Resource

Single haplotype assembly of the human genome from a hydatidiform mole

Karyn Meltz Steinberg,¹ Valerie A. Schneider,² Tina A. Graves-Lindsay,¹ Robert S. Fulton,¹ Richa Agarwala,² John Huddleston,^{3,4} Sergey A. Shirayev,² Aleksandr Morgulis,² Urvashi Surti,⁵ Wesley C. Warren,¹ Deanna M. Church,⁶ Evan E. Eichler,^{3,4} and Richard K. Wilson¹

¹The Genome Institute at Washington University, St. Louis, Missouri 63108, USA; ²National Center for Biotechnology Information, National Library of Medicine, National Institutes of Health, Bethesda, Maryland 20894, USA; ³Department of Genome Sciences, University of Washington, Seattle, Washington 98195, USA; ⁴Howard Hughes Medical Institute, University of Washington, Seattle, Washington 98195, USA; ⁵Department of Pathology and Human Genetics, University of Pittsburgh, Pittsburgh, Pennsylvania 15260, USA; ⁶Personalis, Inc., Menlo Park, California 94025, USA

A complete reference assembly is essential for accurately interpreting individual genomes and associating variation with phenotypes. While the current human reference genome sequence is of very high quality, gaps and misassemblies remain due to biological and technical complexities. Large repetitive sequences and complex allelic diversity are the two main drivers of assembly error. Although increasing the length of sequence reads and library fragments can improve assembly, even the longest available reads do not resolve all regions. In order to overcome the issue of allelic diversity, we used genomic DNA from an essentially haploid hydatidiform mole, CHM1. We utilized several resources from this DNA including a set of end-sequenced and indexed BAC clones and 100× Illumina whole-genome shotgun (WGS) sequence coverage. We used the WGS sequence and the GRCh37 reference assembly to create an assembly of the CHM1 genome. We subsequently incorporated 382 finished BAC clone sequences to generate a draft assembly, CHM1_I.I (NCBI AssemblyDB GCA_000306695.2). Analysis of gene, repetitive element, and segmental duplication content show this assembly to be of excellent quality and contiguity. However, comparison to assembly-independent resources, such as BAC clone end sequences and PacBio long reads, indicate misassembled regions. Most of these regions are enriched for structural variation and segmental duplication, and can be resolved in the future. This publicly available assembly will be integrated into the Genome Reference Consortium curation framework for further improvement, with the ultimate goal being a completely finished gap-free assembly.

[Supplemental material is available for this article.]

The production of a reference sequence assembly for the human genome was a milestone in biology and clearly has impacted many areas of biomedical research (McPherson et al. 2001; International Human Genome Sequencing 2004). The availability of this resource allows us to investigate genomic structure and variation at a depth previously unavailable (Kidd et al. 2008; The 1000 Genomes Project Consortium 2012). These studies have helped make clear the shortcomings of our initial assembly models and the difficulty of comprehensive genome analysis. While the current human reference assembly is of extremely high quality and is still the benchmark by which all other human assemblies must be compared, it is far from perfect. Technical and biological complexity lead to both missing sequences as well as misassembled sequence in the current reference, GRCh38 (Robledo et al. 2002; Eichler et al. 2004; International Human Genome Sequencing 2004; Church et al. 2011; Genovese et al. 2013).

The two most vexing biological problems affecting assembly are (1) complex genomic architecture seen in large regions with highly homologous duplicated sequences and (2) excess allelic diversity (Bailey et al. 2001; Mills et al. 2006; Korbel et al. 2007;

Kidd et al. 2008; Zody et al. 2008). Assembling these regions is further complicated due to the fact that regions of segmental duplication (SD) are often correlated with copy-number variants (CNVs) (Sharp et al. 2005). Regions harboring large CNV SDs have been misrepresented in the reference assembly because assembly algorithms aim to produce a haploid consensus. Highly identical paralogous and structurally polymorphic regions frequently lead to nonallelic sequences being collapsed into a single contig or allelic sequences being improperly represented as duplicates. Because of this complexity, a single, haploid reference is insufficient to fully represent human diversity (Church et al. 2011).

The availability of at least one accurate allelic representation at loci with complex genomic architecture facilitates the understanding of the genomic architecture and diversity in these regions (Watson et al. 2013). To enable the assembly of these regions, we have developed a suite of resources from CHM1, a DNA source containing a single human haplotype (Taillon-Miller et al. 1997; Fan et al. 2002). A complete hydatidiform mole (CHM) is an abnormal product of conception in which there is a very early fetal demise

Corresponding author: rwilson@genome.wustl.edu

Article published online before print. Article, supplemental material, and publication date are at <http://www.genome.org/cgi/doi/10.1101/gr.180893.114>.

© 2014 Steinberg et al. This article is distributed exclusively by Cold Spring Harbor Laboratory Press for the first six months after the full-issue publication date (see <http://genome.cshlp.org/site/misc/terms.xhtml>). After six months, it is available under a Creative Commons License (Attribution-NonCommercial 4.0 International), as described at <http://creativecommons.org/licenses/by-nc/4.0/>.

and overgrowth of the placental tissue. Most CHMs are androgenetic and contain only paternally derived autosomes and sex chromosomes resulting either from dispermy or duplication of a single sperm genome. The phenotype is thought to be a result of abnormal parental contribution leading to aberrant genomic imprinting (Hoffner and Surti 2012). The absence of allelic variation in monospermic CHM makes it an ideal candidate for producing a single haplotype representation of the human genome. There are a number of existing resources associated with the “CHM1” sample, including a BAC library with end sequences generated with Sanger sequencing using ABI 3730 technology (<https://bacpac.chori.org/>), an optical map (Teague et al. 2010), and a BioNano genomic map (see Data access), some of which have previously been used to improve regions of the reference human genome assembly.

BAC clones have historically been used to resolve difficult genomic regions and identify structural variants (Barboudi et al. 2004; Carvalho and Lupski 2008). A BAC library constructed from CHM1 DNA (CHORI-17, CH17) has also been utilized to resolve several very difficult genomic regions, including human-specific duplications at the *SRGAP2* gene family on Chromosome 1 (Dennis et al. 2012). Additionally, the CHM1 BAC clones were used to generate single haplotype assemblies of regions that were previously misrepresented because of haplotype mixing (Watson et al. 2013). Both of these efforts contributed to the improvement of the GRCh38 reference human genome assembly, adding hundreds of kilobases of sequence missing in GRCh37, in addition to providing an accurate single haplotype representation of complex genome regions.

Because of the previously established utility of sequence data derived from the CHM1 resource, we wished to develop a complete assembly of a single human haplotype. To this end, we produced a short read-based (Illumina) reference-guided assembly of CHM1 with integrated high-quality finished fully sequenced BAC clones to further improve the assembly. This assembly has been annotated using the NCBI annotation process and has been aligned to other human assemblies in GenBank, including both GRCh37 and GRCh38. Here we present evidence that the CHM1 genome assembly is a high-quality draft with respect to gene and repetitive element content as well as by comparison to other individual genome assemblies. We will also discuss current plans for developing a fully finished genome assembly based on this resource.

Results

We generated an assembly of the complete hydatidiform mole, CHM1, a genome comprised of 23 chromosomes (1–22 and X and MT) with a total sequence length of 3.04 Gb. Contig N50 length is ~144 kbp and scaffold N50 length is 50 Mbp (Table 1). These N50 statistics were based upon the reference-guided assembly with BAC tiling paths incorporated and are defined as the length of the contig or scaffold where 50% of the assembly is assembled into contigs/scaffolds that are of that length or greater. Compared to other WGS human assemblies, HuRef (J. Craig Venter assembly; GenBank GCA_000002125.2), ALLPATHS (GenBank AEKP00000000.1), and YH_2.0 (GenBank GCA_000004845.2), the CHM1_1.1 assembly has a lower contig number and a higher contig N50, demonstrating that CHM1_1.1 is more contiguous than previously generated individual genome assemblies (Fig. 1). We incorporated high-quality sequence from 382 BAC clones (Supplemental Table S1) to improve the assembly in complex regions where the GRCh37 reference was incorrect (Fig. 2). BAC end sequences themselves were not in-

Table 1. Assembly statistics

Total sequence length	3,037,866,619
Total assembly gap length	210,229,812
Gaps between scaffolds	225
Number of scaffolds	163
Scaffold N50	50,362,920
Number of contigs	40,828
Contig N50	143,936
Total number of chromosomes	23

corporated into the assembly but were used to analyze the quality of the assembly.

We assessed the integrity and fidelity of CHM1 with respect to the GRCh37 reference by analyzing CHORI-17 BAC end sequence mapping to GRCh37. Approximately 95.5% clone ends mapped uniquely concordantly, 4% mapped uniquely discordantly, and the remaining 0.5% mapped to multiple locations. These statistics indicate that the genomic DNA derived from the CHM1 cell line that was used to create the BAC library and Illumina libraries is not grossly rearranged and represents a suitable template for a platinum reference. In addition, analyses from an optical map generated using CHM1 genomic DNA do not show an excess of structural variants that would suggest somatic rearrangement (Teague et al. 2010). SNP genotyping also confirms the haploid content of the cell line, and karyotyping was performed at several stages during passaging to ensure the integrity (Fan et al. 2002).

Assessment of assembly quality

Repetitive element content

The assembly was masked with both WindowMasker (Morgulis et al. 2006) and RepeatMasker (Smit et al. 1996–2010) and 34.29% and 47.21% of the assembly was masked, respectively. This is comparable to the repetitive content of GRCh37 (34.24% and 47.15%, WindowMasker and RepeatMasker, respectively). When the repetitive elements are parsed out by type, the numbers of each element are comparable between GRCh37 and CHM1_1.1 (Table 2).

Gene content

This analysis is based on NCBI *Homo sapiens* annotation run 105 that includes GRCh37.p13, CHM1_1.1, HuRef, and the single chromosome assembly CRA_TCAGchr7v2. For our comparison, we only used the annotations on the original GRCh37 Primary Assembly sequences, as many of the fix patches in patch release 13 are based on CHM1. Using this annotation run provides a better comparison than the original GRCh37 annotation, because the same software and evidence set was used.

Using gene annotation as a proxy for assembly quality, the results indicate that the CHM1_1.1 assembly (39,009 total genes, 19,892 protein coding genes) (Table 3) is of higher quality than the HuRef assembly (38,070 total genes, 19,668 protein-coding genes), though not quite as good as the GRCh37 assembly (39,947 total genes, 20,072 protein-coding genes). The alignment evidence used to support each gene model supports this conclusion. CHM1_1.1 has 21 genes annotated with a “transcription discrepancy” compared to 15 in GRCh37. Interestingly, some genes are problematic in both assemblies, such as *OVGP1* and *MUC19*, suggesting that even in a single haplotype background, complex gene family regions can be difficult to assemble (Supplemental Data: Gene Annotation).

Figure 1. Comparison of contig count and contig N50 between CHM1_1.1 (GenBank GCA_000306695.2) and HuRef (J. Craig Venter assembly; GenBank GCA_000002125.2), ALLPATHS (GenBank AEP000000000.1) and YH_2.0 (GenBank GCA_000004845.2) WGS assemblies. CHM1_1.1 has only 10%–20% of the number of total contigs as the other assemblies and has a contig N50 1.5 to six times larger.

While GRCh37 may have better global gene-annotation metrics, there are regions in which CHM1_1.1 performs better. For example, we identified 549 genes unique to the CHM1 assembly (i.e., absent from the GRCh37.p13 primary assembly; Supple-

mental Table S2). *MUC3B*, a membrane bound mucin that maps to Chromosome 7q22 (NC_018918.2: 100477710-100541651) is annotated only on CHM1_1.1 as predicted from Gnomon gene models. The protein produced by *MUC3B* functions as a major glycoprotein component of mucus gel at the intestinal surface that provide a barrier against foreign particles and microbial organisms. It is part of a tandem duplication involving *MUC3A*, and *MUC3B* is expressed exclusively in the small intestine and colon (Pratt et al. 2000; Kyo et al. 2001). Variants of *MUC3A* have been associated with inflammatory bowel disease, and up-regulation of *MUC3* inhibited adherence of pathogenic *E. coli* in human intestinal cells (Pan et al. 2013). The CHM1 version of *MUC3B* contains four copies of the tandem repeat.

Other clinically relevant CHM1 genes not present in the GRCh37.p13 primary assembly include *KCNJ18* and *DUX4L1*. *KCNJ18* is a member of a large gene family of potassium inwardly rectifying channels located on 17q11.2 (NC_018928.2: 21605469-21617558). It is expressed mostly in skeletal muscle and regulated by thyroid hormone. Mutations in this gene have been associated with thyrotoxic hypokalemic periodic paralysis (MIM 613239) (Ryan et al. 2010). *DUX4L1* encodes a transcription factor comprised of two homeobox domains located within a macrosatellite repeat in the subtelomeric region of 4q (NC_018915.2: 190981943-190983264)

Figure 2. (A) WGS assembly from the first pass (CHM1_1.0; GCF_000306695.1, bronze line) on Chromosome 1p12 (NC_018912.1: 121,050,000–121,400,000) demonstrated a gap (gray box) in the assembly (assembly name: AMYH010000980.1, green lines). Using MEGABLAST, two CH17 clones (AC247039.2 and AC253572.3, red lines) aligned to the region and appeared to span the gap. (B) By incorporating these BAC sequences into the assembly, the gap was subsequently resolved in CHM1_1.1 (NC_018912.2: 121,050,000–121,650,000). The tiling path components, FP325311.11, AC241952.2, AC247039.2, AC253572.3, and AC241377.3 indicate the clone names used to resolve the gap. The clones from A are indicated in red while the other clones are in purple. The final assembly–assembly alignment is indicated in purple, showing the gap resolution.

Table 2. Comparison of repetitive elements in GRCh37 and CHM1_1.1

	CHM1	GRCh37
DNA	482,783	461,751
LINE	1,511,690	1,498,690
LTR	716,720	717,656
Low complexity	376,835	371,543
RNA	778	729
SINE	1,757,213	1,793,723
Satellite	3,099	9,566
Simple repeat	398,210	417,913
rRNA	1,749	1,769
scRNA	1,301	1,340
snRNA	4,306	4,386
srpRNA	1,665	1,481
tRNA	1,852	2,002
Other	19,338	15,581

(Hewitt et al. 1994). Repeat copy-number variation is associated with facioscapulohumeral muscular dystrophy (MIM 158900) (Bosnakovski et al. 2008). Both of these genes are now annotated in GRCh38 with information from the CHM1 data.

Clinical allele analysis

Using data from the NHGRI GWAS catalog and ClinVar (Landrum et al. 2014), we assessed the number of risk alleles present in the CHM1 genome. Most loci could be successfully remapped from GRCh37 to CHM1 (7962/7991 NHGRI GWAS loci and 43,614/48,516 ClinVar loci) using the NCBI Remap tool (<http://www.ncbi.nlm.nih.gov/genome/tools/remap>). The CHM1 genotype matched the “risk” allele at 3284 loci from NHGRI GWAS and 291 loci from ClinVar. CHM1 carries an associated allele for 366 unique pheno-

types out of a total of 1089 unique phenotypes in NHGRI GWAS. Of the 291 matching ClinVar alleles, 22 are categorized as pathogenic. Two of the 22 pathogenic alleles are nonsense mutations that cause autosomal recessive disorders: Hemochromatosis type 2A and Marinesco-Sjogren syndrome. The remaining 18 pathogenic alleles have global minor allele frequencies at least 1% (Supplemental Table S3). Overall, the CHM1 genome does not appear to harbor an excessive number of risk alleles or extremely rare alleles associated with diseases (Supplemental data: Clinical Allele Analysis). A similar analysis performed on the GRCh37 assembly identified 3556 disease susceptibility variants and 15 risk alleles with MAF < 1% (Chen and Butte 2011).

Representation of segmental duplication

Analysis of SDs suggests CHM1_1.1 has good representation of large, duplicated sequences. Using whole-genome assembly comparison (WGAC), we discovered 54,580 pairwise alignments corresponding to 130.9 Mbp of nonredundant duplications or 4.6% of the genome (Supplemental Table S4). Intrachromosomal events comprise a majority of the SDs with 99.7 Mbp in contrast to 57.3 Mbp of interchromosomal duplications. Additionally, intrachromosomal alignments are generally longer and more similar than interchromosomal alignments (Supplemental Fig. S1). Both of these patterns are consistent with our previous WGAC analysis of GRCh37 (Sudmant et al. 2013). Using an alternative approach to detect SD based on read depth analysis (WSSD) (Bailey et al. 2001), we identified 124.6 Mbp of duplicated sequences (4.4% of the genome). These WSSD duplications supported 89.5 of 96.1 Mbp (93%) WGAC duplications that were also ≥ 10 kbp and > 94% identity (Supplemental Fig. S2). Correspondingly, 119.6 Mbp of WSSD duplications (96%) overlapped or occurred within 5 kbp of a WGAC duplication. To determine how CHM1.1 WGAC duplications compared to duplications from GRCh37, we remapped the

Table 3. Gene predictions and models for GRCh37.p13, GRCh37 primary assembly, CHM1_1.1, and HuRef assemblies

Feature	GRCh37.p13	GRCh37.p13 Primary Assembly	CHM1_1.1	HuRef
Genes and pseudogenes	40,158	39,947	39,009	38,070
Protein-coding	20,176	20,072	19,892	19,668
Noncoding	7,667	7,627	7,529	7,151
Pseudogenes	12,315	12,248	11,588	11,251
Genes with variants	15,068	14,994	8,718	8,620
Placed on multiple assembly units	2,665	na	na	na
mRNAs	67,517	64,734	35,142	34,843
Fully supported	67,267	64,514	34,941	34,630
With > 5% ab initio	203	181	162	180
Partial	116	138	602	4,357
Placed on multiple assembly units	2,171	na	na	na
Known RefSeq (NM_)	34,632	34,606	34,367	34,212
Model RefSeq (XM_)	32,885	30,128	775	631
Model RefSeq (XM_) with correction	17	15	21	15
Other RNAs	15,063	14,151	11,408	10,854
Fully supported	13,599	13,123	10,384	9,960
With > 5% ab initio	0	0	0	0
Partial	2,369	2,370	2,401	2,557
Placed on multiple assembly units	279	na	na	na
Known RefSeq (NR_)	6,623	6,618	6,458	6,328
Model RefSeq (XR_)	7,011	6,538	3,946	3,650
CDs	68,035	65,099	35,522	35,173
Fully supported	67,267	64,514	34,941	34,630
With > 5% ab initio	220	197	175	192
Partial	96	118	313	2,966
Known RefSeq (NP_)	34,632	34,605	34,365	34,185
Model RefSeq (XP_)	32,885	30,128	775	631
Model RefSeq (XP_) with correction	17	15	21	15

WGAC alignments from CHM1.1 to GRCh37.p13 with NCBI's remap API. After remapping and omitting coordinates from patches, there were 138 Mbp of CHM1.1 duplications. The two assemblies shared 123 Mbp of duplications corresponding to 89% of CHM1.1 duplications and 90% of GRCh37 (Fig. 3).

Identification of misassemblies

The goal for this project is a completely closed reference assembly containing no gaps. Therefore, it is critical for us to identify the extent of misassembly as well as the specific regions involved for targeted correction. We have already begun the process of loading the assembly and curation regions into the GRC curation database and framework. We performed three separate analyses to assess the integrity of the assembly and identify potential misassemblies.

Identification of heterozygous SNVs

CHM1 is an essentially homozygous resource. Thus, there should be no heterozygous SNVs identified upon aligning the CHM1 reads to GRCh37, and there should be no SNVs identified when these reads are aligned to the CHM1_1.1 assembly. We were therefore interested in using SNV detection to identify potentially misassembled regions in both GRCh37 as well as CHM1_1.1. First, we aligned the Illumina reads from CHM1 libraries to the GRCh37.p13 primary assembly and identified 99,572 heterozygous sites and 2,445,270 homozygous sites. We stratified heterozygous SNVs based on whether they overlapped repetitive or low complexity sequence (Supplemental Table S5). A recent study demonstrated that up to 60% of heterozygous SNVs called from CHM1 Illumina reads aligned to the reference are within low complexity regions (including SDs) of the human genome (Li 2014). We focused on 25,529 heterozygous variants that did not fall within a repetitive

sequence (heterozygous nonrepetitive: HNR variants), as these may be sites of cryptic duplication in the reference sequence or structural variation in CHM1. The HNR variants were overlapped with the RefSeq annotation (Supplemental Table S6) and the functional consequences of each variant were predicted (Fig. 4).

The genes with the most HNR variants were then compared to genes missing copies in the reference and genes with significantly population stratified copy number (high V_{st}) from Sudmant et al. (2010). Genes with known missing copies in the reference assembly, such as *GPRIN2* and *DUSP22*, have 20 and 56 HNR variants, respectively, while high V_{st} genes such as *PDE4DIP* have 267 HNR variants. The gene with the most HNR variants ($N = 618$) is *PRM2* that is part of interchromosomal duplications of Chromosomes 6 and 3 and represents cryptic SDs in the GRCh37 reference genome (Genovese et al. 2013). Additionally, two regions that were incorrectly represented in GRCh37 and subsequently resolved in GRCh38 using the CHM1 derived BAC library, *SRGAP2* (Dennis et al. 2012) and *IGH* (Watson et al. 2013), both had high counts of HNR variants (39 and 54, respectively) providing additional support for the hypothesis that heterozygous calls are indicative of reference assembly errors. The majority of the heterozygous calls are errors that arise during variant detection due to paralogous sequences mapping to low complexity regions.

We then aligned the Illumina reads from the CHM1 libraries to the CHM1_1.1 assembly. A total of 86,544 SNVs were called, and 79% of these variants overlap repetitive sequence (RepeatMasker and WGAC) (Supplemental Table S7). There is a significant enrichment of variants in repetitive sequence compared to sequence not annotated as repetitive (1000 permutations, simulation based P -value < 0.001). Thirty-four regions totaling 49 Mb have SNV density per kilobase two standard deviations higher than the mean SNV density per kilobase of 0.03 (Supplemental Table S8). Sixty-

four percent of the bases in SNV-rich regions are annotated as repetitive. There are 294 unique RefSeq genes in SNV rich regions including the beta-defensin gene cluster on Chromosome 8 and *NBPF1* on Chromosome 1 (Supplemental Table S9). These regions are highly duplicated and the variant calls likely represent paralogous sequence variants.

CH17 BAC ends mapped to CHM1_1.1 assembly

We aligned the BAC end sequences derived from the CHORI-17 BAC library to the CHM1_1.1 assembly. As this is the same DNA source as the assembly, there should be no structural variation. The majority of placements were concordant (96.22%), suggestive of a high-quality assembly; however, regions with multiple discordant alignments may represent assembly errors. A query set of 306,838 BAC end sequences representing 158,396 unique clones from the CH17 BAC library was aligned to the CHM1_1.1 assembly (Supplemental Table S10). We identified 1192 regions with 3927 unique clones that likely contained assembly errors based on an unexpected size distribution of the aligned BACs. Among unique dis-

Figure 3. (A) Comparison of segmental duplications (SDs) in GRCh37 and CHM1.1 assemblies predicted by WGAC analysis by chromosome. The duplication content is comparable between GRCh37 and CHM1_1.1 assemblies indicating good assembly quality. (B) Venn diagram of SDs in GRCh37 and CHM1_1.1 assemblies shows that most duplications are shared between the assemblies.

Figure 4. Functional consequences of CHM1 heterozygous variants not in repetitive sequence (HNR variants). Approximately 97% of HNR variants are intergenic or intronic. Of the remaining 3% of other variants, ~48% are in the 3' or 5' UTR, 17% are silent, and 35% are coding (missense, nonsense, essential splice site).

cordant clones, 2840 suggested a deletion in the assembly and 443 suggested additional sequence in the CHM1_1.1 assembly not represented in the BAC resource. The regions demonstrating insertion may be due to instability in BAC clones. On average, there are significantly more bases in SDs (based on WGAC) in the single discordant and multiple mapped clone ends compared to the single concordant clone ends (mean: 0.24, 0.96, 0.04; standard deviations: 0.18, 0.14, 0.02, respectively, for single discordant, multiple, and single concordant; Student's *t*-test, two tailed $P = 0$ for each comparison). The remaining unique discordant placements were comprised of incorrectly oriented ends, indicating that the assembly and clone sequences are inverted relative to one another.

CH17 clones with discordant placements on the CHM1_1.1 assembly may be used to identify regions misassembled due to errors in the reference or genomic variation. For example, in the SMA duplication region at 5q13.3 (NC_018916.2) (Supplemental Fig. S3), the GRCh37 reference chromosome represents a single resolved SMA haplotype (Schmutz et al. 2004). However, many CH17 clones aligning to the corresponding region of the CHM1_1.1 assembly have discordant placements that are characterized by inversions and size discrepancies, suggesting that the CHM1_1.1 assembly does not faithfully represent the CHM1 genome at this locus. This observation is consistent with the known variability of this genomic region in the human population, which is associated with its complex SD structure (Ogino et al. 2004). It should be noted, however, that the clone placements located within the local BAC assemblies in this region are largely concordant, whereas those associated with WGS contigs are discordant. This result demonstrates how the use of sequence from large insert clones can resolve regions too complex for even the reference-guided assembly of WGS contigs. Assembly with additional BAC clones will likely be required to close the existing gaps and fully resolve the CHM1 genomic sequence in these complex regions.

Alignment of a long read data set

To identify errors in the CHM1_1.1 genome assembly (GCF_000306695.2) introduced as a consequence of errors in the GRCh37 primary assembly unit that was used to guide its assem-

bly, we aligned CHM1 PacBio reads, which became available after the CHM1_1.1 assembly was complete and annotated, to the CHM1_1.1 assembly. We hypothesized that these alignments in such regions of CHM1_1.1 would exhibit one or more of the following characteristics: (1) low coverage with respect to coverage in surrounding regions; (2) sharp boundaries at which alignment coverage drops off; or (3) inversions. Low coverage is often associated with highly fragmented assembly regions, which are themselves hallmarks of assembly problems (though they may not necessarily reflect errors introduced by GRCh37). Sharp boundaries could occur at component boundaries (indicative of GRCh37 tiling path errors) or within assembly components (indicative of component assembly errors in GRCh37). Although other assembly features (i.e., repetitive elements or structural variation) can also result in read alignments having similar characteristics, such regions should be enriched for assembly errors.

To identify CHM1_1.1 assembly errors corresponding to unrecognized GRCh37 errors, we focused on CHM1_1.1 assembly sites where alignment coverage dropped off sharply. To this end, we produced a list of regions where there were PacBio aligned reads that met the above criteria, and we refer to these reads as "cliffs." We focused on bins where the cliff count is greater than or equal to 10 and the depth is less than $2\times$ the coverage (< 108) to eliminate artifacts from repetitive elements. There are 274 loci where cliffs are within 1 kbp of the component boundary and 2109 loci where cliffs are > 1 kbp from the component boundary (Supplemental Data: PacBio). Using this approach we are able to clearly visualize regions with assembly errors such as the one on Chromosome 11, where two tiling path components are inverted in the CHM1_1.1 assembly and require correction (Fig. 5).

Discussion

There has been a dramatic decrease in sequence cost with a concomitant increase in throughput leading to the availability of thousands of sequenced genomes and exomes. However, analysis of individual genomes depends upon the availability of a high-quality reference assembly. Despite the high quality of the human

Figure 5. Overview of the Chr 11 (NC_018922.2) 1.9-Mb region, exhibiting three alignment bins with a large number of PacBio “cliff” reads where the alignment coverage dropped off sharply. WGS component (light green lines) boundaries flanked by such reads are marked with red dashed lines. The ends of each component at the boundary are labeled with letters to show orientation. Pairs of alignments corresponding to three different PacBio reads are marked in yellow, green, and dark blue. These alignments overlap by < 10% on each of the reads. The split alignments for these three reads suggest that the two WGS components marked in purple should be inverted and translocated as indicated by the arrow at the top of the image. The other PacBio reads in these bins exhibit the same pattern of split alignments, which supports the proposed reordering and orientation of the WGS components. The bottom light green lines show a proposed tiling path with the orientation corrected; the letters indicate where each end of the initial tiling path components should be placed.

reference assembly, many groups have described shortcomings of this resource, including remaining gaps, single-nucleotide errors, or gross misassembly due to complex haplotypic variation (Eichler et al. 2004; Doggett et al. 2006; Kidd et al. 2010; Chen and Butte 2011; The 1000 Genomes Project Consortium 2012). Both gaps and misassembled regions often arise because the DNA sequence used for the assembly was from multiple diploid sources containing complex structural variation. Because such loci often contain medically relevant gene families, it is important to resolve variation at these sites, as the structural and single-nucleotide diversity is likely associated with clinical phenotypes (Eichler et al. 2004). Thus, to resolve structurally complex regions and provide a more effective reference resource for such loci, we combined WGS data and BAC sequences from a haploid DNA source to create a single haplotype assembly of the human genome.

Haplotype information is critical to interpreting clinical and personal genomic information as well as genetic diversity and ancestry data, and most previously sequenced individual human genomes are not haploresolved. The current reference human genome sequence represents a mosaic that further complicates haplotyping; within a BAC clone there is a single haplotype representation, but haplotypes can switch at BAC clone junctions. By utilizing an essentially haploid DNA source, we resolved a single haplotype across complex regions of the genome where the reference genome contained a mixture of haplotypes from various sources and/or contained unresolved gaps. For example, a gap on Chromosome 4p14 in GRCh37 (Chr 4: 40296397–40297096) was completely resolved using CHM1 WGS data. The gap was flanked by repetitive elements that were not traversed by a clone. This region has subsequently been updated with a complete tiling path in GRCh38.

The addition of high-quality BAC sequence to our assembly was vital to resolving gaps. For example, in GRCh37 at Chromosome 15q25.2 there was a 79-kbp gap due to over-expansion of a hypervariable region. This region contains many *GOLGA6L* core duplication genes (Jiang et al. 2007) and highly identical SDs. RP11 BAC clones on one side and RP13 BAC clones on the other side flanked the gap. Using the BAC-based sequence-resolved CH17 haplotype, and the gap was filled in GRCh38 (Supplemental Fig. S4). A preliminary analysis of PacBio data shows this region remains unresolved even using long read sequencing. This underscores the importance of curation and employing multiple sequencing strategies to obtain an accurate genome representation.

Despite the high quality of the CHM1_1.1 assembly, we did identify regions that require further improvement. Some of these problems are due to the repetitive nature of the loci, while others are due to using GRCh37 to guide the CHM1_1.1 assembly. The availability of diverse, assembly independent resources, including the recently released long read data set from PacBio provide a pathway for problem identification and correction. Unfortunately, the CHM1_1.1 assembly was completed before the PacBio data were released; however, we anticipate that the next version of the assembly will incorporate this valuable data set. Additionally, a preliminary analysis of data from BioNano Genomics also demonstrates that a second, independent source of data can be useful for improving the assembly (Supplemental Data: BioNano Analysis). The GRC has established the infrastructure to support assembly curation and the development of highly refined reference assemblies, as evidenced by the release of two successive human genome assemblies (GRCh37 and GRCh38) and a mouse genome assembly (GRCm38). We have already begun using these resources to improve the CHM1_1.1 assembly.

We chose a reference-guided assembly method rather than performing a de novo assembly of the short WGS reads. An analysis of a de novo assembly from short reads using the SOAP algorithm (Li et al. 2008) found significant contamination and missing sequences (Alkan et al. 2011). In general, the de novo assemblies were ~16% shorter than the reference genome, and > 99% of previously validated duplications were missing translating to more than 2300 missing coding exons. Another human assembly from massively parallel sequences using the ALLPATHS-LG (Gnerre et al. 2011) showed improvements over the SOAP assembly but still only covered ~40% of SDs. As described above, the gene and repetitive element coverage of the CHM1_1.1 assembly is comparable to GRCh37. We did not do a formal comparison to GRCh38 because many of the CHM1 BAC tiling paths are used in both assemblies, meaning they are no longer completely independent. Approximately 29 Mb of clone sequence and 134 kbp of WGS sequence from the CHM1_1.1 assembly has been incorporated into the GRCh38 primary assembly, while > 13 Mb of clone sequence has been utilized for alternative sequence representations. The somewhat fragmented nature of the CHM1_1.1 assembly means it is not ready to become the Primary assembly in the GRCh series of reference assemblies; however, our goal is to improve this assembly so that it could serve this role.

A single haplotype reference assembly will not be sufficient for alignment and variant detection in large-scale human genomic studies. Individual-specific sequences between a random pair of human individuals ranges between 1.8 and 4 Mb (Li et al. 2010). The GRC formalized the concept of multiallelic representation of complex genome regions with the release of GRCh37. The newest reference genome GRCh38 contains 261 alternative sequence representations at 178 regions, many of which were resolved using

the CHM1 data. A recent study provides the basis for a more formal graph representation (Paten et al. 2014) but a great deal of tool development needs to occur before we can formally move to such an assembly representation. While this development occurs, the current multiallelic reference provides data that allow us to explore complex genomic regions. This is important for analyzing genes that are only represented on the alternate loci, as they need to somehow be represented in the assembly. In the current GRCh38 there are 153 genes present only on the alternate loci, some of which are medically relevant such as *HLA-DRB3* and *LILRA3*. A linear genome representation would not be able to correctly handle these cases. The use of the single haplotype CHM1 resource has proven quite valuable in resolving several complex regions of the human genome. In many of these cases, the GRCh37 representation was the mixture of several haplotypes and not likely found in any individual. We plan on continuing to develop this resource in an effort to ensure that we have at least one correct representation of all loci in the human genome.

Methods

Cell line

CHM1 cells were grown in culture from one such conception at Magee-Women's Hospital (Pittsburgh, PA) after parental consent and IRB approval. Cryogenically frozen cells from this culture were grown and transformed using human telomerase reverse transcriptase (TERT) to develop a cell line. This cell line retains a 46,XX karyotype and complete homozygosity. It was subsequently used for genomic research by multiple investigators and was also used to prepare a BAC library (CHORI 17; <https://bacpac.chori.org/>) for further research.

Illumina sequencing

We performed whole-genome shotgun sequencing on the CHM1 DNA. KAPA qPCR (KK4854, Kapa Biosystems) was performed on the LightCycler 480 System (Roche Life Sciences) to quantify the libraries and determine the appropriate concentration to produce optimal recommended cluster density on a HiSeq 2000 V2 or V3 2 × 100-bp sequencing run. HiSeq 2000 V2 and V3 runs were completed according to the manufacturer's recommendations (Illumina). We generated > 617 Gb of sequence used for the assembly. The average insert size was 315 bp for three libraries, 3 kbp for three libraries, and 8 kbp for two libraries.

Assembly

Assembly of the CHM1 genome used deep coverage WGS sequence reads generated using the Illumina HiSeq platform. These data are publicly available in NCBI's Sequence Read Archive (SRA; <http://www.ncbi.nlm.nih.gov/sra>) under project SRP017546. The project has ten experiments, of which one was a pilot experiment using 25-bp unpaired reads while the remaining nine were all paired-end reads. Eight experiments had a total of 31 runs and were used in producing the assembly (Supplemental Table S11). Reads were aligned to the GRCh37 primary assembly using SRPRISM v2.3 aligner.

A reference-guided assembly was produced using ARGO v1.0. This assembly has 2,818,728,129 bp in 47,737 contigs with N50 of 139,647 bp. Both SRPRISM and ARGO were developed at NCBI, but are not yet published. SRPRISM v2.5 is available at <ftp://ftp.ncbi.nlm.nih.gov/pub/agarwala/srprism/>. Briefly, SRPRISM creates an index on the reference genome and uses the index to find locations on the genome to do extensions. It has resource

requirements and performance characteristics comparable to the fastest available aligners, yet provides explicit criteria for search sensitivity and reports all results that have the same quality. ARGO uses conservative heuristics that take into account the insert size and read orientation to produce a most likely sequence for the assembly.

In addition to the Illumina data, we selected BAC clones that overlapped SDs and corresponded to regions in SDs, known structural variation, and loci around gaps. Two hundred ninety six clones in 45 tiling paths and 104 singleton clones were provided. By mapping the clone information to the reference-guided assembly using BLAST and manually reviewing the alignments to decide the best location in the assembly to incorporate the clone sequence, some of the worst regions of the assembly were significantly improved. Four singleton clones could not be used as they are significantly diverged from the assembly. Fourteen additional clones were redundant with other clones. Clone AC243629.2 has an internal expansion that was discovered after the assembly release and has now been subsequently removed. After incorporating clone information, the assembly had 2,846,046,639 bp in 41,406 contigs with an N50 of 143,718 bp. Prior to submission of the assembly to GenBank, the contigs were subsequently filtered to remove some WGS that was redundant to one of the clone paths, to remove small WGS contigs at chromosome termini, to trim terminal Ns from WGS contigs, and to accommodate a newly finished clone component and then scaffolded according to alignment with the GRCh37 primary assembly.

The reference assembly process includes (1) taking the aligned reads and building an insert size histogram implied by the alignments, (2) determining valid insert size range and orientation per run, (3) filtering alignments to retain alignments that are consistent with the insert size range and orientation, (4) subdividing the reference into regions of matches, mismatches, and gaps depending on whether all reads aligning to the reference position match the base, at least some reads do not match the base, or the base is not aligned to by any read, respectively, and (5) producing sequence for each region when possible. Sequence for matching regions is reproduced as it is in the reference. Sequence for regions with mismatches is produced by doing a multiple alignment of the portions of reads that are aligning to that region with the sequence with most copies in the multiple alignment retained as the sequence for the region. Sequence for gaps between two consecutive contigs on the reference is produced by using reads whose one mate aligns to either contig and the other mate has a 30-mer matching the end of one of the contigs such that the mate results in extension of the contig and preserves insert size and orientation constraints. This process may or may not result in filling the gap completely.

Gene annotation

The CHM1_1.1 assembly was masked using RepeatMasker (Smit et al. 1996–2010) and annotated using the NCBI Eukaryotic Genome Annotation Pipeline. Briefly, the assembly is masked using RepeatMasker and then aligned to a set of same-species RefSeq transcripts and genomic sequences to directly annotate the gene, RNA, and CDS features. The assembly is also aligned to Gnomon gene prediction models. Gnomon is a two-step gene prediction program that assembles overlapping alignments into “chains” followed by a prediction step that extends the chains into complete models and creates full ab initio models, using a hidden Markov model (HMM). If the RefSeq and Gnomon models are predicted to have the same splice pattern, the RefSeq transcripts are given precedence. Gnomon predictions are included in the final set of annotations if they do not share all splice sites with a RefSeq transcript and if they meet certain quality thresholds. Novel genes

identified were remapped to GRCh37 using NCBI’s remap tool (<http://www.ncbi.nlm.nih.gov/genome/tools/remap>); these mappings are listed in Supplemental Table S12.

Segmental duplication annotation

We applied whole-genome assembly comparison (WGAC) and read depth CNV (WSSD) methods to discover SDs in the CHM1.1 reference assembly. For WGAC analysis, we eliminated all repetitive sequences from the assembly as annotated by RepeatMasker, identified alignments > 1 kbp and with higher than 90% identity, and refined alignments into pairwise duplication calls as previously described (Bailey et al. 2001). Duplication and RepeatMasker files are in Supplemental Data: Duplication Analysis.

For read depth CNV analysis, we aligned Illumina whole-genome shotgun (WGS) reads from 11 lanes (SRR642629, SRR642634, SRR642635, SRR642638, SRR642639, SRR642640, SRR642641, SRR642642, SRR642643, SRR642683, SRR642746) with mrsFAST (v. 2.5.0.4) (Hach et al. 2010) and called raw copy number across 1-kbp windows as previously described (Sudmant et al. 2010). From these raw copy number calls, we identified duplications as regions with copy number ≥ 3 and ≥ 10 kbp of non-repeat, non-gap sequence.

Assembly–assembly alignment

We aligned the CHM1_1.1 assembly to the GRCh37 and HuRef assemblies using the two-phase NCBI pipeline. Aligning the two assemblies using BLAST generates the first phase alignments and any locus on the query assembly must have 0 or 1 alignment to the target assembly. Additionally, we use in-database masking through precomputed WindowMasker masked regions. BLAST alignments are then trimmed and post-processed to remove low quality and spurious alignments. Chromosome to chromosome alignments are favored over chromosome to scaffold or scaffold-to-scaffold alignments. Alignments based on common components are then merged into the longest, consistent stretches possible, resulting in a set of alignments called the “Common component set.” We then eliminate the remaining BLAST alignments that are redundant with the common component alignments. The remaining alignments are then merged independently of the common component alignments and redundant alignments are removed. The two alignment sets are then combined into a single set of alignments and then sorted to select the “First pass set,” which are ranked to favor, in order: (1) common component alignments, (2) chromosome to chromosome alignments, (3) alternate to alternate alignments, (4) chromosome to alternate alignments, and (5) count of identities. Finally, only alignments with nonconflicting query/subject ranges are kept for the First Pass set. Conflicting alignments are reserved for evaluation in the “Second Pass.” In order to capture duplicated sequences, we do a “Second Pass” to capture large regions (> 1 kb) within an assembly that have no alignment, or a conflicting alignment, in the First Pass. In the “Second Pass” alignments, a given region in the query assembly can align to more than one region in the target assembly.

Variant analysis

For variant analyses, Illumina reads from CHM1 genomic DNA were mapped to the GRCh37 primary assembly reference using BWA version 0.5.9 (Li and Durbin 2009). Single-nucleotide variants (SNVs) were called using both SAMtools (Li et al. 2009) and VarScan v2.2.9 (Koboldt et al. 2012). Variants were filtered to remove false positives due to alignment and sequencing errors using the values in Supplemental Table S13. The Illumina reads were then aligned to the CHM1_1.1 assembly and variants were

called using the same parameters as above. We overlapped the variants with RefSeq and Gnomon gene annotations as well as SDs (based on WGAC) and RepeatMasker. SNV density per kilobase and transition:transversion ratio (Ts:Tv) were calculated in 1 MB non-overlapping windows using vcftools version 0.1.11 (<http://vcftools.sourceforge.net/>).

BAC end sequence mapping

BAC end sequences from the CH17 BAC library generated from the CHM1 cell line were aligned to the CHM1_1.1, GRCh37 and GRCh38 assemblies and clone placements generated as described in Schneider et al. (2013). BAC end mappings are provided in Supplemental Data: BAC end mapping. On the CHM1_1.1 assembly, the average insert length = 208,638 and the standard deviation = 20,197. On GRCh37, the average insert length = 208,637 and the standard deviation = 20,149. BAC ends from single concordant, single discordant, and multiply mapped clones were evaluated for SD content and overlapped with gene annotation from RefSeq and Gnomon from the CHM1_1.1 assembly.

Clinical allele analysis

We obtained data from the NHGRI GWAS catalog using the UCSC Genome Browser track intersected with the dbSNP137 track. If the risk allele was reported on the negative strand, we were able to use the dbSNP137 information to correctly assign risk alleles to the positive strand. Additionally, we downloaded the vcf file containing the ClinVar data from NCBI. We took the unique union of risk alleles from both sources and remapped them to CHM1_1.1 coordinates using NCBI default remap parameters. If there were two or more locations we chose the preferred mapping or discarded both. We then compared the risk allele at each locus with the CHM1 genotype.

PacBio alignment

We obtained CHM1 reads from the Pacific Biosciences website and aligned them to the CHM1_1.1 assembly using BLASR with the following parameters (-nproc 4 -sam -clipping soft -bestn 2 -minMatch 12 -affineAlign -sortRefinedAlignments). Read statistics are outlined in Supplemental Table S14, and a graph of the read length distribution is in Supplemental Figure S5. To call cliff regions, we required that a PacBio read must have two and only two alignments on CHM1_1.1, both alignments must be on the same CHM1_1.1 sequence, and one of the two alignments must meet the criteria of "Score $\leq -2.0 \times \text{ReadLength}$." We also required query coverage of the smaller of two segments be $\geq 10\%$ and that the smaller alignment must still involve at least 10% of the PacBio read. Two alignments could not overlap each other by $> 10\%$ and unique coverage $\geq 50\%$. Coverage drop-offs that occurred within 1 kbp of a CHM1_1.1 boundary were flagged.

The PacBio reads used for this analysis aligned to the CHM1_1.1 assembly at an average coverage depth of $54\times$. As expected, coverage at regions containing repetitive sequence was notably higher. To improve our likelihood of detecting examples of misassemblies, we restricted our review of this list to sites where surrounding coverage did not indicate the presence of repetitive sequence and the drop-off in coverage was roughly equivalent to surrounding coverage.

Data access

All Illumina sequence, assembly, and clone data have been submitted to the NCBI GenBank (<http://www.ncbi.nlm.nih.gov/genbank>) and Sequence Read Archive (SRA; <http://www.ncbi.nlm.nih.gov/sra>). The

GenBank Assembly ID for the CHM1_1.1 assembly is GCA_000306695.2. Accessions for clones can be found in Supplemental Table S1. Illumina sequencing data can be found under study accession SRP017546, and individual library accessions are listed in Supplemental Table S11. Gene annotation is available from the NCBI genome annotation FTP site (ftp://ftp.ncbi.nlm.nih.gov/genomes/H_sapiens/ARCHIVE/ANNOTATION_RELEASE.105/). PacBio data are available from <http://datasets.pacb.com/2014/Human54x/fast.html> and have also been submitted to SRA under accession number SRX533609. The BioNano Genomics map contigs are available in the Supplemental Material and at http://genome.wustl.edu/pub/supplemental/KMS_genome_research_2014/. Additional analyses and data sets are available in the Supplemental Material and at figshare (<http://dx.doi.org/10.6084/m9.figshare.1091429>).

Competing interest statement

E.E.E. is on the scientific advisory board (SAB) of DNAnexus, Inc. and was an SAB member of Pacific Biosciences, Inc. (2009–2013) and SynapDx Corp. (2011–2013).

Acknowledgments

We thank Pieter de Jong for the creation of the CHORI-17 BAC library used extensively in this project. We acknowledge Nathan Bouk for his expertise in sequence alignment and insightful discussions of alignment data. We acknowledge the efforts of the production and finishing groups at The Genome Institute, particularly Aye Wollom, Susie Rock, Milinn Kremitzki, and Derek Albrecht. We are greatly appreciative of BioNano Genomics and Dr. Pui Kwok for generously sharing the genomic map and alignments to the CHM1_1.1 assembly. E.E.E. is an investigator of the Howard Hughes Medical Institute. This work was supported by NIH Grants 2R01HG002385 and 5P01HG004120 to E.E.E.

References

- The 1000 Genomes Project Consortium 2012. An integrated map of genetic variation from 1,092 human genomes. *Nature* **491**: 56–65.
- Alkan C, Sajjadian S, Eichler EE. 2011. Limitations of next-generation genome sequence assembly. *Nat Methods* **8**: 61–65.
- Bailey JA, Yavor AM, Massa HF, Trask BJ, Eichler EE. 2001. Segmental duplications: organization and impact within the current human genome project assembly. *Genome Res* **11**: 1005–1017.
- Barbouti A, Stankiewicz P, Nusbaum C, Cuomo C, Cook A, Hoglund M, Johansson B, Hagemeijer A, Park SS, Mitelman F, et al. 2004. The breakpoint region of the most common isochromosome, i(17q), in human neoplasia is characterized by a complex genomic architecture with large, palindromic, low-copy repeats. *Am J Hum Genet* **74**: 1–10.
- Bosnakovski D, Xu Z, Gang EJ, Galindo CL, Liu M, Simsek T, Garner HR, Agha-Mohammadi S, Tassin A, Coppee F, et al. 2008. An isogenetic myoblast expression screen identifies DUX4-mediated FSHD-associated molecular pathologies. *EMBO J* **27**: 2766–2779.
- Carvalho CM, Lupski JR. 2008. Copy number variation at the breakpoint region of isochromosome 17q. *Genome Res* **18**: 1724–1732.
- Chen R, Butte AJ. 2011. The reference human genome demonstrates high risk of type 1 diabetes and other disorders. *Pac Symp Biocomput* **2011**: 231–242.
- Church DM, Schneider VA, Graves T, Auger K, Cunningham F, Bouk N, Chen HC, Agarwala R, McLaren WM, Ritchie GR, et al. 2011. Modernizing reference genome assemblies. *PLoS Biol* **9**: e1001091.
- Dennis MY, Nettle X, Sudmant PH, Antonacci F, Graves TA, Nefedov M, Rosenfeld JA, Sajjadian S, Malig M, Kotkiewicz H, et al. 2012. Evolution of human-specific neural SRGAP2 genes by incomplete segmental duplication. *Cell* **149**: 912–922.
- Doggett NA, Xie G, Meincke LJ, Sutherland RD, Mundt MO, Barbani NS, Davy BE, Robinson ML, Rudd MK, Weber JL, et al. 2006. A 360-kb interchromosomal duplication of the human HYDIN locus. *Genomics* **88**: 762–771.

- Eichler EE, Clark RA, She X. 2004. An assessment of the sequence gaps: unfinished business in a finished human genome. *Nat Rev Genet* **5**: 345–354.
- Fan JB, Surti U, Taillon-Miller P, Hsie L, Kennedy GC, Hoffner L, Ryder T, Mutch DG, Kwok PY. 2002. Paternal origins of complete hydatidiform moles proven by whole genome single-nucleotide polymorphism haplotyping. *Genomics* **79**: 58–62.
- Genovese G, Handsaker RE, Li H, Altemose N, Lindgren AM, Chambert K, Pasaniuc B, Price AL, Reich D, Morton CC, et al. 2013. Using population admixture to help complete maps of the human genome. *Nat Genet* **45**: 406–414.
- Gnerre S, Maccallum I, Przybylski D, Ribeiro FJ, Burton JN, Walker BJ, Sharpe T, Hall G, Shea TP, Sykes S, et al. 2011. High-quality draft assemblies of mammalian genomes from massively parallel sequence data. *Proc Natl Acad Sci* **108**: 1513–1518.
- Hach E, Hormozdiari F, Alkan C, Hormozdiari F, Birol I, Eichler EE, Sahinalp SC. 2010. mrsFAST: a cache-oblivious algorithm for short-read mapping. *Nat Methods* **7**: 576–577.
- Hewitt JE, Lyle R, Clark LN, Valleley EM, Wright TJ, Wijmenga C, van Deutekom JC, Francis F, Sharpe PT, Hofker M, et al. 1994. Analysis of the tandem repeat locus D4Z4 associated with facioscapulohumeral muscular dystrophy. *Hum Mol Genet* **3**: 1287–1295.
- Hoffner L, Surti U. 2012. The genetics of gestational trophoblastic disease: a rare complication of pregnancy. *Cancer Genet* **205**: 63–77.
- International Human Genome Sequencing. 2004. Finishing the euchromatic sequence of the human genome. *Nature* **431**: 931–945.
- Jiang Z, Tang H, Ventura M, Cardone MF, Marques-Bonet T, She X, Pevzner PA, Eichler EE. 2007. Ancestral reconstruction of segmental duplications reveals punctuated cores of human genome evolution. *Nat Genet* **39**: 1361–1368.
- Kidd JM, Cooper GM, Donahue WF, Hayden HS, Sampas N, Graves T, Hansen N, Teague B, Alkan C, Antonacci F, et al. 2008. Mapping and sequencing of structural variation from eight human genomes. *Nature* **453**: 56–64.
- Kidd JM, Sampas N, Antonacci F, Graves T, Fulton R, Hayden HS, Alkan C, Malig M, Ventura M, Giannuzzi G, et al. 2010. Characterization of missing human genome sequences and copy-number polymorphic insertions. *Nat Methods* **7**: 365–371.
- Koboldt DC, Zhang Q, Larson DE, Shen D, McLellan MD, Lin L, Miller CA, Mardis ER, Ding L, Wilson RK. 2012. VarScan 2: somatic mutation and copy number alteration discovery in cancer by exome sequencing. *Genome Res* **22**: 568–576.
- Korbel JO, Urban AE, Affourtit JP, Godwin B, Grubert F, Simons JF, Kim PM, Palejev D, Carriero NJ, Du L, et al. 2007. Paired-end mapping reveals extensive structural variation in the human genome. *Science* **318**: 420–426.
- Kyo K, Muto T, Nagawa H, Lathrop GM, Nakamura Y. 2001. Associations of distinct variants of the intestinal mucin gene MUC3A with ulcerative colitis and Crohn's disease. *J Hum Genet* **46**: 5–20.
- Landrum MJ, Lee JM, Riley GR, Jang W, Rubinstein WS, Church DM, Maglott DR. 2014. ClinVar: public archive of relationships among sequence variation and human phenotype. *Nucleic Acids Res* **42**: D980–D985.
- Li H. 2014. Towards better understanding of artifacts in variant calling from high-coverage samples. *Bioinformatics* **30**: 2843–2851.
- Li H, Durbin R. 2009. Fast and accurate short read alignment with Burrows-Wheeler transform. *Bioinformatics* **25**: 1754–1760.
- Li R, Li Y, Kristiansen K, Wang J. 2008. SOAP: short oligonucleotide alignment program. *Bioinformatics* **24**: 713–714.
- Li H, Handsaker B, Wysoker A, Fennell T, Ruan J, Homer N, Marth G, Abecasis G, Durbin R, 1000 Genome Project Data Processing Subgroup. 2009. The Sequence Alignment/Map format and SAMtools. *Bioinformatics* **25**: 2078–2079.
- Li R, Li Y, Zheng H, Luo R, Zhu H, Li Q, Qian W, Ren Y, Tian G, Li J, et al. 2010. Building the sequence map of the human pan-genome. *Nat Biotechnol* **28**: 57–63.
- McPherson JD, Marra M, Hillier L, Waterston RH, Chinwalla A, Wallis J, Sekhon M, Wylie K, Mardis ER, Wilson RK, et al. 2001. A physical map of the human genome. *Nature* **409**: 934–941.
- Mills RE, Luttig CT, Larkins CE, Beauchamp A, Tsui C, Pittard WS, Devine SE. 2006. An initial map of insertion and deletion (INDEL) variation in the human genome. *Genome Res* **16**: 1182–1190.
- Morgulis A, Gertz EM, Schaffer AA, Agarwala R. 2006. WindowMasker: window-based masker for sequenced genomes. *Bioinformatics* **22**: 134–141.
- Ogino S, Wilson RB, Gold B. 2004. New insights on the evolution of the SMN1 and SMN2 region: simulation and meta-analysis for allele and haplotype frequency calculations. *Eur J Hum Genet* **12**: 1015–1023.
- Pan Q, Tian Y, Li X, Ye J, Liu Y, Song L, Yang Y, Zhu R, He Y, Chen L, et al. 2013. Enhanced membrane-tethered mucin 3 (MUC3) expression by a tetrameric branched peptide with a conserved TFLK motif inhibits bacteria adherence. *J Biol Chem* **288**: 5407–5416.
- Paten B, Novak A, Haussler D. 2014. Mapping to a reference genome structure. arXiv:1404.5010.
- Pratt WS, Crawley S, Hicks J, Ho J, Nash M, Kim YS, Gum JR, Swallow DM. 2000. Multiple transcripts of MUC3: evidence for two genes, MUC3A and MUC3B. *Biochem Biophys Res Commun* **275**: 916–923.
- Robledo R, Orru S, Sidoti A, Muresu R, Esposito D, Grimaldi MC, Carcassi C, Rinaldi A, Bernini L, Contu L, et al. 2002. A 9.1-kb gap in the genome reference map is shown to be a stable deletion/insertion polymorphism of ancestral origin. *Genomics* **80**: 585–592.
- Ryan DP, da Silva MR, Soong TW, Fontaine B, Donaldson MR, Kung AW, Jongjaroenprasert W, Liang MC, Khoo DH, Cheah JS, et al. 2010. Mutations in potassium channel Kir2.6 cause susceptibility to thyrotoxic hypokalemic periodic paralysis. *Cell* **140**: 88–98.
- Schmutz J, Martin J, Terry A, Couronne O, Grimwood J, Lowry S, Gordon LA, Scott D, Xie G, Huang W, et al. 2004. The DNA sequence and comparative analysis of human chromosome 5. *Nature* **431**: 268–274.
- Schneider VA, Chen HC, Clausen C, Meric PA, Zhou Z, Bouk N, Husain N, Maglott DR, Church DM. 2013. Clone DB: an integrated NCBI resource for clone-associated data. *Nucleic Acids Res* **41**: D1070–D1078.
- Sharp AJ, Locke DP, McGrath SD, Cheng Z, Bailey JA, Vallente RU, Pertz LM, Clark RA, Schwartz S, Segraves R, et al. 2005. Segmental duplications and copy-number variation in the human genome. *Am J Hum Genet* **77**: 78–88.
- Smit A, Hubley R, Green P. 1996–2010. RepeatMasker Open-2.0.
- Sudmant PH, Kitzman JO, Antonacci F, Alkan C, Malig M, Tsalenko A, Sampas N, Bruhn L, Shendure J, Genomes P, et al. 2010. Diversity of human copy number variation and multicopy genes. *Science* **330**: 641–646.
- Sudmant PH, Huddleston J, Catacchio CR, Malig M, Hillier LW, Baker C, Mohajer K, Kondova I, Bontrop RE, Persengiev S, et al. 2013. Evolution and diversity of copy number variation in the great ape lineage. *Genome Res* **23**: 1373–1382.
- Taillon-Miller P, Bauer-Sardina I, Zakeri H, Hillier L, Mutch DG, Kwok PY. 1997. The homozygous complete hydatidiform mole: a unique resource for genome studies. *Genomics* **46**: 307–310.
- Teague B, Waterman MS, Goldstein S, Potamouis K, Zhou S, Reslewic S, Sarkar D, Valouev A, Churas C, Kidd JM, et al. 2010. High-resolution human genome structure by single-molecule analysis. *Proc Natl Acad Sci* **107**: 10848–10853.
- Watson CT, Steinberg KM, Huddleston J, Warren RL, Malig M, Schein J, Willsey AJ, Joy JB, Scott JK, Graves TA, et al. 2013. Complete haplotype sequence of the human immunoglobulin heavy-chain variable, diversity, and joining genes and characterization of allelic and copy-number variation. *Am J Hum Genet* **92**: 530–546.
- Zody MC, Jiang Z, Fung HC, Antonacci F, Hillier LW, Cardone MF, Graves TA, Kidd JM, Cheng Z, Abouelleil A, et al. 2008. Evolutionary toggling of the MAPT 17q21.31 inversion region. *Nat Genet* **40**: 1076–1083.

Received July 2, 2014; accepted in revised form October 7, 2014.

Single haplotype assembly of the human genome from a hydatidiform mole

Karyn Meltz Steinberg, Valerie A. Schneider, Tina A. Graves-Lindsay, et al.

Genome Res. 2014 24: 2066-2076 originally published online November 4, 2014

Access the most recent version at doi:[10.1101/gr.180893.114](https://doi.org/10.1101/gr.180893.114)

Supplemental Material <http://genome.cshlp.org/content/suppl/2014/10/17/gr.180893.114.DC1>

References This article cites 47 articles, 11 of which can be accessed free at:
<http://genome.cshlp.org/content/24/12/2066.full.html#ref-list-1>

Creative Commons License This article is distributed exclusively by Cold Spring Harbor Laboratory Press for the first six months after the full-issue publication date (see <http://genome.cshlp.org/site/misc/terms.xhtml>). After six months, it is available under a Creative Commons License (Attribution-NonCommercial 4.0 International), as described at <http://creativecommons.org/licenses/by-nc/4.0/>.

Email Alerting Service Receive free email alerts when new articles cite this article - sign up in the box at the top right corner of the article or [click here](#).

Affordable, Accurate
Sequencing.

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