications in the bakeoff assemblies. Supporting the latter idea, more than half (~67%) of the total of 120 genes missing among the four reference assemblies were in repetitive gene families, where false haplotype duplications are more easily generated (REF) (**Supplementary Table 6**). These genes include the MHC HLA immune cluster, keratin associated proteins (KRTAP), olfactory receptors (ORs), RNA 18S and 5-8S genes, several Long Intergenic Non-Protein Coding RNA (LINC), and over 30 micro RNAs (MIRs). Most of the genes absent in the final references were absent in all of the Trio-based assemblies, indicating real biological loss differences among individuals and haplotypes of an individual. There were about a dozen genes present in GRCh38 and the asm17 referenced based assembly, but not in any of the other HG002 assemblies or CHM13, illustrating the possible bias of reference-based assembly methods.



**Fig. 5 | Genes unaligned and thus presumed absent in the four main reference assemblies compared.** n = # of genes absent in reference assemblies. Values in 4-way Venn diagram are the number of shared or uniquely absent genes among the 4 assemblies.

### Greater diversity between haplotypes

Given more complete HPRC HG002 assemblies relative to GRCh38, and diploid relative to the CHM13 haploid, we performed a detailed analysis of heterozygosity between the haplotypes, following approaches we used on a VGP trio-based marmoset assembly<sup>46</sup>. We noted a remarkably high amount of heterozygosity (2.06% of total bp; ~2.6 million SNVs; ~631K small SVs < 50 bp; 10.4K large SVs ≥ 50 bp; **Fig. 4b**; **Supplemental Table 8**), or 3,263,480 variants total. Most of the additional variation was in the newly assembled centromeres, where there was a sharp peak in SNV, indels, inversions, and translocations (**Fig. 4b**). When not including the centromeres, heterozygosity was within the normal expected range 1.2% (REF). Most of the remaining gaps were also in the centromeres; although Chrs 4, 11, 12, 18 had 0 or 1 gap each for one or both haplotypes, indicating that the increased diversity is not primarily due to assembly errors. The SVs included 59 large inversions, on Chrs 1, 6, 9, 14, 16, with the largest being a 15 Mb inversion on Chr 9. There were 38 translocations and 7,330 copy number variations (CNVs). Of the CNVs, 260 were in protein coding genes, 3.5-fold higher than the average of 75 determined from less complete short-read assemblies from the 1000 genomes project<sup>47</sup> but more similar to rates reported for structural variants of



human genomes using long-reads (citations, multiple). Of these, 9 genes had remarkable differences between haplotypes (**Fig. 4b**), and included: the neuroblastoma breakpoint family member 19 (NBPF19) present in 1 copy in the maternal haplotype, but 46 in the paternal haplotype; FAM90A8P also present at 1 copy in the maternal, but 17 in the paternal; and the NPIPA7 brain enriched gene with 22 copies in the maternal and 8 in paternal haplotypes. All three genes are thought to be primate specific or more rapidly evolving in primates<sup>48–50</sup>. Among the 12,736 SNVs located in CDS, 5,844 were nonsynonymous and 6,397 (50%) were synonymous, changing the amino acid sequence between haplotypes (**Fig. 4d**). The increased diversity at the centromeres, although expected (REF), was not seen at this level in a recent marmoset trio genome assembly using a similar type of analysis<sup>46</sup>. We believe the difference is that for the marmoset trio, CLR reads were used, leading to potential collapsing of the centromeric repeats, whereas for the human HG002 assembly, the higher accuracy HiFi reads were used, allowing separation of highly similar repeats and more centromere diversity to be revealed. Overall, a more complete human diploid genome assembly reveals a greater amount of genetic diversity than otherwise expected.

### A look towards the future

This study allowed us to determine which current approaches yield the highest quality metrics for *high-quality diploid maternal and paternal assemblies*. Key factors were the use of trio-parental sequence data to sort haplotype sequences in the child, a graph-based approach to resolving these haplotypes during the assembly process rather than before or after it, and combining different sequence data types and assembly tools where one approach captures information missed by another. Haplotype binning of reads before assembly (e.g. Trio HiCanu) could mispartition reads, leading to lower metrics. Haplotype local phasing during a graph-based assembly (e.g. hifiasm) could prevent such errors in partitioning reads.

These findings confirm and advance upon those recently reported by the VGP, HGSVC and T2T assembly pipelines. For the initial VGP pipeline, we used Pacbio CLR, which was less accurate relative to the more recent Pacbio HiFi read technology. This meant that short read polishing to increase accuracy of the long read-based contigs was no longer necessary. More accurate long reads allowed generation of larger contigs, reduction of collapsing repetitive sequences in the centromeres, and increased haplotype phasing accuracy. For the latter, instead of FALCON-Unzip that had produced a more complete pseudohaplotype and a fragmented alternate haplotype, hifiasm, DipAsm, PGAS and CrossStitch produce two comparable pseudohaplotypes. Comparisons of the same samples sequenced by both PacBio CLR and HiFi by the HGSVC showed that HiFi produced phased genome assemblies that were significantly more accurate than those generated with CLR (Ebert Ref #12) and approaches that incorporated parent-child trio data significantly reduced phasing errors (Chaisson, 2019 PMID: 30559433). An advance adopted from the T2T approach used on CHM13 was development of a tool for automated incorporation of polished ONT reads for gap filling, but for both haplotypes, independently. We also made advances on the Trio assembly approach, by not only haplotype phasing the starting long reads and Hi-C, but the ONT reads and the Bionano optical maps. These advancements lead to nearly completely phased haplotypes, and near complete assembly of chromosome arms. All the components of the current pipeline developed here are available on the Galaxy computational platform, and in modular form with different steps that can be included or excluded. What remains is developing diploid assembly methods that prevent the remaining collapses and gaps in the centromeric satellite arrays, large human satellite arrays, and short arms of the acrocentric chromosomes.

Based on the findings in this study, the HPRC has decided to use the trio graph-partitioning approach of hifiasm to generate the base contigs of the first 47 individuals (94 haplotype genomes) that will contribute to generating the first human pangenome reference (BioProject PRJNA730822). Similarly, the trio and a non-trio version of hifiasm followed by the scaffolding approach used here for HPRC HG002 have been adopted by other large-scale sequencing projects, such as the VGP, the Earth Biogenome Project (EBP), and the Darwin Tree of Life Project (DTOL). Improvements have also been made to some of the

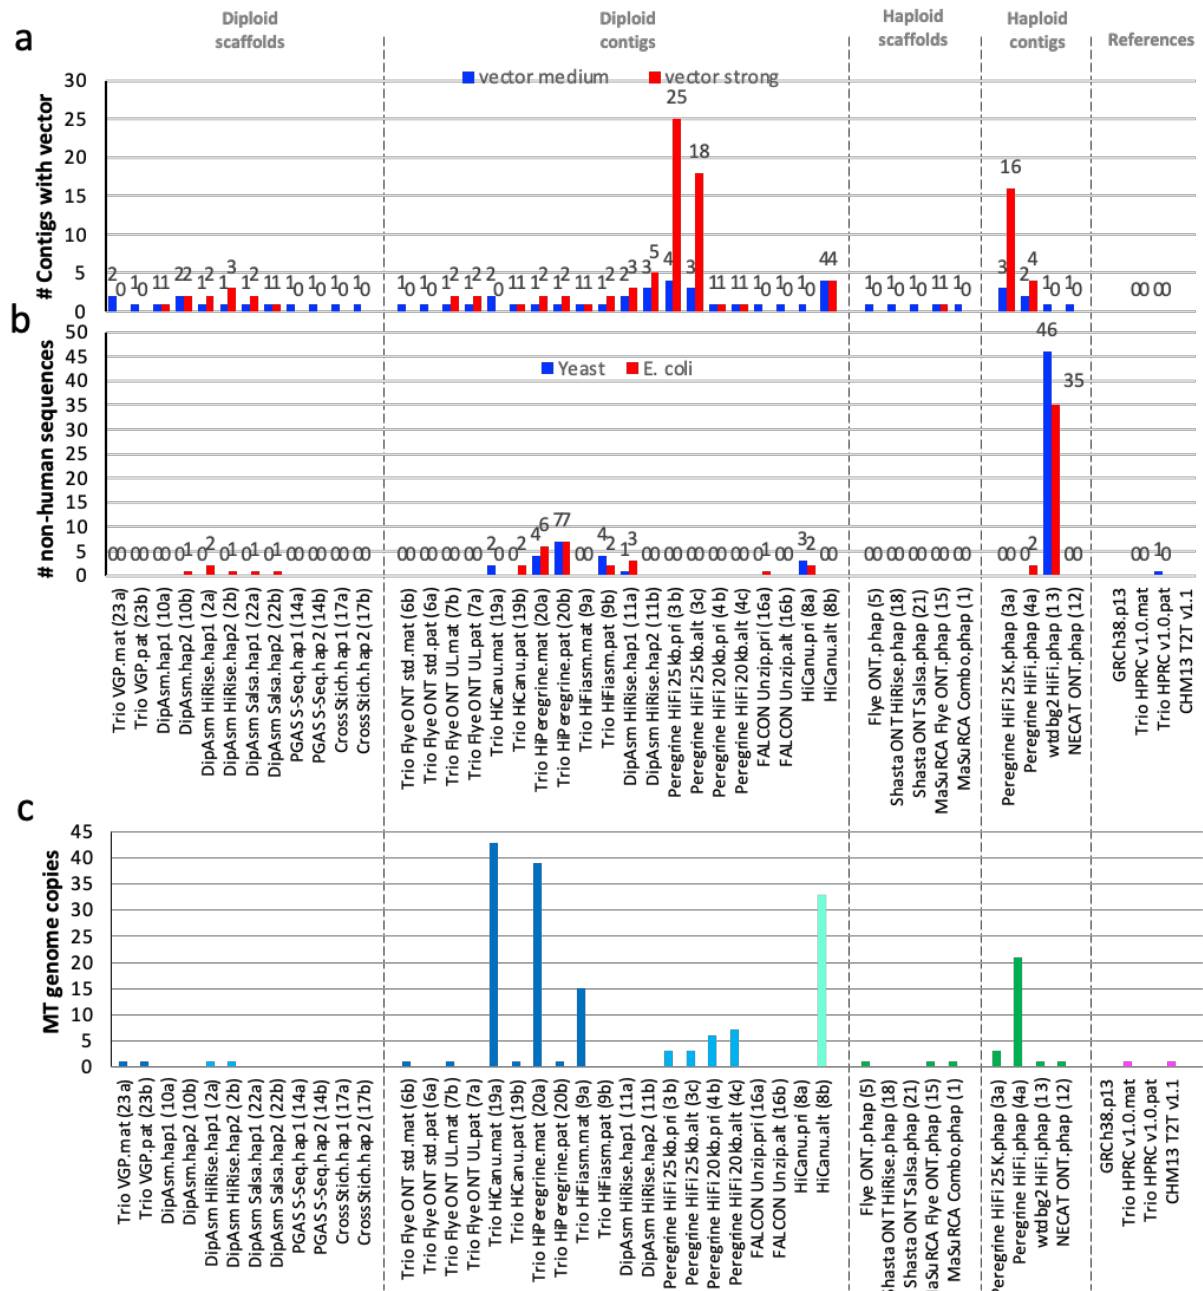
other assembly algorithms since the versions tested here(REF) including on HG002; but thus far, the trio-graph-based approach with trio-based scaffolding still yields the best combination of the highest quality metrics.

Future efforts will be necessary to develop a phasing method that does not require parental sequence data that works as well as a trio method. This will make it possible to generate equivalently good diploid reference assemblies for human and non-human organisms where parental data may not be available. Towards this end, using Hi-C or Strand-seq data for haplotype phasing are promising alternatives, as both data types contain within-chromosome haplotype information of an individual. Two methods have successfully used Hi-C, FALCON Phase<sup>20</sup> and hifiasm (Hi-C)<sup>51</sup>, and another has used Strand-Seq<sup>52</sup> to generate maternal and paternal phased long-read based human genome assemblies with relatively few switch errors, including on HG002. As with trio binning, these approaches appear to work best when phasing is integrated with the assembly process, but further improvements are necessary to match or surpass the quality metrics seen with a parental trio-graph-based approach.

We used ONT-UL reads to fill in GA repeat-rich and other challenging sequence gaps in the HiFi-based contigs. As an alternative, the PacBio CLR reads that do not make it to HiFi accuracy levels contain some of the GA-rich repetitive sequences, and could be used to fill some of these gaps, as currently done in the VGP 2.1 pipeline on the Galaxy platform ([https://training.galaxyproject.org/training-material/topics/assembly/tutorials/vgp\\_genome\\_assembly/tutorial.html](https://training.galaxyproject.org/training-material/topics/assembly/tutorials/vgp_genome_assembly/tutorial.html)). The remaining handful of gaps per chromosome in the HG002 assemblies are mainly restricted to the hard to assemble regions around segmental duplications, centromeres, telomeres, rDNA arrays, and other complex repeats, many with differences between haplotypes. Direct integration of ONT-UL data within the assembly graph was necessary for finishing these regions in the T2T CHM13 assembly<sup>5</sup>. Thus, integration of both HiFi and ONT-UL data in a diploid assembly graph, combined with long-range phasing information from trios, Hi-C, or Strand-seq may soon enable automated T2T diploid genome assembly. For those who wish to contribute assemblies to the human pangenome references, we encourage them to utilize our recommended process to obtain the highest quality assemblies possible, and contribute new methods to further improve the quality and completeness of human and other assemblies.


Using currently available methods and technologies, we were able to complete over 98.5% of both HG002 haplotypes. The new biological discoveries made here demonstrate that even with a single individual, additional genetic diversity contributing to the human population can be found. Using these methods for the generation of additional diploid human genomes and creation of a human reference pangenome we believe will enable a more complete picture of human genetic diversity, greater accuracy for precision medicine, and a greater understanding of the biology of genomes.

## **Extended Data Figures**



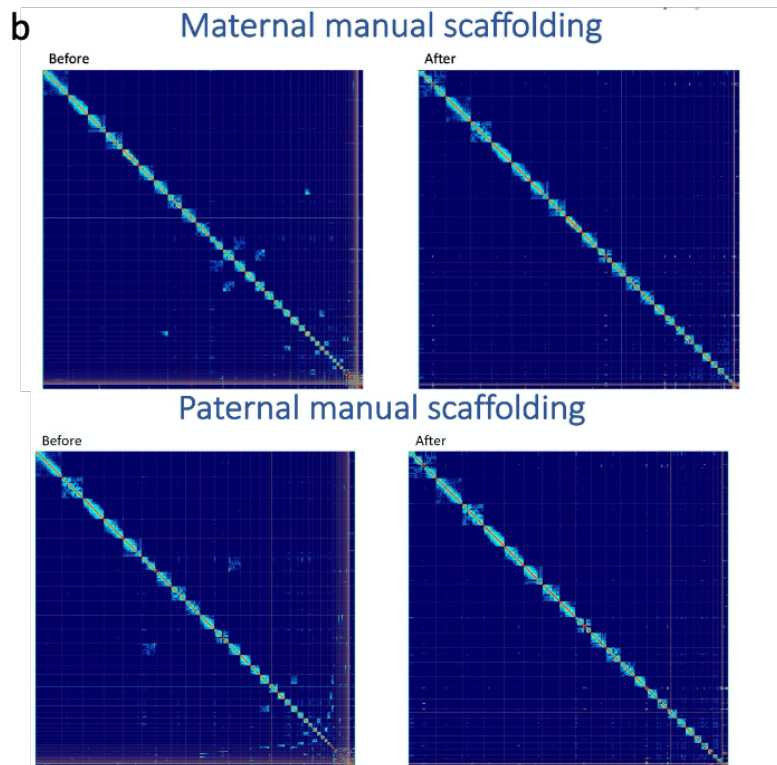
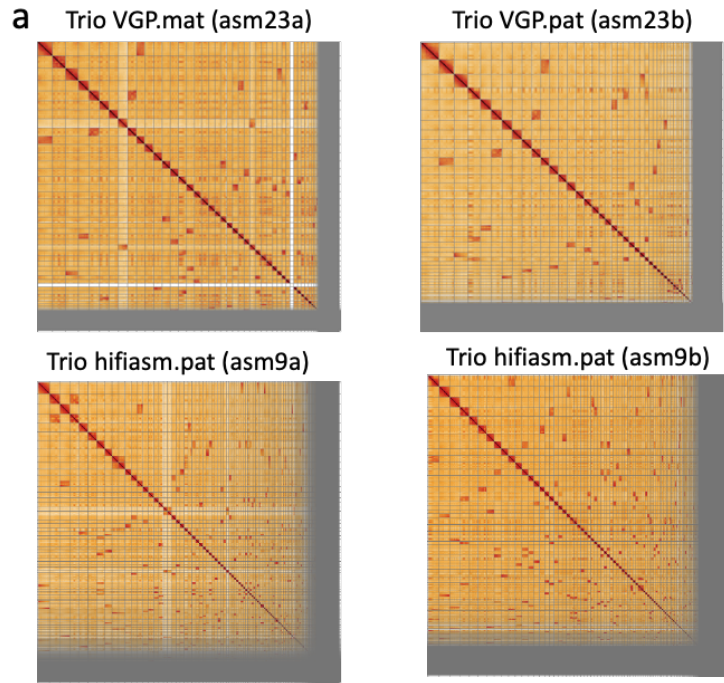
**Extended Data Fig. 1 | Non-human and organelle genomes found in the human genome assemblies.** **a**, The number of contigs that had remaining library clone vector sequences in each assembly. **Medium is XX-XX% identity; strong is XX-100% identity.** **b**, The number of contigs with non-human yeast and *E.coli* sequences. Values above columns are the specific numbers. **c**, The number of endogenous mitochondrial genome copies found in each assembly. The bar color code is the same as described in Fig. 1.






**Extended Data Fig. 2 | Assembly size and false duplication metrics.** **a**, Percent assembly sizes of expected maternal with Chr X (3,054,832,041 bp) or paternal with Y (2,995,432,041 bp) for trio-based assemblies, or simply relative to maternal size for all other assemblies. **b**, Total summed length of all contigs. **c**, Total summed length of scaffolds, with proportion contributed by Ns (red) in gaps. **d**, estimated percent of assembly size that is due to false duplications based on k-mer values for each haplotype. Abbreviations and color coding explanations are the same as in Fig. 1 legend.







**Extended Data Fig. 4 | Hi-C contact maps. a**, Example Hi-C contact maps for bakeoff maternal (mat) and paternal (pat) haplotype assemblies. The Trio VGP scaffolded assembly has several dozen large joins and many small ones to make from the off-diagonal signals. The Trio hifiasm contig only assembly as expected has many more needed. **b**, Reference HPRC HG002 assemblies for each haplotype before and after manual curation, showing less off diagonal signals and no major scaffolds/contigs not placed in chromosomes after curation.



**Extended Data Fig. 5 | Curation structural analyses.** **a**, Missed contig joins that should have been brought together in the same scaffold. **b**, Misjoins of contigs that should not have been brought together. **c**, other scaffold errors, including false inversions and translocations. **d**, Contig errors, including chimeric contigs. ND = not done, as it would be redundant with another assembly approach. NA = Not applicable, as these references were curated by a different process, with nevertheless many errors corrected. Other abbreviations and color coding explanations are the same as in Fig. 1 legend.




**Extended Data Fig. 6 | Variant benchmarking.** **a**, True positive percent of known SNVs found between HG002 haplotypes in each assembly. **b**, True positive percent of known small indels found between HG002 haplotypes in each assembly. For the diploid assemblies, comparisons were made between the two haplotypes (maternal vs paternal for the trio assemblies; haplotype 1 vs haplotype 2 for the non-trio assemblies). For the haploid assemblies, we scored as TP if at least one of the variants were found. Abbreviations and color coding explanations are the same as in Fig. 1 legend.



**Extended Data Fig. 7 | Haplotype phasing metrics.** **a**, Haplotype switch errors within scaffolds and/or contigs of each assembly (lower % is more accurate). Values written in the graphs for the haploid assemblies (greens) are off the chart, in order to not mask the lower switch error values of most other assemblies. **b**, Total Gb of each assembly that has been haplotype phased (~3.0 is the theoretical maximum of the maternal haplotype; 2.9 for the paternal). **c**, Haplotype phasing completeness according to parental k-mer statistics for each assembly. A complete phased and complete assembly will have both maternal (blue) and paternal (red) each at 100% without mixture from the other. The trio approaches had nearly full phase separation, whereas the non-trio approaches nearly had half and half separation because there was not an attempt to phase across contigs or scaffolds/chromosomes belonging to the same maternal or paternal haplotypes. Combined values over 100% indicate a mixture of haplotype presumably due to false duplications; although values under 100% could still have false duplications. Other abbreviations and color coding explanations are the same as in Fig. 1 legend.








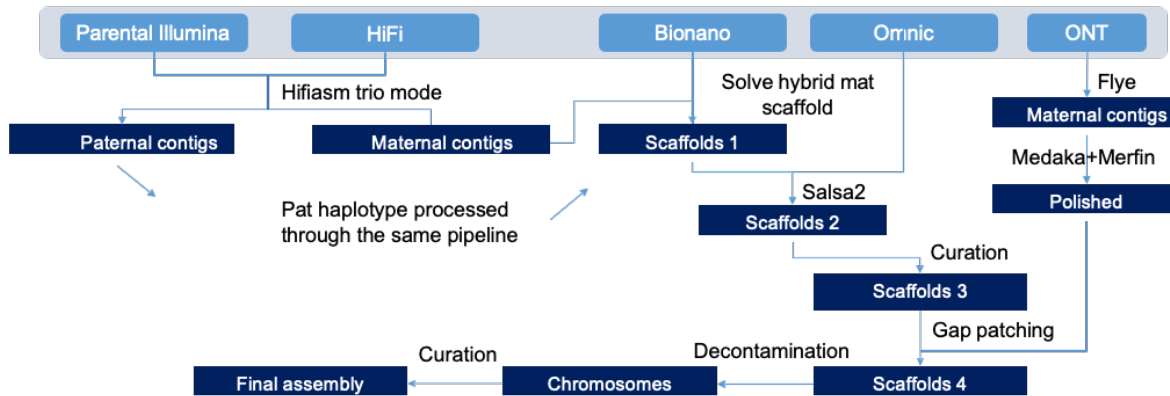
**Extended Data Fig. 9 | Annotation benchmarking.** **a**, Side-by-side comparisons of gene transcripts that did not align to each assembly (blue) versus those that were split between two or more scaffolds/contigs (red). **b**, Number of genes that had less than 95% the length covered in the assembly. **c**, Genes in the assemblies with overlapping transcripts due to possible collapse in the assemblies. **d**, Genes requiring frameshift corrections to make a complete protein. Values written in the graphs for those off the chart, in order to not mask the lower values of most other assemblies. Abbreviations and color coding explanations are the same as in Fig. 1 legend.



### a, Orientation vs. whole genome alignments across assemblies



### b, HPRC Trio pipeline v1.0 used to generate first HPRC HG002 references

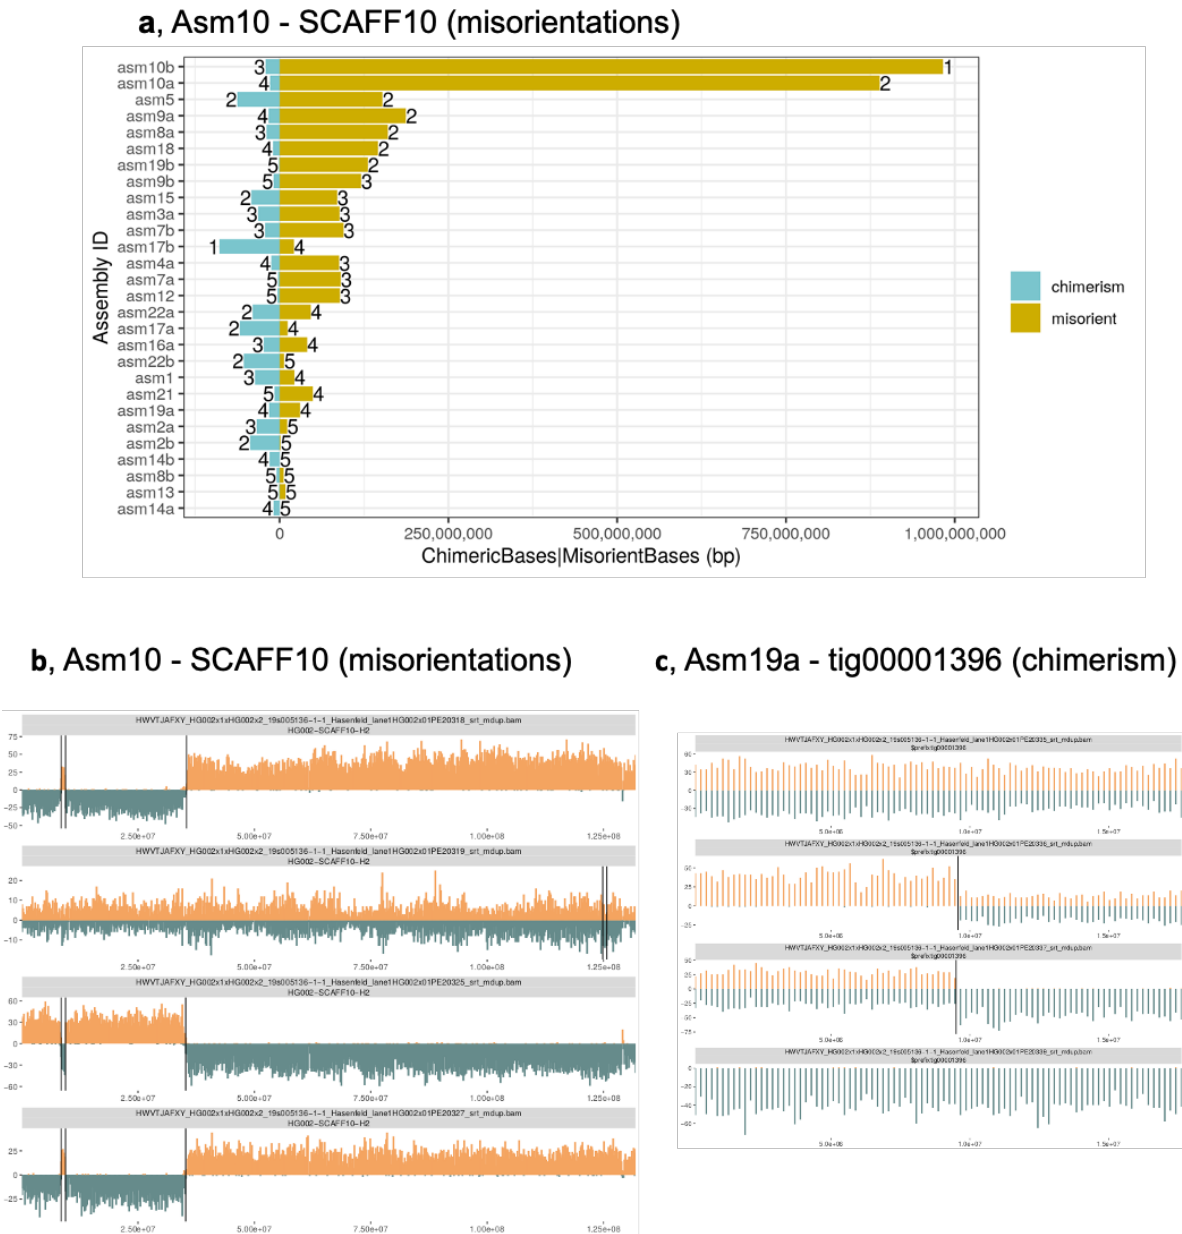


### c, Complete sequence of both C4A and C4B genes and their haplotypes.




**Extended Data Fig. 10 | Pangenome alignment and generation of high-quality platinum HG002 dioploid assemblies.** **a**, Output of graph-based alignment of all chromosomes concatenated from all 45 HG002 assemblies of this study (both haplotypes of dioploid assemblies). Red vs Black, are same orientation. Connections at the bottom, is a clustering of the alignments. **b**, HPRC v1.0 pipeline developed to produce the reference platinum quality HG002 maternal and paternal assemblies of this study. All steps shown are highlighted for the maternal data, but the paternal data went through the same pipeline. **c**, Alignment of four paralogs of the C4A and C4B genes form the HPRC-HG002 haplotypes with CRCh38 and CHM13, showing complete assemblies with no gaps, and the variants in these genes were 100% concordant with the GIAB HG002 v4.2.1 benchmark.

Supplementary Figures



**Figure 2. Strand-Seq identification of large scaffolding and contig errors.** **a**, Quantification of chimeric and misorientations errors for each assembly. Only contigs or scaffolds  $\geq 1$  Mp were analyzed, and chimeric or collapses contigs  $\geq 1$  Mp and misorientations  $\geq 400$  kb were analyzed. Number at each bar sets the quality category for each assembly based on 0.25, 0.5, 0.75 quantiles **b**, Raw data examples of errors detected with Strand-seq analyses, using sequence coverage on the + (orange) and - (green) strands.



Supplementary Fig. 2.

[illegible][illegible]

32



arising during cell growth. For paternal (HG003) and maternal (HG004) samples, DNA was extracted from cell lines publicly available as GM24149 (RRID:CVCL\_1C54) and GM24143 (RRID:CVCL\_1C48), respectively, and Illumina sequencing of DNA from NIST RM 8392 (containing vials of HG002, HG003, and HG004) was used.

#### *Genome sequencing:*

The following sequence data used for this study (HG002 Data Freeze v1.0) are available in a github link: [https://github.com/human-pangenomics/HG002\\_Data\\_Freeze\\_v1.0](https://github.com/human-pangenomics/HG002_Data_Freeze_v1.0). DNA samples were extracted from large homogenized growths of B-lymphoblastoid cell lines of HG002, HG003, and HG004 from the Coriell Institute for Medical Research.

#### *Illumina reads.* Paired-end reads

Whole genome data, TruSeq (LT) libraries, 300x PCR-free paired end 150bp + 40x, PCR-free paired end 250bp on Illumina HiSeq 2500, from GIAB<sup>22</sup>. HG002 was sequenced to 51.7x coverage, HG003 to 69.1x, and HG004 to 70.6x.

*Long-molecule linked reads.* 10X Genomics reads: Chromium Genome Platform from 10X Genomics was sequenced to two depths: 51.7x coverage, and a deeper 84.4x coverage (300Gb) dataset. Additional data is available from BioProject: PRJNA527321. TELL-seq linked reads: These reads were made available from another study<sup>53</sup>.

*PacBio reads.* DNA was sheared to ~20 kb with a Megaruptor 3, libraries were prepared with SMRTbell Express Template Prep Kit 2.0 and size-selected with SageELF to the targeted size (15 kb, 19 kb, 20 kb, or 25 kb), and sequenced on the Sequel II System with Chemistry 2.0 (15 kb or 20 kb libraries; 36x and 16x coverage, respectively), pre-2.0 Early Access Chemistry (15 kb, 19 kb, and 25 kb libraries; 24x, 14x and 11x coverage, respectively), and Sequel System with Chemistry 3.0 (15 kb libraries; 28x coverage). PacBio Continuous Long Reads (CLR): Libraries were prepared with SMRTbell Express Template Prep Kit 2.0, size selected to a target size (>30 kb), and sequenced on a Sequel II System with Chemistry 1.0 and Chemistry 2.0 to >60X fold coverage from 2 SMRT Cells.

*ONT reads.* ONT r9 unsheared PromethION; ONT ultralong GridION; Oxford Nanopore Unsheared PromethION data: Totaling 658x coverage of unsheared sequencing (28 PromethION Flow cells). With a total of ~51x 100kb+ reads. Oxford Nanopore (Ultra-long GridION data): Totaling ~52x coverage of unsheared sequencing (106 MinION Flow cells). With a total of ~15x 100kb+ reads.

*Hi-C linked reads.* Two different Hi-C data sets were made with two distinct protocols, to reach as uniform coverage across the genome as possible: Dovetail Omni-C (named Hi-C1) and Arima Genomics High Coverage Hi-C (named Hi-C2) protocols. For Hi-C1, about 100,000 cultured HG002 cells were processed for proximity ligation libraries without restriction enzymes. High coverage (69x) sequencing was done on a Nova-Seq (250bp PE). For Hi-C2, two libraries were prepared from two cell culture replicates, and sequenced with 2x150 bp and 2x250 bp Illumina reactions each. The combination of restriction enzymes represent 10 possible cut sites: ^GATC, G^ANTC, C^TNAG, T^TAA. The "^" is the cut site on the + DNA strand, and the 'N' can be any of the 4 genomic bases.

*Strand-seq.* Strand-specific libraries were generated as previously described (Ebert et al. 2021) from 192 barcoded single-cells and sequenced on a NextSeq Illumina platform. The 192 barcoded single-cell libraries were pooled for sequencing of the HG002 sample. Raw demultiplexed fastq files from the paired-end sequencing run (80 bp read lengths) were uploaded for each single cell library. These data can be found at

[https://s3-us-west-2.amazonaws.com/human-pangenomics/index.html?prefix=HG002/hpp\\_HG002\\_NA24385\\_son\\_v1/Strand\\_seq/](https://s3-us-west-2.amazonaws.com/human-pangenomics/index.html?prefix=HG002/hpp_HG002_NA24385_son_v1/Strand_seq/).

*Optical maps.* BioNano DLE1 data was collected (Molecules >150 kb): 317 Gbp Read N50 (Molecules >150 kb): 323 kb provided by BioNano Genomics and the Genome in a Bottle (GIAB) Consortium

### Genome assembly pipelines tested:

The assembly bakeoff was an open public science experiment and evaluation, where researchers of the HRPC and anyone in the scientific community could contribute, with the goal of creating the highest quality *de novo* assembly possible, of one or both haplotypes, using an automated process. We did this by contacting known assembly experts, sending out announcements on consortium email list (e.g. HRPC, VGP, T2T, HGSVC, and GIAB), and announcements on HRPC associated websites (<https://humanpangenome.org/hg002/>; [https://github.com/human-pangenomics/HG002\\_Data\\_Freeze\\_v1.0](https://github.com/human-pangenomics/HG002_Data_Freeze_v1.0)). We grouped the assembly pipelines tested into five categories according to whether contigs only or contigs and scaffolds were generated, and whether the contigs and/or scaffolds were haplotype phased or merged as a pseudohaplotype (**Table 1**). The assemblies were further classified by whether parental trio data was used, and whether they were reference based or *de novo* (**Table 1**). The assemblies were assigned ID numbers based on the order received by the consortium evaluation team, and in no part reflect order of assembly metric quality. All but two assemblies (asm3, asm23) that used PacBio data used the recommended downsampled HiFi SMRT cell runs from the 15 kb and 20 kb insert libraries totalling ~34X coverage ([https://github.com/human-pangenomics/HG002\\_Data\\_Freeze\\_v1.0#hg002-data-freeze-v10-recommended-downsampled-data-mix](https://github.com/human-pangenomics/HG002_Data_Freeze_v1.0#hg002-data-freeze-v10-recommended-downsampled-data-mix)). Asm3 used the 19 kb, 20 kb, and 25 kb insert libraries. Asm23 used Pacbio CLR. The specific method details for each assembly pipeline, under each of the five major categories, are described below.

### Diploid scaffold assemblies

#### Trio binning phasing VGP pipeline 1.6 (asm23):

This assembly was based on a modified version of the VGP trio pipeline 1.6<sup>8</sup>. All data types (Pacbio CLR, 10XG linked-reads, Bionano maps, and Hi-C2 reads) were haplotype-binned or filtered by haplotype. Briefly, CLR reads were binned (hapUnknownFraction=0.01) and assembled into contigs using HiCanu<sup>14</sup> v1.8. NA24143 (maternal HG004) and NA24149 (paternal HG003) 250bp PE Illumina reads were used for binning. CLR coverage of the child (HG002) was 74x and 72x for the maternal and paternal haplotypes, respectively. To polish the contigs, the binned CLR reads were used for each respective haplotype with Arrow (variantCaller v2.3.3). The two haplotype contigs were then purged from each other using purge\_dups v1.0 (PMID: 31971576), conducted in the haploid mode (calcuts -d1) and only JUNK and OVLP were removed. To these contigs, Bionano molecules were aligned and assigned to the haplotype bin with higher alignment confidence. Bionano molecules aligning equally well to both parental haplotype contigs (alignment score discrepancies of less than  $10^{-2}$ ) were randomly split into two clusters equally and assigned to the bins. The same method of splitting the molecules was used for molecules aligning to neither of the parental assemblies ([https://github.com/jlee-BNG/Bionano\\_Trio\\_binning](https://github.com/jlee-BNG/Bionano_Trio_binning) - currently private). Binned Bionano molecules were then assembled to haploid assemblies. Cross-checking was then performed by aligning the paternal and maternal Bionano assemblies to the parental assemblies to identify regions where both parents shared the same allele, and the best allele was picked for the next round of trio binning and assemblies. 10XG and Hi-C reads were filtered for k-mers of the alternate haplotype using Meryl (<https://github.com/marbl/meryl/tree/master/src/meryl>) and a custom script part of the VGP trio pipeline 1.6 was used to exclude read pairs containing k-mers only found in the other haplotype. With this prepared data, three rounds of Scaffolding were then conducted on each haplotype, sequentially with the binned 10XG reads using Scaff10x v4.2, binned Bionano maps with Solve v3.4, and binned Hi-C linked reads with Salsa v2.2. The assemblies were not further polished since they already reached Q40 as judged by Merqury.

Compute time was not tracked. The Source code is available at: <https://github.com/VGP/vgp-assembly/tree/master/pipeline>.

#### DipAsm contig and scaffolding pipeline (asm10):

This assembly is based on a protocol similar to DipAsm reported in<sup>26</sup>. Pacbio HiFi reads were first assembled into unphased contigs using Peregrine. Contigs were grouped and ordered into scaffolds with Hi-C2 data. The HiFi reads were then mapped to scaffolds using minimap2 and heterozygous SNPs called using DeepVariant<sup>54</sup>. The heterozygous SNP calls were phased with both HiFi and Hi-C2 data using HapCUT2<sup>55</sup> and WhatsHap<sup>56</sup>. The reads were then partitioned based on their phase using a custom script. The partitioned reads were re-assembled into phased contigs using Peregrine. The contigs were then ordered and joined together with 100 Ns to produce phased scaffolds. Compute time was not tracked. The source code is available at: <https://github.com/shilpagarg/DipAsm>

#### Dovetail DipAsm variant pipeline (asm2 and asm22):

The Dovetail pipeline used is a variation of the DipAsm pipeline described in<sup>26</sup>. The main difference is that DipAsm used HiFi reads for SNP calling with DeepVariant and the Dovetail protocol used Omni-C reads (Hi-C1) for SNP calling with FreeBayes. In particular, Pacbio HiFi reads were assembled into contigs using the Peregrine assembler with default parameters. These contigs were joined into chromosome-scale scaffolds using Dovetail Omni-C data and either HiRise (Dovetail Genomics; asm2) or Salsa2<sup>15</sup> (asm22) scaffolders. Omni-C reads were then aligned to scaffolds and haplotype SNPs were called using FreeBayes. These SNPs were then phased with HapCUT2 and Omni-C long-range links to obtain chromosome-scale phased blocks. These phased SNPs were used to partition HiFi and Omni-C reads into two haplotypes. Reads for which the partitioning could not be done ambiguously were assigned to both haplotypes. Phased HiFi reads for each haplotype were assembled again with Peregrine and scaffolded with haplotype-specific Omni-C reads to obtain chromosome-scale phased scaffolds. Compute time was not tracked; all the tools were run on AWS EC2 with c5d.9xlarge instance type. Source code for HiRise is proprietary. Source for SALSA2 is available at: <https://github.com/marbl/SALSA/commit/974589f3302b773dcf0f20c3332fe9daf009fb93>

#### PGAS pipeline (asm14):

The recent Phased Genome Assembly using Strand-seq (PGAS) diploid genome assembly pipeline is described in<sup>52</sup>. First, a non-haplotype resolved (“squashed”) contig assembly was generated from PacBio HiFi reads using Peregrine v0.1.5.5 ([github.com/cschin/Peregrine](https://github.com/cschin/Peregrine)). Illumina short reads from the Strand-seq data<sup>37</sup> were aligned against this squashed assembly to identify contigs that most likely originate from the same chromosome based on similar Watson/Crick strand inheritance patterns<sup>57</sup>. This information was then used to cluster the contigs into roughly chromosome-scale clusters, which helps avoid chimeric chromosome assemblies, allows for parallelization of the assembly pipeline, and facilitates phasing. Next, heterozygous single nucleotide variants were identified based on long-read alignments against the clustered assembly with DeepVariant v0.9.0. To obtain chromosome-scale haplotypes, integrative phasing with WhatsHap<sup>56</sup> was performed, combining local dense phase information derived from long reads with global sparse phase information inferred from Strand-seq alignments. Next phased heterozygous SNVs were used to assign each HiFi read to its corresponding haplotype (“haplo-tagging”) or remain in the fraction of haplotype-unassigned reads. The haplotags were used to split the HiFi reads into two haploid read sets, which, together with the haplotype-unassigned reads, were the input to assemble two haplotype contig sequences per chromosome-scale cluster with Peregrine v0.1.5.5. After polishing the contigs for two rounds with Racon v1.4.10<sup>58</sup> and the haploid long-read datasets, the per chromosome cluster assemblies were merged to create a genome-scale diploid assembly. The final round of scaffolding of each haplotype was performed with the short reads from the Strand-seq data, on HiFi contigs with a minimum size of 500 kb. This size thresholding was necessary since the contig order can only be inferred from strand-state changes resulting from sister chromatid exchanges (SCE; a process during DNA replication where two sister

chromatids break, rejoin, and physically exchange regions of the parental strands). SCEs are low-frequency events that are thus less likely to produce a traceable signal with decreasing contig size. The complete assembly pipeline run required less than 2000 CPU hours on a 3-node cluster (3 x 36C, 1.4 TB RAM) with a peak RAM usage of around 600 GB (squashed Peregrine assembly). Source code is available at: [github.com/ptrebert/project-diploid-assembly](https://github.com/ptrebert/project-diploid-assembly); pipeline parameter version 8

#### CrossStitch (asm17):

The assembly was produced using CrossStitch, a reference-based pipeline for diploid genome assembly. SNPs and small indels were called with respect to GRCh38 for HG002 from alignments of unbinned 30x PacBio HiFi reads. Variant calling was performed on this BAM using DeepVariant v0.9<sup>54</sup> and the PacBio model. A full set of commands and parameters are available on the PacBio case study: <https://github.com/google/deepvariant/blob/r0.9/docs/deepvariant-pacbio-model-case-study.md>. Larger structural variants (SV) were called by running Sniffles v1.0.11 (with parameters -s 10 -l 10 --min\_homo\_af 0.7) on minimap2 v2.17 alignments of the HiFi reads and refining these calls with Iris v1.0.1 (<https://github.com/mkirsche/Iris>). Then, the SNVs and small indel variants (less than 30 bp) called from DeepVariant were phased using HapCUT2 v1.1 on the ONT + Hi-C alignments, and these phase blocks were used to assign a haplotype to each HiFi read. SV phasing was performed by observing the reads supporting each heterozygous SV call and assigning the variant to the haplotype which the majority of the reads came from. Finally, vcf2diploid (<https://github.com/abyzovlab/vcf2diploid>) from the AlleleSeq algorithm<sup>59</sup> was used to incorporate small variant and SV calls into a template consisting of the GRCh38 reference genome sequence, producing the final assembly for HG002. The end-to-end assembly took less than two days on a high memory machine at JHU using at most 40 cores at a time. Peak RAM utilization was <100GB

Source code is available at: <https://github.com/schatzlab/crossstitch> (commit e49527b)

#### ***Diploid contig assemblies***

##### Trio binning flye ONT pipeline (asm6 and asm7):

Following a trio based assembly approach<sup>19</sup>, Parental Illumina 21-mers were counted in the child, maternal, and paternal read sets (full sets, not subset coverage recommendations). Haplotype-specific mers were created using Merqury v1.0<sup>60</sup> and Meryl v1.0 (<https://github.com/marbl/meryl>) with the command: hapmers.sh.sh mat.k21.meryl pat.k21.meryl child.k21.meryl. These short reads were then used to bin ONT standard long (asm6) or ultralong > 100kb (asm7) reads into their maternal and paternal specific haplotypes. The ONT recommended subset reads were then assigned using splitHaplotigs in Canu v2.0<sup>14</sup> with the command: splitHaplotype -cl 1000 -memory 32 -threads 28 -R HG002\_ucsc\_ONT\_1t100kb.fastq.gz \ -R HG002\_giab\_ULfastqs\_guppy3.2.4.fastq.gz \ -H ./0-kmers/haplotype-DAD.meryl 6 ./haplotype-DAD.fasta.gz \ -H ./0-kmers/haplotype-MOM.meryl 6 ./haplotype-MOM.fasta.gz \ -A ./haplotype-unknown.fasta.gz.

Flye v2.7-b1585<sup>27</sup> was then run on the binned reads to generate maternal and paternal contigs with the command fly --threads 128 --min-overlap 10000 --asm-coverage 40 -out\_dir <MOM/DAD> --genome-size 3.1g --nano-raw haplotype-<DAD/MOM>.fasta.gz. Flye sometimes inserts gaps when it isn't certain of a repeat sequence, and thus some contigs appear as scaffolds. But the assembly is still contig-level. No base-level polishing (with short or long reads) was conducted on the assembly. The ONT standard Flye runs took approximately 1200 CPU hrs (20 wall clock) and 500 Gb of memory. The ONT ultralong assemblies took approximately 3000 CPU hrs (60 wall clock) and 800 Gb of memory. Source codes for Canu, Mercury, and Flye are available at: <https://github.com/marbl/canu>; <https://github.com/marbl/merqury>; <https://github.com/fenderglass/Flye>

##### Trio binning HiCanu contig pipeline (asm19):



Following a trio based assembly approach<sup>19</sup>, parental Illumina 21-mers were counted in the child, maternal, and paternal read sets (full sets, not subset coverage recommendations). Haplotype-specific mers were created using Merqury v1.0<sup>60</sup> and Meryl v1.0 (<https://github.com/marbl/meryl>) with the command: hapmers.sh mat.k21.meryl pat.k21.meryl child.k21.meryl. The HiFi recommended 34x subset reads were then assigned to using splitHaplotigs in Canu v2.0<sup>14</sup> with the command: splitHaplotype -cl 1000 -memory 32 -threads 28 -R m64012\_190920\_173625.fastq.gz -R m64012\_190921\_234837.fastq.gz -R m64011\_190830\_220126.Q20.fastq.gz -R m64011\_190901\_095311.Q20.fastq.gz -H ./0-kmers/haplotype-DAD.meryl 6 ./haplotype-DAD.fasta.gz -H ./0-kmers/haplotype-MOM.meryl 6 ./haplotype-MOM.fasta.gz -A ./haplotype-unknown.fasta.gz. Any reads which were unclassified were randomly divided into two bins. The resulting maternal and paternal read sets were independently assembled with HiCanu<sup>14</sup> v2.0 with the commands:

```
canu -p 'asm' 'gridOptions=--time=4:00:00 --partition=quick,norm' 'gridOptionsCns=--time=30:00:00 --partition=norm' 'genomeSize=3.1g' 'gfaThreads=48' 'batOptions=-eg 0.01 -sb 0.01 -dg 6 -db 6 -dr 1 -ca 50 -cp 5' -pacbio-hifi haplotype-[DAD|MOM].fasta.gz haplotype-unknown-batch[1|2].fastq.gz. Source code: https://github.com/marbl/canu, https://github.com/marbl/merqury
```

Publication citation: <https://doi.org/10.1101/2020.03.14.992248>

All runs used the "quick" partition of the NIH Biowulf cluster (<https://hpc.nih.gov>). HiCanu required approximately 1400 CPU hours per haplotype (19 wall clock hours) and no single job required more than 64 GB of memory.

#### Trio binning Peregrine contig pipeline (asm20):

The same binned reads as for asm19 were used for this assembly. The reads were assembled with Peregrine v0.1.5.3+0.gd1eebc.dirty with the command `yes yes | python3 Peregrine/bin/pg_run.py asm \ input.list 24 24 24 24 24 24 24 24 \`

```
--with-consensus --shimmer-r 3 --best_n_ovlp 8 \
--output ./
```

with the input.list specifying appropriate haplotype input reads. Compute time was not tracked.

Source code: <https://github.com/cschin/Peregrine>, <https://github.com/marbl/merqury>

#### Trio phasing Hifiasm contig pipeline (asm9):

Hifiasm finds alignments between HiFi reads and corrects sequencing errors observed in alignments<sup>28</sup>. It labels a corrected read with its inferred parental origin using parent-specific 31-mers counted from parental short reads. HiFi reads in long homozygous regions do not harbor parent-specific 31-mers and are thus unlabeled. Hifiasm then builds a string graph from read overlaps that carries read labels. It traces paternal and maternal reads in the graph to generate paternal and maternal contigs, respectively. We collected paternal 31-mers from short reads with “yak count -b37 -o sr-pat.yak sr-pat.fq.gz” (and similarly for maternal) and assembled HiFi reads with “hifiasm -1 sr-pat.yak -2 sr-mat.yak hifi-reads.fq.gz”. The assembly took 305 CPU hours. Source code is available at: <https://github.com/chhy123/hifiasm/releases/tag/v0.3>.

DipAsm contig pipeline (asm11): The assembly pipeline mimics the DipAsm steps explained for asm10. The pipeline takes as input HiFi and Hi-C datasets and outputs the phased contigs. Initially, the pipeline produces unphased contigs using Peregrine and then these unphased contigs are scaffolded to produce chromosome-scale sequences using HiRise. Afterwards, the heterozygous SNPs are called and are phased using HiFi and Hi-C data. These phased SNPs are informative sites to partition HiFi reads to haplotypes on the chromosome level. The phased reads are then assembled using Peregrine to produce phased contigs.

#### Peregrine contig pipeline (asm3 and asm4):

The Peregrine assembler<sup>31</sup> was used to generate contigs on the HiFi reads, using either the full coverage sequence (asm3) consisting of 19Kb, 20kb, and longer 25Kb read libraries or downsampled to 34X and

shorter 15 kb reads (asm 4). A module was written to separate likely true variant sites from differences between reads caused by sequencing errors. This was done by using the overlap data from the peregrine assembler overlapping modules with additional alignment analysis. The variants of the read overlapped data were analyzed to get a subset of variants that should belong to the same haplotypes. The reads with the same set of variants were considered to be haplotype consistent and the overlaps between those haplotype consistent reads were considered for constructing the contig assembly. Overlaps between different haplotypes or different repeats from the analysis results were ignored. It is expected that the generated contigs are from single haplotypes in those regions, which have enough heterozygous variants. Compute time was not tracked. Source code is available at: [https://github.com/cschin/Peregrine\\_dev/commit/93d416707edf257c4bcb29b9693c3fda25d97a29](https://github.com/cschin/Peregrine_dev/commit/93d416707edf257c4bcb29b9693c3fda25d97a29)

### ***Partial diploid contig assemblies***

#### **FALCON-Unzip contig pipeline (asm16):**

FALCON-Unzip<sup>18</sup> version 2.2.4-py37hed50d52\_0 was run on reads from 4 SMRT Cells from two HiFi libraries (15 kb and 20 kb, 34x coverage total) reads with "input\_type = preads, length\_cutoff\_pr = 8000, ovlp\_daligner\_option = -k24 -h1024 -e.98 -l1500 -s100, ovlp\_HPCdaligner\_option = -v -B128 -M24, ovlp\_DBSplit\_option = -s400, overlap\_filtering\_setting = --max-diff 200 --max-cov 200 --min-cov 2 --n-core 24 --min-idx 98 --ignore-indels" for the initial contig assembly and default parameters for unzipping haplotypes. The assembly took 2,540 CPU-core hours on nodes with Intel Xeon processor E5-2600 v4. Source code is available at: [https://anaconda.org/bioconda/pb-falcon/2.2.4/download/linux-64/pb-falcon-2.2.4-py37hed50d52\\_0.tar.bz2](https://anaconda.org/bioconda/pb-falcon/2.2.4/download/linux-64/pb-falcon-2.2.4-py37hed50d52_0.tar.bz2)

#### **HiCanu purge dups contig phasing pipeline (asm8):**

HiCanu<sup>14</sup> v2.0 with the command `canu -p 'asm' 'gridOptions=--time=4:00:00 --partition=quick,norm' 'gridOptionsCns=--time=30:00:00 --partition=norm ' 'genomeSize=3.1g' 'gfaThreads=48' -pacbio-hifi m64012_190920_173625.fastq.gz m64012_190921_234837.fastq.gz m64011_190830_220126.Q20.fastq.gz m64011_190901_095311.Q20.fastq.gz` Purge dups (PMID: 31971576) was used to remove alternate haplotypes, GitHub commit b5ce3276773608c7fb4978a24ab29fdd0d65f1b5, with the thresholds of 5 7 11 30 22 42. Purge dups introduces gaps near the purged sequenced regions of the contigs, and thus some contigs appear as scaffolds. But the assembly is still contig-level. HiCanu required approximately 1800 CPU hours and no single job required more than 64 GB of memory (22 wall clock hours). Purge dups required 40 CPU hours and <1GB of memory.

### ***Merged haploid contig and scaffold assemblies***

#### **Flye ONT pipeline (asm5):**

Flye v2.7b-b1579<sup>27</sup> was used to assemble (downsampled) ONT reads into contigs, using the default parameters with extra "--asm-coverage 50 --min-overlap 10000" options. Two iterations of the Flye polishing module were applied using ONT reads, followed by two polishing iterations using HiFi reads. Finally, Flye graph-based scaffolding module was run on the polished contigs, which generated 54 scaffold connections and slightly improved the assembly contiguity. Assembly took ~5000 CPU hours and polishing (ONT + HiFi) took ~3000 CPU hours. Peak RAM usage was ~900 Gb. The pipeline was run on a single computational node with two Intel Xeon 8164 CPUs, with 26 cores each and 1.5 TB of RAM. Source code is available at: <https://github.com/fenderglass/Flye/>. Commit id: ec206f8

#### **Shasta ONT + HiC (asm18 and asm21):**

The Shasta assembler<sup>25</sup> was used to assemble ONT reads into contigs. The contigs were polished using PEPPER (<https://github.com/kishwarshafin/pepper>) which also uses only the ONT reads. The contigs were scaffolded with Omni Hi-C (Hi-C1) using HiRise (asm18) or Salsa v2.0 (asm21). Compute time was not tracked. Source code is available at: <https://github.com/chanzuckerberg/shasta>

### Flye and MaSuRCA (asm15):

A subset of downsampled ONT UL data that contained ~38x genome coverage of 120kb reads or longer was used. The ONT reads were assembled into contigs using the Flye assembler<sup>27</sup> v2.5. The contigs were polished with downsampled 30x coverage of PacBio HiFi 15kb and 20kb reads, using POLCA, a tool distributed with the Maryland Super-Read Celera Assembler (MaSuRCA)<sup>29</sup>. To scaffold and assign the assembled contigs to chromosomes, a reference-based scaffolding method was used embodied in the chromosome\_scaffolder script included in MaSuRCA. GRCh38.p12 was used as a reference (without the ALT scaffolds) for scaffolding, with chromosome\_scaffolder option enabled, which allows it to fill in the gaps in scaffolds, where possible, with GRCh38 sequence, in lowercase letters. The final assembly was named JHU\_HG002\_v0.1. Compute time was not tracked. Source code is available at: <https://github.com/alekseyzimin/masurca>

### MaSuRCA (asm1):

MaSuRCA v3.3.1<sup>29</sup> with default parameters was run on the combined Illumina, ONT, and PacBio HiFi data to obtain a set of contigs designated the Ash1 v0.5 assembly. The ONT and PacBio data were an earlier release, from 2018, and the read lengths were shorter than the later release used by most other methods in this evaluation. After initial scaffolding, MaSuRCA was used to remove redundant haplotype-variant scaffolds by aligning the assembly to itself and looking for scaffolds that were completely covered by another larger scaffold and that were >97% identical to the larger scaffold. To scaffold and assign the assembled contigs to chromosomes, we used a reference-based scaffolding method embodied in the chromosome\_scaffolder script included in MaSuRCA. We used the GRCh38.p12 as the reference (without the ALT scaffolds), and we enabled an option in chromosome\_scaffolder that allows it to fill in gaps with GRCh38 sequence, using lowercase letters. Finally, we examined SNVs reported at high frequency in an Ashkenazi population from the Genome Aggregation Database (gnomAD). GnomAD v3.0 contains SNV calls from short-read whole-genome data from 1,662 Ashkenazi individuals. At 273,866 heterozygous SNV sites where HG002 contained the Ashkenazi major allele and where our assembly used a minor allele, we replaced the allele in Ash1 with the Ashkenazi major allele. A publication of the final curated asm1 assembly has been made<sup>61</sup>. Compute time was not tracked. Source code is available at: <https://github.com/alekseyzimin/masurca>

### ***Merged haploid contig assemblies***

#### wtdbg2 (asm13)

The standard wtdbg2 assembly pipeline<sup>32</sup> was applied on HiFi-reads. Parameters '-k 23 -p 0 -S 0.8 --no-read-clip --aln-dovetail -l' were customized to improve the contiguity. Source code is available at: <https://github.com/ruanjue/wtdbg2> commit:d6667e78bbde00232ff25d3b6f16964cc7639378

#!/bin/bash

```
wtdbg2 -k 23 -p 0 -AS 4 -s 0.8 -g 3g -t 96 --no-read-clip --aln-dovetail -l -fo dbg -i  
../rawdata/SRR10382244.
```

```
fasta -i ../rawdata/SRR10382245.fasta -i ../rawdata/SRR10382248.fasta -i ../rawdata/SRR10382249.fasta
```

```
wtpoa-cns -t 96 -i dbg.ctg.lay.gz -fo dbg.raw.fa
```

```
minimap2 -I64G -ax asm20 -t96 -r2k dbg.raw.fa ../rawdata/SRR10382244.fasta  
../rawdata/SRR10382245.fasta
```

```
../rawdata/SRR10382248.fasta ../rawdata/SRR10382249.fasta | samtools sort -m 2g -@96 -o dbg.bam
```

```
samtools view -F0x900 dbg.bam | wtpoa-cns -t 96 -d dbg.raw.fa -i - -fo dbg.cns.fa
```

Publication citation: 10.1101/530972

Compute time ?????.

wtdbg2: real 20349.731 sec, user 1178897.390 sec, sys 18351.090 sec, maxrss 194403704.0

kB, maxvsize 209814736.0 kB

wtpoa-cns(1): real 3350.517 sec, user 260551.730 sec, sys 1040.200 sec, maxrss 9978492.0

kB, maxvsize 15839032.0 kB  
wtpoa-cns(2):real 2181.084 sec, user 149528.810 sec, sys 815.380 sec, maxrss 11134244.0 kB,  
maxvsize 16012012.0 kB  
others: unknown

#### NECAT Feng Luo group (asm12):

We used the NECAT assembler<sup>30</sup> to assemble ONT reads of HG002, which contained about 53x coverage excluding ONT UL reads. The command "necat.pl config cfg" was first used to generate the parameter file 'cfg'. The default values in 'cfg' were replaced with the following parameters: "GENOME\_SIZE=3000000000, THREADS=64, PREP\_OUTPUT\_COVERAGE=40, OVLP\_FAST\_OPTIONS=-n 500 -z 20 -b 2000 -e 0.5 -j 0 -u 1 -a 1000", OVLP\_SENSITIVE\_OPTIONS=-n 500 -z 10 -e 0.5 -j 0 -u 1 -a 1000, CNS\_FAST\_OPTIONS=-a 2000 -x 4 -y 12 -l 1000 -e 0.5 -p 0.8 -u 0, CNS\_SENSITIVE\_OPTIONS=-a 2000 -x 4 -y 12 -l 1000 -e 0.5 -p 0.8 -u 0, TRIM\_OVLP\_OPTIONS=-n 100 -z 10 -b 2000 -e 0.5 -j 1 -u 1 -a 400, ASM\_OVLP\_OPTIONS=-n 100 -z 10 -b 2000 -e 0.5 -j 1 -u 0 -a 400, CNS\_OUTPUT\_COVERAGE=40". The command "necat.pl bridge cfg" was run to generate the final contigs. It took ~12555 CPU hours (error correction 2500 h, assembling 8123 h, bridging 1216 h, polishing 716 h) on a 4-core 24-thread Intel(R) Xeon(R) 2.4 GHz CPU (CPU E7-8894[v4]) machine with 3 TB of RAM. Source code is available at: <https://github.com/xiaochuanle/NECAT>. Commit id: 47c6c23.

#### **HPRC-HG002 references:**

HiFi reads with adaptors were removed with HiFiAdapterFilt (<https://github.com/sheinasim/HiFiAdapterFilt>). Maternal and paternal contigs were generated from 160X coverage of the remaining HiFi reads using hifiasm v0.14.1 in trio mode. The maternal and paternal HiFi contigs were then separately scaffolded with paternal and maternal Trio-Bionano Solve v1.6 maps, respectively. Conflicts between the Bionano maps and Pacbio HiFi-based contigs were manually reviewed by three experts and decisions were made to accept or reject the cuts proposed by Bionano Solve. Further scaffolding of the paternal and maternal scaffolds was done with Salsa v2.3 and OmniC Hi-C reads excluding reads from the other haplotype with Meryl; that is, paired reads with *k*-mers only seen in one parent were removed prior to mapping and scaffolding. The resulting scaffolds were manually curated to ensure the proper order and orientation of contigs within the scaffolds, leading to additional joins and breaks. In parallel, ONT-UL reads were haplotype-binned with Canu v2.1 using Illumina short reads of the parents, and then were assembled into their respective maternal and paternal contigs using Flye v2.8.3-b1695 with --min-overlap 10000. Bases were recalled with Guppy v4.2.2 prior to the assembly. The contigs were polished calling variants with Medaka (<https://github.com/nanoporetech/medaka>). The variants were filtered with Merfin using *k*-mers derived from Illumina short reads. Bcftools was used to apply the variants. The resulting ONT-UL assembly was used to patch gaps of the scaffolded HiFi-based assembly, using custom scripts ([https://github.com/gf777/misc/tree/master/HPRC%20HG002/for\\_filling](https://github.com/gf777/misc/tree/master/HPRC%20HG002/for_filling)). Finally, a decontamination and an additional round of manual curation were conducted.

#### **Trio binning with optical mapping:**

Using the Bionano Direct Label and Stain (DLS) chemistry and the Saphyr machine, high coverage of Bionano optical maps were generated of the HG002 (son), HG003 (father), and HG003 (mother) trio of samples and each assembled into diploid assemblies (Table [xx - molecule and assembly metrics](#)) with Bionano Solve 3.6. To separate the paternal and maternal alleles in the child assembly, the child molecules were aligned to the father and mother assemblies and assigned to the bin with higher alignment confidence.

Molecules that aligned equally well (alignment score difference  $\leq 10^{-2}$ ) to the parents were split into two clusters equally and assigned to the bins. Similarly, molecules that aligned to neither of the parents were split into two clusters equally and assigned to the bins. This resulted in **xx**G of the child molecules binned to the paternal allele and **xx**G binned to the maternal allele. These binned molecules were then assembled into the paternal haplotype (3.00 Gbp in length with N50 of 70.36 Mbp) and maternal haplotype (2.98 Gbp in length with N50 of 70.40 Mbp) using Bionano Solve 3.6.

The above trio-binning method separates the paternal and maternal alleles in most loci except for where the parents share an allele. To further improve the separation of the parental alleles, a cross-checking step is performed. In the cross-checking step, the binned paternal and maternal assemblies are aligned to both the father and mother assemblies (using `cross_check_alignment.py` with `RefAligner 11741` and `optArguments_customized_for_diploid_reference.xml`). Using these alignments, regions where the parents shared an allele but were homozygous in one and heterozygous in the other were identified, and the corresponding coordinates of these regions in the binned assemblies were identified and then broken (using `haplotype_segregation_cross_check_rscript.R`, `haplotype_segregation_cut_step.py`). Breaking at these regions allowed further separation of alleles in the next round of trio binning using the binned and cross-checked assemblies as anchors. In the cross-checking step, **xx** breaks are performed on the paternal binned assembly and **xx** breaks in the maternal binned assembly. For the trio-binning after cross-checking, **xx**G of the probe molecules were binned to the paternal haplotype and **xx**G binned to the maternal haplotype. These binned molecules were then assembled into the paternal haploid assembly (2.98 Gbp with N50 of 79.22 Mbp) and maternal haploid assembly (2.96 G with N50 of 66.60 Mbp) using Bionano Solve 3.6.

All molecules (original and binned) and assemblies can be found in [https://bionanowithcentrexit-my.sharepoint.com/:f:/g/personal/jlee\\_bionanogenomics\\_com/Eq9gz0in15hKnC5CMrhUYkwBmjY0uiC\\_gPhu5QDR82AThw?e=u5Eyzd](https://bionanowithcentrexit-my.sharepoint.com/:f:/g/personal/jlee_bionanogenomics_com/Eq9gz0in15hKnC5CMrhUYkwBmjY0uiC_gPhu5QDR82AThw?e=u5Eyzd)

All trio binning and cross-checking scripts can be found in [https://github.com/andypang0714/Bionano\\_Trio\\_binning v2.0](https://github.com/andypang0714/Bionano_Trio_binning_v2.0).

## Evaluation methods:

**Contamination, manual curation** (missed-joins, miss-joins, inversion, order errors, chromosome assignments). Sex chromosomes and mitochondria.

### *Continuity metrics*

Assembly continuity statistics were collected using `asm_stats.sh` from the VGP pipeline [<https://github.com/VGP/vgp-assembly/tree/master/pipeline/stats>, REF Rhie et al, Nature], using genome size of 3Gb for calculating NG50 values. All N-bases were considered as gaps.

### *Completeness, phasing, and base call accuracies.*

We collected 21-mers from Illumina reads of HG002 (250bp paired-end) and the parental genomes (HG003 and HG004) using Meryl [REF Rhie et al, Genome Biology], and used Merqury [REF Rhie et al, Genome Biology] to calculate QV, completeness and phasing statistics. Like continuity metrics, phase block NG50 was obtained using a genome size of 3Gb. False duplications were post-calculated using `false_duplications.sh` in Merqury and spectra-cn histogram files for each haploid-representation of the assemblies.

### *Collapsed analyses*



We calculated collapsed and expandable sequences using previously described methods<sup>62</sup>. In brief we aligned downsampled HiFi reads from HG002 independently to each assembly of HG002 and defined collapsed bases as regions in the assembly with greater than expected coverage (mean + at least three standard deviations) that were at least 15kb in length. We performed analyses with common repeats collapses included and excluded, defining common repeats collapses as sequences that were over 75% common repeat elements as identified by RepeatMasker (v4.1.0) and TRF (v4.09). This filter removed many collapses corresponding to alpha satellite and human satellite to get a better estimate of collapsed segmental duplications. Further, we defined expandable Mbp as the estimate of how much sequence would be in the collapsed regions had they been correctly assembled. We estimated this by multiplying the length of each collapse against the read depth divided by the average genome coverage. Code used for this analysis is available at <https://github.com/mrvollger/SDA> (commit 23fa175).

### ***Strand-seq analyses***

In order to evaluate structural accuracy of each assembly, we first aligned Strand-seq data from HG002 to each assembly using BWA-MEM (version 0.7.15-r1140) REF ([Li and Durbin 2010](#)) with the default parameters. Subsequently, all secondary and supplementary alignments were removed using SAMtools (version 1.9) REF ([Li et al. 2009](#)) and duplicate reads were marked using Sambamba (version 0.6.8) REF ([Tarasov et al. 2015](#)). Duplicated reads and reads with mapping quality less than 10 were removed prior to subsequent Strand-seq data analysis. To evaluate structural and directional contiguity of each assembly, we used R package SaaRclust REF ([Porubsky et al. 2021](#)) with the following parameters: bin.size = 200000; step.size = 200000; prob.th=0.25; bin.method = 'fixed'; min.contig.size = 100000; min.region.to.order = 500000; ord.method = 'greedy'; num.clusters = 100; remove.always.WC = TRUE; mask.regions = FALSE; and max.cluster.length.mbp = 300. SaaRclust automatically reports contigs that likely contain a misassembly and marks them as either misorientation (change in directionality of a piece of contig) or chimerism (regions of a contig that originate from different chromosomes). To reduce false positive calls, we report only misoriented and chimeric regions that are at least 400 kb and 1 Mb in length, respectively. Current version of the R package SaaRclust is available at <https://github.com/daewoooo/SaaRclust> (devel branch).

### ***Variation benchmark analysis***

We used v4.2.1 GIAB Benchmark variants with GA4GH, with v3.0 stratifications, which enabled comparative performance assessment inside and outside challenging genomic regions such as segmental duplications, homopolymers, tandem repeats<sup>34</sup>. Benchmarking tools from GIAB and GA4GH enable performance to be stratified by type of error (e.g., genotyping errors) and genome context (e.g., segmental duplications). Variants were first called using the dipcall assembly variant calling pipeline (<https://github.com/lh3/dipcall>)<sup>39</sup>. Dipcall first aligns an assembly to the GRCh38 reference genome ([ftp://ftp.ncbi.nlm.nih.gov/genomes/all/GCA/000/001/405/GCA\\_000001405.15\\_GRCh38/seqs\\_for\\_alignment\\_pipelines.ucsc\\_ids/GCA\\_000001405.15\\_GRCh38\\_no\\_alt\\_analysis\\_set.fna.gz](ftp://ftp.ncbi.nlm.nih.gov/genomes/all/GCA/000/001/405/GCA_000001405.15_GRCh38/seqs_for_alignment_pipelines.ucsc_ids/GCA_000001405.15_GRCh38_no_alt_analysis_set.fna.gz)) using minimap2 (<https://github.com/lh3/minimap2>)<sup>63</sup>. We used optimized alignment parameters -z200000,10000 to improve alignment contiguity, as this is known to improve variant recall in regions with dense variation, such as the MHC<sup>64</sup>. Dipcall uses the resulting alignment to generate a bed file with haplotype coverage and call variants. All filtered variants except those with the GAP2 filter were removed. GAP2 filtered variants occurred particularly in primary-alternate assemblies in homozygous regions where the alternate contig was missing. These GAP2 variants were kept as filtered to give separate performance metrics, and treated as a homozygous variant with respect to GRCh38 by changing the genotype (GT) field from 1| to 1|1. The resulting variant calls were benchmarked using hap.py v0.3.12 with the RTG Tools (v3.10.1) vcfeval comparison engine (<https://github.com/Illumina/hap.py>)<sup>40</sup>. Earlier versions of hap.py and vcfeval do not output lenient regional variant matches to the FP.al field. The hap.py comparison was performed with the v4.2.1 Genome In A Bottle HG002 small variant benchmark vcf and bed ([42](https://ftp-</a></p>
</div>
<div data-bbox=)

trace.ncbi.nlm.nih.gov/ReferenceSamples/giab/release/AshkenazimTrio/HG002\_NA24385\_son/NISTv4.2.1/GRCh38/<sup>23</sup> and V3.0 of the GIAB genome stratifications (doi:10.18434/M32190). To improve reproducibility and transparency, Snakemake (<https://snakemake.readthedocs.io/en/stable/>)<sup>65</sup> was used for pipeline construction and execution (<https://github.com/jmcdani/HPRC-assembly-benchmarking>). The extensive performance metrics output by hap.py in the extended.csv files were summarized in the following metrics for Completeness, Correctness, and Hard Regions.

Completeness metrics were calculated from SNV FN rate or Recall to assess how much of the benchmark does the callset cover, where 100% means capturing all variants and 0% means capturing none. These completeness metrics were calculated at different stringencies with SNP.Recall or as a true positive (TP). `SNP.Recall\_ignoreGT` is a measure of how well the assembly captures at least one of the variant alleles, and is considered TP if at least one allele in a variant was called correctly, regardless of whether genotype was correct. This is calculated from  $(\text{SNP.TRUTH.TP} + \text{SNP.FP.gt}) / \text{SNP.TRUTH.TOTAL}$  for the row with “ALL” in the FILTER column. `SNP.Recall` is a measure of how well the assembly represents genotypes, and counted as TP if the variant and genotype are called correctly. When only one contig is present, we assume the region is homozygous. This is calculated from METRIC.Recall for TYPE=SNP, SUBTYPE=\*, SUBSET=\*, FILTER=ALL. `SNP.Recall.fullydiploid` is a measure of how well the assembly represents both haplotypes correctly, requiring that exactly one contig from each haplotype align to the location (contigs smaller than 10kb are ignored by dipcall by default). This is calculated from METRIC.Recall for TYPE=SNP, SUBTYPE=\*, SUBSET=\*, FILTER=PASS.

Correctness metrics were calculated from the false positive rate for SNVs and INDELs, converted into a phred scaled per base error rate. Each SNP and INDEL was counted as a single error on one haplotype regardless of size and genotype. `QV\_dip\_snp\_indel` is the error rate in all benchmark regions, calculated as  $-10 * \log_{10}((\text{SNP.QUERY.FP} + \text{INDEL.QUERY.FP}) / (\text{Subset.IS\_CONF.Size} * 2))$ . `NoSegDup.QV\_dip\_snp\_indel` is the same as QV\_dip\_snp\_indel except that it excludes segmental duplication regions.

Hard Regions metrics were calculated for particularly difficult-to-assemble regions like segmental duplications. `Segdup.QV\_dip\_snp\_indel` is the same as the `QV\_dip\_snp\_indel` Correctness metric, but only for segmental duplication regions.

To benchmark SVs, we aligned the final HG002 assembly to GRCh37 and used truvari v3.1.0 to benchmark variants against the GIAB v0.6 benchmark in v0.6.2 benchmark regions.

### ***BUSCO analyses***

Busco completeness for the 41 assemblies was calculated with BUSCO v3.1.0 using the mammalia\_odb9 lineage set.

### ***Annotations***

Human RefSeq transcripts of type “known” (with NM or NR prefixes<sup>66</sup>) were queried from NCBI Entrez on December 8, 2021, and aligned to the 43 assembled haplotypes and to GRCh38 (GCF\_000001405.26) and CHM13 T2T v1.1 (GCA\_009914755.3). The coding transcripts and non-coding transcripts longer than 300 bp were first aligned with BLAST (evalue of 0.0001, word size 28 and best-hits options best\_hit\_overhang=0.1 and best\_hit\_score\_edge=0.1) to the genomes masked with RepeatMasker ([www.repeatmasker.org](http://www.repeatmasker.org))<sup>67</sup> or Windowmasker<sup>68</sup>. Sets of results obtained with both masking methods were passed to the global alignment algorithm Splign<sup>69</sup> (75% min exon identity, 50% min compartment identity and 20% min singleton identity) to refine the splice junctions and align exons missed by BLAST. Sequences for which no alignment with coverage higher than 95% of the query and sequences with unaligned overhangs at the 5’ or 3’ end were re-aligned with Blast and Splign to the unmasked genome, and then submitted to the same filter. Non-coding transcripts shorter than 300 bp were aligned with BLAST to the unmasked genome (evalue of 0.0001, word size 16, 98% identity and best-hits options best\_hit\_overhang=0.1 and best\_hit\_score\_edge=0.1) and then with Splign (75% min exon identity, min compartment identity and min singleton identity) and submitted to the same filter as the other transcripts.



The alignments for each transcript were then ranked based on identity and coverage. Transcripts that aligned best to GRCh38 sex chromosomes were filtered out prior to computing the statistics below.

Several measures for assembly completeness and correctness, and for sequence accuracy were compared across all assembly haplotypes: 1) **Unaligned genes**, either due to one or more child transcript alignment being absent or too low in sequence identity; 2) **Split genes**, across two or more scaffolds; 3) **Low coverage genes**, with less than 95% of the coding sequence in the assembly; 4) **Dropped genes**, most often due to collapsed regions in the assembly; and 5) **Genes with frameshifted CDSs**, where the best-ranking alignment requires insertions or deletions to compensate for suspected insertions or deletions in the genomic sequence that cause frameshift errors. For category 4, since each RefSeq transcript is associated with a single gene<sup>70</sup> and genes are not expected to overlap, unless explicitly known to, collapsed regions were identified as loci where transcripts from multiple genes co-aligned, and measured as the count of genes for which the best alignment of a child transcript needed to be dropped to resolve the conflict.

### Pangenomic assessment of the assemblies

We performed pairwise alignments for all chromosomes of all 45 assemblies with the wfmash sequence aligner (<https://github.com/ekg/wfmash>, commit 09e73eb), requiring homologous regions at least 300 kb long and nucleotide identity of at least 98%. We used the alignment between all assemblies to build a pangenome graph with the seqwish variation graph inducer (<https://github.com/ekg/seqwish>, commit ccfebf0), ignoring alignment matches shorter than 79 bp (to remove possible spurious relationships caused by short repeated homologies). To obtain chromosome-specific pangenome graphs, contigs were partitioned by aligning all of them with wfmash against the GRCh38 and CHM13 reference sequences. Graph statistics, visualizations, and pairwise Jaccard similarities and Euclidean distances between haplotypes were obtained with the ODGI toolkit<sup>71</sup> (commit 67a7e5b). We performed the multidimensional analyses in the R development environment (version 3.6.3), equipped with the following packages: tidyverse (version 1.3.0), RColorBrewer (version 1.1.2), ggplot2 (version 3.3.3), ggrepel (version 0.9.1), and stats (version 3.6.3). Specifically, we applied the Classical Multidimensional Scaling on the Euclidean distance matrix to perform the Principal Component Analysis (PCA). Pangenome graphs at selected *loci* were built and visualized by using the PGGB pipeline (commit 5d26011) and the ODGI toolkit. Code and links to data resources used to build the pangenome graphs, perform the multidimensional analyses, and produce all the figures can be found at the following repository: [https://github.com/AndreaGuarracino/HG002\\_assemblies\\_assessment](https://github.com/AndreaGuarracino/HG002_assemblies_assessment).

### Heterozygosity analysis

To call the full spectrum of heterozygosity between the two haploid sequences, we directly compared two haploid assemblies using Mummer (v4.0.0rc1) with the parameters of “nucmer -maxmatch -l 100 -c 500”. SNP and small indels were generated by “delta-filter -m -i 90 -l 100” and followed by “dnadiff”. Several custom scripts were used to analyze the Mummer output, as described in our prior marmoset study (<https://github.com/comery/marmoset>). We employed SyRi (v1.5, [<https://doi.org/10.1186/s13059-019-1911-0> ]) to detect SVs from Mummer alignments using default parameters. SVs in which more than half the feature consisted of gaps were dropped. For CNVs, we only kept local tandem contractions/expansions, whole-genome copy number changes were not included in these results. to avoid false positives caused by assembly issues and insufficient detection power, we only kept translocations where haplotypes reciprocally share the best alignment.

### Data and Software Availability

1. Available location of the karyotype data
2. Available locations of the raw sequence data
3. Available locations of the assemblies

## Consent

1. Informed consent obtained by the Personal Genome Project, which permits open sharing of genomic data, phenotype information, and redistribution of cell lines and derived products.

## Author Contributions

## Competing Interest

E.E.E. is a scientific advisory board (SAB) member of Variant Bio, Inc.

## Funding

The primary source of funding for this study was from NHGRI grant U01HG010971 (<https://grants.nih.gov/grants/guide/rfa-files/rfa-hg-19-004.html>). Additional funding: HHMI (E.D.J.), NHGRI grants HG002385 and HG010169 (E.E.E.) .....The computational resources and personnel support for the DipAsm assemblies were HG010906/HG/NHGRI NIH HHS and NNF21OC0069089 to SG. NCI U01CA253481 to MCS; NSF DBI-1627442 to MCS. T.M. and P.E. (asm14 team) acknowledge the support by the BMBF-funded de.NBI Cloud within the German Network for Bioinformatics Infrastructure (de.NBI) (031A532B, 031A533A, 031A533B, 031A534A, 031A535A, 031A537A, 031A537B, 031A537C, 031A537D, 031A538A), and the German Research Foundation (391137747 to T.M.). E.D.J., E.E.E. D.H. are investigators of the Howard Hughes Medical Institute. This work was supported in part by the National Center for Biotechnology Information of the National Library of Medicine (NLM), National Institutes of Health, and by National Institute of Standards and Technology intramural research funding.

## Acknowledgements

Certain commercial equipment, instruments, or materials are identified to specify adequately experimental conditions or reported results. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the equipment, instruments, or materials identified are necessarily the best available for the purpose. T.M. and P.E. (asm14 team) acknowledge the computational infrastructure and support provided by the Centre for Information and Media Technology at Heinrich Heine University Düsseldorf.

## Supplementary Material

### Supplementary Note 1: Estimation of diversity representation

Based on theoretical estimates from SNPs in the 1000G project, and using Kruglyak et al 2000 formula<sup>72</sup>, to detect 100% of minor alleles greater than 2% in the human world population or ~97% of those above 1%, one would need to generate complete sequence for ~350 diploid genomes (700 haplotypes) from diverse populations, the number thus far funded. To obtain 99% of those above 1%, one would need ~450 genomes (900 haplotypes). Minor allele frequency is the frequency at which the second most common allele occurs in a population. This calculation will differ for structural variants, estimated at 50 genomes for alleles above 5%. These values will need to be revised in an iterative fashion, as more relatively complete and error-free genomes are generated. For example, as validated in this study, haplotype diversity is uneven across the genome, e.g. where centromeres have much greater diversity than most other genic parts of the genome. Also the technology used will make a difference. These values were determined based mostly on short-read only assemblies, which do not assemble centromeres well, and also not GC-rich or repeat regions well.

### Supplementary Note 2: Additional reads with adaptors that needed to be removed.

After the standar method of removing adaptors in the sequence reads, we screened the presence of any remaining adaptors in the HG002 Pacbio libraries. We found that 0.3-0.4% of reads (~2,000) still contained at least 1 adapter per SMRT cell movie, and that adds up to a 1.7% frequency overall (all contaminated reads/all reads). Such adaptor contaminated reads were greatly reduced with chemistry 3.0 relative to 2.0 and pre2.0. The PacBio 45 bp blunt adapter was the sequence that was not properly clip in the sequence processing:

```
> gnl|uv|NGB00972.1:1-45 Pacific Biosciences Blunt Adapter  
ATCTCTCTCTTTTCCTCCTCCTCCGTTGTTGTTGTTGAGAGAGAT
```

These were due to reads with adapter 2-dimers (more prevalent in chem pre2.0 than in chem2.0). The location of the adapters were mostly (90%) at the beginning of the reads (first 1/3rd of the read length), and 87% within the first 100bp; 0.1% in the middle of the read (2/3 of the read); and the remaining 9.5% at the end of the read (last third of the read), and 9.3 % within the last 100bp. To identify reads with reads of all these categories, we used HiFiAdapterFilt (<https://github.com/sheinasim/HiFiAdapterFilt>). Some reads were entirely made of adaptors, which presumably occurs by self ligation in the library preparation. We removed all such reads with adaptors for the final HG002 assembly, even if they had human sequence attached, because using cutup versions of these reads led to lower assembly metrics relative to without them. Later we switched to CutAdapt<sup>35</sup>, as the results had about 90% overlap with HiFiAdapterFilt, but Cutadapt was about 2 times faster and slightly more sensitive (5-10%).

### References

1. Lander, E. S. *et al.* Initial sequencing and analysis of the human genome. *Nature* **409**, 860–921 (2001).
2. Schneider, V. A. *et al.* Evaluation of GRCh38 and de novo haploid genome assemblies demonstrates the enduring quality of the reference assembly. *Genome Res.* **27**, 849–864 (2017).
3. Sherman, R. M. & Salzberg, S. L. Pan-genomics in the human genome era. *Nat. Rev. Genet.* **21**, 243–254 (2020).
4. Logsdon, G. A., Vollger, M. R. & Eichler, E. E. Long-read human genome sequencing and its applications. *Nat. Rev. Genet.* **21**, 597–614 (2020).
5. Nurk, S. *et al.* *The complete sequence of a human genome.*

<http://biorxiv.org/lookup/doi/10.1101/2021.05.26.445798> (2021)

doi:10.1101/2021.05.26.445798.

6. Wang, T. *et al.* The Human Pangenome Project: A global resource to map genomic diversity. (submitted).
7. 1000 Genomes Project Consortium *et al.* A global reference for human genetic variation. *Nature* **526**, 68–74 (2015).
8. Rhie, A. *et al.* Towards complete and error-free genome assemblies of all vertebrate species. *Nature* **592**, 737–746 (2021).
9. Kim, J. *et al.* False gene and chromosome losses affected by assembly and sequence errors. *bioRxiv* 2021.04.09.438906 (2021) doi:10.1101/2021.04.09.438906.
10. Ko, B. J. *et al.* Widespread false gene gains caused by duplication errors in genome assemblies. *bioRxiv* 2021.04.09.438957 (2021) doi:10.1101/2021.04.09.438957.
11. Kelley, D. R. & Salzberg, S. L. Detection and correction of false segmental duplications caused by genome mis-assembly. *Genome Biol* **11**, R28 (2010).
12. Ebert, P. *et al.* Haplotype-resolved diverse human genomes and integrated analysis of structural variation. *Science* **372**, eabf7117 (2021).
13. Amarasinghe, S. L. *et al.* Opportunities and challenges in long-read sequencing data analysis. *Genome Biol.* **21**, 30 (2020).
14. Nurk, S. *et al.* HiCanu: accurate assembly of segmental duplications, satellites, and allelic variants from high-fidelity long reads. *Genome Res.* **30**, 1291–1305 (2020).
15. Ghurye, J. *et al.* Integrating Hi-C links with assembly graphs for chromosome-scale assembly. *PLoS Comput Biol* **15**, e1007273 (2019).

16. Durand, N. C. *et al.* Juicebox Provides a Visualization System for Hi-C Contact Maps with Unlimited Zoom. *Cell Syst* **3**, 99–101 (7).
17. Bocklandt, S., Hastie, A. & Cao, H. Bionano Genome Mapping: High-Throughput, Ultra-Long Molecule Genome Analysis System for Precision Genome Assembly and Haploid-Resolved Structural Variation Discovery. *Adv. Exp. Med. Biol.* **1129**, 97–118 (2019).
18. Chin, C. S. *et al.* Phased diploid genome assembly with single-molecule real-time sequencing. *Nat Methods* **13**, 1050–1054 (2016).
19. Koren, S. *et al.* De novo assembly of haplotype-resolved genomes with trio binning. *Nat Biotechnol* (2018) doi:10.1038/nbt.4277.
20. Kronenberg, Z. N. *et al.* Extended haplotype-phasing of long-read de novo genome assemblies using Hi-C. *Nat Commun* **12**, 1935 (4).
21. Ball, M. P. *et al.* A public resource facilitating clinical use of genomes. *Proc. Natl. Acad. Sci. U. S. A.* **109**, 11920–11927 (2012).
22. Zook, J. M. *et al.* Extensive sequencing of seven human genomes to characterize benchmark reference materials. *Sci. Data* **3**, 160025 (2016).
23. Wagner, J. *et al.* *Benchmarking challenging small variants with linked and long reads.* <http://biorxiv.org/lookup/doi/10.1101/2020.07.24.212712> (2020)  
doi:10.1101/2020.07.24.212712.
24. Chaisson, M. J. *et al.* Resolving the complexity of the human genome using single-molecule sequencing. *Nature* **517**, 608–11 (2015).
25. Shafin, K. *et al.* Nanopore sequencing and the Shasta toolkit enable efficient de novo assembly of eleven human genomes. *Nat. Biotechnol.* **38**, 1044–1053 (2020).

26. Garg, S. *et al.* Chromosome-scale, haplotype-resolved assembly of human genomes. *Nat. Biotechnol.* **39**, 309–312 (2021).
27. Kolmogorov, M., Yuan, J., Lin, Y. & Pevzner, P. A. Assembly of long, error-prone reads using repeat graphs. *Nat. Biotechnol.* **37**, 540–546 (2019).
28. Cheng, H., Concepcion, G. T., Feng, X., Zhang, H. & Li, H. Haplotype-resolved de novo assembly using phased assembly graphs with hifiasm. *Nat. Methods* **18**, 170–175 (2021).
29. Zimin, A. V. *et al.* The MaSuRCA genome assembler. *Bioinforma. Oxf. Engl.* **29**, 2669–2677 (2013).
30. Chen, Y. *et al.* Efficient assembly of nanopore reads via highly accurate and intact error correction. *Nat. Commun.* **12**, 60 (2021).
31. Chin, C.-S. & Khalak, A. *Human Genome Assembly in 100 Minutes*.  
<http://biorxiv.org/lookup/doi/10.1101/705616> (2019) doi:10.1101/705616.
32. Ruan, J. & Li, H. Fast and accurate long-read assembly with wtdbg2. *Nat. Methods* **17**, 155–158 (2020).
33. Formenti, G. *et al.* Complete vertebrate mitogenomes reveal widespread repeats and gene duplications. *Genome Biol* **22**, 120 (2021).
34. Silkaitis, K. & Lemos, B. Sex-biased chromatin and regulatory cross-talk between sex chromosomes, autosomes, and mitochondria. *Biol. Sex Differ.* **5**, 2 (2014).
35. Aganezov, S. *et al.* A complete reference genome improves analysis of human genetic variation. <http://biorxiv.org/lookup/doi/10.1101/2021.07.12.452063> (2021)  
doi:10.1101/2021.07.12.452063.
36. Howe, K. *et al.* Significantly improving the quality of genome assemblies through curation.



*Gigascience* **10**, (2021).

37. Falconer, E. *et al.* DNA template strand sequencing of single-cells maps genomic rearrangements at high resolution. *Nat. Methods* **9**, 1107–1112 (2012).
38. Sanders, A. D., Falconer, E., Hills, M., Spierings, D. C. J. & Lansdorp, P. M. Single-cell template strand sequencing by Strand-seq enables the characterization of individual homologs. *Nat. Protoc.* **12**, 1151–1176 (2017).
39. Li, H. *et al.* A synthetic-diploid benchmark for accurate variant-calling evaluation. *Nat. Methods* **15**, 595–597 (2018).
40. Krusche, P. *et al.* Best practices for benchmarking germline small-variant calls in human genomes. *Nat. Biotechnol.* **37**, 555–560 (2019).
41. Wagner, J. *et al.* *Towards a Comprehensive Variation Benchmark for Challenging Medically-Relevant Autosomal Genes*. <http://biorxiv.org/lookup/doi/10.1101/2021.06.07.444885> (2021) doi:10.1101/2021.06.07.444885.
42. Hui, J., Shomorony, I., Ramchandran, K. & Courtade, T. A. Overlap-based genome assembly from variable-length reads. in *2016 IEEE International Symposium on Information Theory (ISIT)* 1018–1022 (IEEE, 2016). doi:10.1109/ISIT.2016.7541453.
43. Formenti, G. *et al.* *Merfin: improved variant filtering and polishing via k-mer validation*. <http://biorxiv.org/lookup/doi/10.1101/2021.07.16.452324> (2021) doi:10.1101/2021.07.16.452324.
44. Olson, N. D. *et al.* *precisionFDA Truth Challenge V2: Calling variants from short- and long-reads in difficult-to-map regions*. <http://biorxiv.org/lookup/doi/10.1101/2020.11.13.380741> (2020) doi:10.1101/2020.11.13.380741.

45. Zook, J. M. *et al.* A robust benchmark for detection of germline large deletions and insertions. *Nat. Biotechnol.* **38**, 1347–1355 (2020).
46. Yang, C. *et al.* Evolutionary and biomedical insights from a marmoset diploid genome assembly. *Nature* (2021) doi:10.1038/s41586-021-03535-x.
47. Handsaker, R. E. *et al.* Large multiallelic copy number variations in humans. *Nat. Genet.* **47**, 296–303 (2015).
48. Vandepoele, K., Van Roy, N., Staes, K., Speleman, F. & van Roy, F. A novel gene family NBPF: intricate structure generated by gene duplications during primate evolution. *Mol. Biol. Evol.* **22**, 2265–2274 (2005).
49. Bosch, N. *et al.* Characterization and evolution of the novel gene family FAM90A in primates originated by multiple duplication and rearrangement events. *Hum. Mol. Genet.* **16**, 2572–2582 (2007).
50. Cantsilieris, S. *et al.* An evolutionary driver of interspersed segmental duplications in primates. *Genome Biol.* **21**, 202 (2020).
51. Cheng, H. *et al.* Robust haplotype-resolved assembly of diploid individuals without parental data. *ArXiv210904785 Q-Bio* (2021).
52. Porubsky, D. *et al.* Fully phased human genome assembly without parental data using single-cell strand sequencing and long reads. *Nat. Biotechnol.* **39**, 302–308 (2021).
53. Chen, Z. *et al.* Ultralow-input single-tube linked-read library method enables short-read second-generation sequencing systems to routinely generate highly accurate and economical long-range sequencing information. *Genome Res.* **30**, 898–909 (2020).
54. Poplin, R. *et al.* A universal SNP and small-indel variant caller using deep neural networks.

- Nat. Biotechnol.* **36**, 983–987 (2018).
55. Edge, P., Bafna, V. & Bansal, V. HapCUT2: robust and accurate haplotype assembly for diverse sequencing technologies. *Genome Res.* **27**, 801–812 (2017).
56. Porubsky, D. *et al.* Dense and accurate whole-chromosome haplotyping of individual genomes. *Nat. Commun.* **8**, 1293 (2017).
57. Ghareghani, M. *et al.* Strand-seq enables reliable separation of long reads by chromosome via expectation maximization. *Bioinformatics* **34**, i115–i123 (2018).
58. Vaser, R., Sović, I., Nagarajan, N. & Šikić, M. Fast and accurate de novo genome assembly from long uncorrected reads. *Genome Res.* **27**, 737–746 (2017).
59. Rozowsky, J. *et al.* AlleleSeq: analysis of allele-specific expression and binding in a network framework. *Mol. Syst. Biol.* **7**, 522 (2011).
60. Rhie, A., Walenz, B. P., Koren, S. & Phillippy, A. M. Merqury: reference-free quality, completeness, and phasing assessment for genome assemblies. *Genome Biol.* **21**, 245 (2020).
61. Shumate, A. *et al.* Assembly and annotation of an Ashkenazi human reference genome. *Genome Biol.* **21**, 129 (2020).
62. Vollger, M. R. *et al.* Long-read sequence and assembly of segmental duplications. *Nat Methods* **16**, 88–94 (2019).
63. Li, H. Minimap2: pairwise alignment for nucleotide sequences. *Bioinforma. Oxf. Engl.* **34**, 3094–3100 (2018).
64. Chin, C.-S. *et al.* A diploid assembly-based benchmark for variants in the major histocompatibility complex. *Nat. Commun.* **11**, 4794 (2020).

65. Köster, J. & Rahmann, S. Snakemake--a scalable bioinformatics workflow engine. *Bioinforma. Oxf. Engl.* **28**, 2520–2522 (2012).
66. O’Leary, N. A. *et al.* Reference sequence (RefSeq) database at NCBI: current status, taxonomic expansion, and functional annotation. *Nucleic Acids Res.* **44**, D733-745 (2016).
67. Smit, A. F. A., Hubley, R. & Green, P. Repeatmasker. *Open* **3.0**, (1996).
68. Morgulis, A., Gertz, E. M., Schäffer, A. A. & Agarwala, R. WindowMasker: window-based masker for sequenced genomes. *Bioinforma. Oxf. Engl.* **22**, 134–141 (2006).
69. Kapustin, Y., Souvorov, A., Tatusova, T. & Lipman, D. Splign: algorithms for computing spliced alignments with identification of paralogs. *Biol. Direct* **3**, 20 (2008).
70. Brown, G. R. *et al.* Gene: a gene-centered information resource at NCBI. *Nucleic Acids Res.* **43**, D36-42 (2015).
71. Guarracino, A., Heumos, S., Nahnsen, S., Prins, P. & Garrison, E. *ODGI: understanding pangenome graphs*. <http://biorxiv.org/lookup/doi/10.1101/2021.11.10.467921> (2021) doi:10.1101/2021.11.10.467921.
72. Kruglyak, S., Durrett, R., Schug, M. D. & Aquadro, C. F. Distribution and abundance of microsatellites in the yeast genome can Be explained by a balance between slippage events and point mutations. *Mol. Biol. Evol.* **17**, 1210–1219 (2000).