

# *All of Us* Research Program

## Genomic Research Data Quality Report

*All of Us* Curated Data Repository (CDR) release C2022Q4R9

<b>Overview</b>	<b>4</b>
<b>Executive Summary</b>	<b>6</b>
<b>Introduction</b>	<b>7</b>
<b>Arrays</b>	<b>8</b>
Consistency across Genome Centers	8
Single Sample QC	8
Sex Concordance	9
Method	9
Results	10
Call Rate	10
Method	10
Results	10
Cross-Individual Contamination Rate	11
Method	12
Results	12
<b>Short Read Whole Genome Sequencing (srWGS)</b>	<b>13</b>
Consistency across Genome Centers	13
Single Sample QC	13
Fingerprint Concordance	14
Method	14
Results	15
Sex Concordance	15
Method	15
Results	16
Cross-Individual Contamination Rate	16
Method	16
Results	17
Coverage	18
Method	18
Results	19
Short-read WGS SNP & Indel Joint Callset QC	19
Sample Hard Threshold Flag	20
Method	21

Results	21
Sample Population Outlier Flag	21
Method	21
Results	21
Variant Hard Threshold Filters	22
Method	22
Results	23
Allele-Specific Variant Quality Score Recalibration (AS-VQSR)	23
Method	23
Sensitivity and Precision Evaluation	24
Method	24
Results	25
<b>Short Read Structural Variants (srWGS SVs)</b>	<b>26</b>
<b>Long Reads</b>	<b>27</b>
Data generation for IrWGS	27
Single Sample QC for IrWGS	27
Fingerprint Concordance	29
Method	29
Results	29
Sex Concordance	30
Method	30
Results	30
Cross-individual Contamination Rate	31
Method	31
Results	31
Coverage	32
Method	32
Results	32
Read Length Median	33
Method	33
Results	33
Outlier Sample Filtering	34
Method	34
Results	34
Joint Callset QC for the IrWGS SNP/Indel callsets	35
Relatedness	36
Method	36
Results	36
Sample Population Outlier	36
Method	36
Results	36

<b>Known Issues</b>	<b>38</b>
Known Issue #1: Small subset of samples missing corresponding CDR data	38
Known Issue #2: 11 samples were affected by a sample swap incident	38
Known issue #3: Array samples (N=416) from previous release are missing in this release	39
Known Issue #4: Single array sample missing from Array Hail MT and PLINK files	39
Known Issue #5: Larger than expected changes in ancestry predictions from previous release	40
Known issue #6: Ancestry prediction has higher error rates for Middle Eastern ancestry	40
Known Issue #7: VDS issue with GT	41
Known Issue #8: AS_VQSLOD is incorrect in the VDS and dropped in callset data with all participants	41
Known Issue #9: Smaller callset ChrM VCFs are empty	42
Known issue #10: srWGS SNP & Indel VDS and VCFs from the Cohort Browser will have extraneous INFO field (AS_YNG)	42
Known issue #11: srWGS SNP & Indel variant calls on chromosome Y need additional filtering	42
Known Issue #12: QUAL information has been removed for srWGS SNP & Indel variants	43
Known Issue #13: srWGS callset using new convention for genotype filtering flag	43
Known Issue #14: Data processing issue affecting array data	44
Known Issue #15: Array and srWGS data with bone marrow transplant history	44
<b>FAQ</b>	<b>45</b>
<b>References</b>	<b>46</b>
<b>Appendix A: Ancestry</b>	<b>49</b>
<b>Appendix B: Self-reported race/ethnicity</b>	<b>53</b>
<b>Appendix C: Data type availability with genomic data</b>	<b>54</b>
<b>Appendix D: Genome Centers and Data and Research Center</b>	<b>56</b>
<b>Appendix E: Array processing overview</b>	<b>57</b>
<b>Appendix F: Self-reported sex assigned at birth</b>	<b>60</b>
<b>Appendix G: All of Us Hereditary Disease Risk genes</b>	<b>61</b>
<b>Appendix H: DRAGEN invocation parameters</b>	<b>62</b>
<b>Appendix I: Samples used in the Sensitivity and Precision Evaluation</b>	<b>64</b>
<b>Appendix J: High quality site determination (srWGS)</b>	<b>65</b>
<b>Appendix K: Relatedness (srWGS)</b>	<b>66</b>
<b>Appendix L: Plots of the first principal component against population outlier QC metrics</b>	<b>67</b>
<b>Appendix M: Long Read Workflow Diagrams</b>	<b>70</b>
<b>Appendix N: IrWGS analysis versions and parameters</b>	<b>74</b>

# Overview

This document details the *All of Us* Genome Centers (GC) and Data and Research Center (DRC) quality control (QC) steps for genomic data in the research pipeline. This pipeline removes or flags samples and variants in the genomic data that fail quality thresholds. We apply these QC steps in the research pipeline before we release the genomic data for research use. We, the *All of Us* DRC, only describe QC processes that are performed analytically (i.e., after the sample has been genotyped and sequenced). All descriptions and results are limited to the v7 data release made available in the Researcher Workbench April 20, 2023, which contains 312,945 genotyping array (“array”) samples, 245,394 short read whole genome sequencing (srWGS) samples with single nucleotide polymorphism, insertion, and deletion variant calls (SNPs and Indels), 97,940 srWGS samples with structural variant (SV) calls, and 1,027 long read whole genome sequencing (lrWGS) samples with SNP, Indel, and SV calls. The srWGS SV samples and lrWGS samples are a subset of the srWGS SNP and Indel samples, which in turn are a subset of the array data. The samples in the genomic data correspond to the *All of Us* Curated Data Repository (CDR) release C2022Q4R9 (“v7”), though please see [Known Issue #1](#), as 20 array samples (less than 0.01%), six srWGS samples (less than 0.01%), and three srWGS SV samples (less than 0.01%) are missing their corresponding CDR data. These pipelines are automated unless otherwise noted. This document covers the genomic data types including small variants (SNPs and Indels), structural variants, raw data, and auxiliary data. Small variants are available for array samples, srWGS samples, and lrWGS samples. Structural variants are available for srWGS samples and lrWGS samples. **Please note: As of June 17, 2024, we have released an additional set of v7 srWGS SVs, increasing the srWGS samples that have SV data to 40%. The QC steps for these 97,940 srWGS samples with SV calls is available as a [separate article](#).**

Audience: This document is intended for researchers using, or considering the use of, the genomic data in the Researcher Workbench (RW). This document assumes knowledge of sequencing, genotype arrays, common genomic data QC approaches, and the variant file formats released in *All of Us*. We recommend that at a minimum researchers read the [Known Issues](#) and the [FAQ](#) section below, even if they are not as concerned with the QC process.

## Notes:

- Details of the processing (e.g., algorithms) are out of scope for this document.
- The locations of raw data are in the ‘[Controlled CDR directory document](#)’, published on the User Support Hub [\[1\]](#). Auxiliary data sample lists are also published on the User Support Hub.
- The genomic data mentioned in this document requires Controlled Tier access to view. To register for access, please go to <https://www.researchallofus.org/register/>
- A small number of array and srWGS SNP & Indel samples are missing their corresponding CDR and Cohort Builder data and thus the sample counts are not the same. Please see [Known Issue #1](#) for more details.



# Executive Summary

On April 20, 2023, the *All of Us* Research Program released the genomic data of 312,945 array samples, 245,394 srWGS samples, and 1,027 lrWGS samples in the Researcher Workbench (RW) for use by researchers registered for Controlled Tier access. There are over 1.8M array SNP and Indel sites, over 1.03B srWGS SNP and Indel sites, over 64 million long read SNP and Indel sites on grch38\_noalt, and over 73 million long read SNP and Indel sites on T2T-CHM13v2.0. There are over 1,506,805 SVs called on the 97,940 srWGS sample cohort and SVs called for each lrWGS sample. In addition to variant calls, raw data (IDAT files for array data, CRAM files for srWGS data, BAM files for lrWGS data) and auxiliary files (predicted ancestry, relatedness/kinship scores, functional annotation, and flagged samples) are available in the RW through Controlled Tier access. Quality control processes, performed both independently and across samples, indicate that these data are ready for general analysis. We suggest researchers, at a minimum, read the [Known Issues](#) and [FAQ](#) sections below before using the data.

# Introduction

*All of Us* is collecting biospecimens and generating genomic data for all participants who have consented among its target of 1,000,000 participants. As the program continues, the DRC will periodically release genomic data - in sync with planned CDR release timelines. This document describes the third release of genomic data to *All of Us* researchers (v7) made available in the RW on April 20, 2023. The genomic data contains 312,945 array samples, 245,394 srWGS samples, 97,940 srWGS samples with SV calls, and 1,027 lrWGS samples from a diverse set of participants (see [Appendix A](#) and [Appendix B](#)). Genomic data can be joined with other data types (e.g. survey data) for analysis ([Appendix C](#)), though please see [Known Issue #1](#). In this document, we describe the QC processes applied to the array, srWGS, and lrWGS data. We describe which processes were performed at the GCs and which were performed at the DRC (see [Appendix D](#)), but for most researchers this demarcation has no practical significance.

This document is organized by data type and describes the QC processes performed. For each data type, we will outline the consistency, single sample QC, and joint callset QC.

1. Consistency is the uniformity of protocols at each GC that reduce the probability of batch effects and normalize the data across GCs. Descriptions in this document, for both QC and sample processing, apply to all GCs unless otherwise noted.
2. Single sample QC are the QC processes for each sample independently to catch major errors. If a sample fails these tests, it is excluded from the release and not reported in this document. We also use these tests to confirm internal consistency between the GCs and the DRC. These tests detect sample swaps, cross-individual contamination, and sample preparation errors.
3. Joint callset QC are the processes executed on the joint callset, which use information across samples to flag samples and variants. The QC steps are performed after single sample QC, during creation of the joint callset. The flagged samples and variants are not removed from the callset unless otherwise specified.

We have also performed data validation experiments, such as replicating GWAS results, but the results are shown in other, upcoming documentation (see the User Support Hub [\[1\]](#) or Tutorial Workspaces in the RW, the RW requires authorization to access).

# Arrays

There are 312,945 array samples in the v7 release. The SNP and Indel variants from array samples are available in VCF, Hail, and PLINK formats. In addition, raw Array data is available in IDAT format. The data is described in the [‘How the All of Us Genomic data are organized’](#) article on the User Support Hub [1]. The QC process for array data includes consistency and single sample QC steps. Array data is not joint-called so no joint callset QC was performed.

## Consistency across Genome Centers

The genome centers (GCs) established a consistent sample and data processing protocol for array data generation to attenuate the likelihood of batch effects across GCs.

The GCs generate variant calls (VCFs) that are submitted to the DRC. The GCs use the same lab protocols, scanners, software, and input files:

- GCs generate raw intensity data (.idat) using the same hardware (iSCAN scanners from Illumina). These files will still contain biases across GCs.
- GCs normalize the raw intensity data onto the same scale. This process yields a normalization transform for probe intensities, which are one of the inputs for variant calls. The array cluster definition file (.egt) was updated between this release and the prior release. This update was done to reduce variation across GCs. Each GC used the newly defined clusters to generate variant calls as well as reprocessing array samples from the prior release.
- GCs use identical pipelines to generate VCFs, including identical pipeline versions and input parameters, where applicable. As a result, the VCFs contain the same information, regardless of GC, including metadata about inputs.

Please see [Appendix E](#) for details.

## Single Sample QC

For array samples, we perform sex concordance, call rate tests, and test cross-individual contamination. These tests are designed to detect sample swaps and sample preparation errors and are performed at the GCs. The list of specific QC processes and an overview of the results can be found in [Table 1](#). Some srWGS QC processes, such as [Fingerprint Concordance](#), use array data.

For more details about the array single sample QC process, including preparation, see [Appendix E](#).

**Table 1 -- Array Single Sample QC processes**

QC process	Passing criteria	Error modes addressed	v7 release results
------------	------------------	-----------------------	--------------------



Sex concordance	Sex call is concordant with self-reported sex at birth. OR Self-reported sex at birth reported as "Other" or was not reported	-Sample swaps	All array samples are concordant.  *Other refers to a participant self-reporting "Intersex", "I prefer not to answer", or "none of these fully describe me"
Call rate	> 0.98 (> 98%)	-Sample contamination -Sample preparation error	All array samples meet the threshold.
Cross-individual contamination rate	No passing criteria	-Sample contamination from another individual	For arrays, we only report the contamination rate, but do not filter array samples, since the call rate is a proxy for high levels of contamination.

## Sex Concordance

We checked the computed sex against the self-reported sex assigned at birth for concordance. We used gencall to determine the computed sex and CDR data for the self-reported sex assigned at birth ([Appendix F](#)). If the two sources were not concordant, we assumed a potential sample swap, removed the sample, and investigated the source of the swap.

### Method

We call the gencall tool [\[2\]](#) v3.0.0 to make a call on the sex of the sample from the array data. We use the Picard 2.26.0 tool, CollectArraysVariantCallingMetrics [\[3\]](#), to perform the actual concordance check against the self-reported sex assigned at birth. If we do not have a "male" or "female" for the sex assigned at birth, because the participant reported it as "Intersex", "I prefer not to answer", "none of these fully describe me", or skipped the question, we passed the sex concordance check for that sample, regardless of the information from gencall. The sex assigned at birth data from the CDR is described in [Appendix F](#).

To generate sex calls from the array, we call gencall from the Illumina Array Analysis Platform Genotyping Command Line Interface (iaap-cli):

Parameter	Value	Notes
Tool name	"gencall"	
Manifest file	Bead pool manifest (BPM)	Illumina-supplied file that contains metadata (alleles, mapping information, source, etc.) for all of the probes on the genotyping array.
Cluster file	Cluster file (EGT)	Used for normalization of intensities across GCs
-f	Location of the IDAT (.idat) files	
-i	"1"	Algorithm version

<code>--gender-estimate-call-rate-threshold</code>	<code>-0.1</code>	This effectively disables the sex estimation.
--	-------------------	---

To ensure concordance with the self-reported sex assigned at birth, we call `CollectArraysVariantCallingMetrics` with the following parameters from the Picard toolkit:

Parameter	Value
Tool name	"CollectArraysVariantCallingMetrics"
INPUT	Array single sample VCF
DBSNP	"gs://gcp-public-data--broad-references/hg38/v0/Homo_sapiens_assembly38.db_snp138.vcf"

## Results

Since we catch sex concordance failures before including a sample in the release, all array samples in the v7 release passed a sex concordance check. Note that 2.09% of array samples passed the sex concordance check solely because they did not answer "male" or "female" on the self-reported sex assigned at birth question. [Appendix F](#) has more details on this CDR question and responses.

## Call Rate

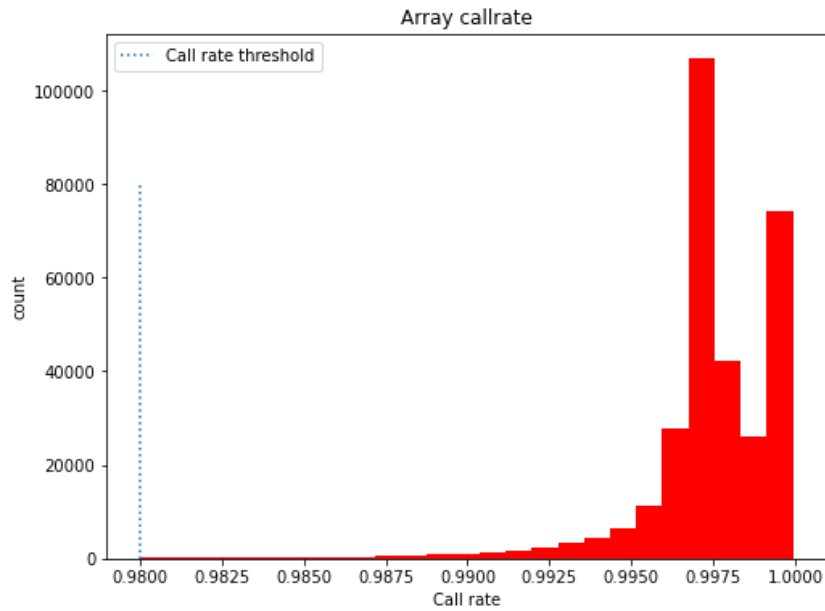
### Method

The call rate is the number of successful variant calls divided by the number of probes. We invoke the gencall tool [\[2\]](#) v3.0.0, as described above in the [Sex Concordance](#) QC process. The gencall tool generates both sex calls and the call rate. We also invoke `CollectArraysVariantCallingMetrics` with the same parameters as the above section to extract the call rate metric from the VCF header.

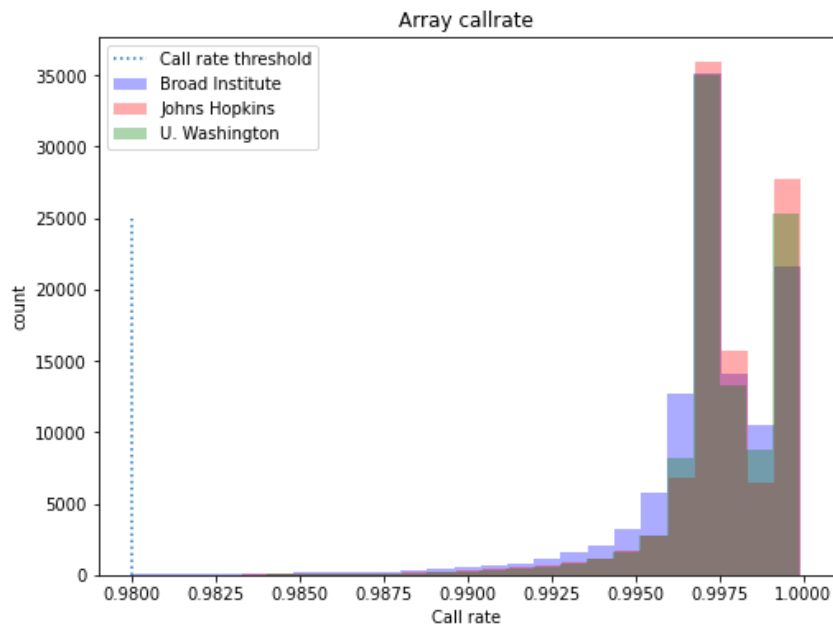
We applied a threshold of 0.98 to the call rate for inclusion in the v7 release.

### Results

As seen in [Figure 1](#), we did not include any samples that were below the call rate threshold of 0.98. See [Figure 2](#) for cross-GC call rate frequencies. Please note that differences in call rates between males and females will cause a double peak in call rate frequencies, since sites on chrY will have a lower call rate for females.



**Figure 1** -- Histogram of the array call rate for the v7 release.



**Figure 2** -- Call rate across each GC.

### Cross-Individual Contamination Rate

For all samples, we estimate the proportion of data coming from an individual other than the one being processed, referred to as the contamination rate. For array samples, as the contamination rate increases, we expect a lower call rate. We fail array samples for a call rate that does not meet the threshold.

## Method

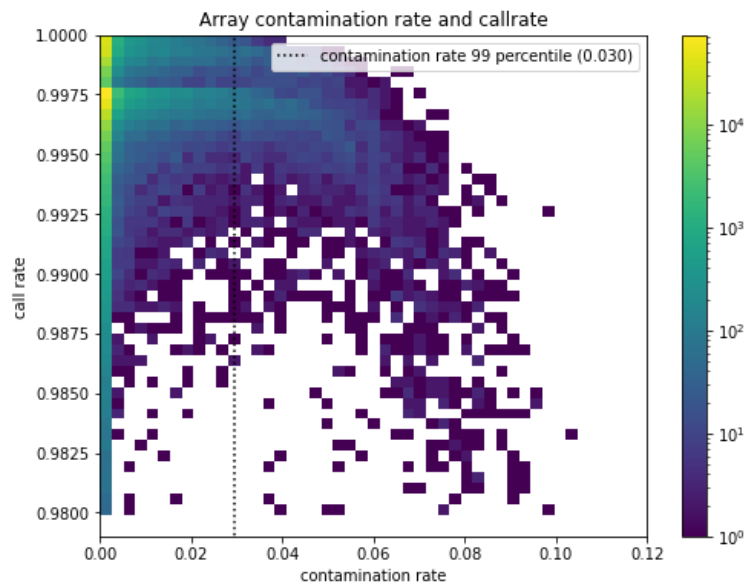
We use BAFRegress [4] to estimate the contamination rate in our array data. We do not use the cross-individual contamination rate to filter array samples, and we do not process the corresponding WGS aliquots for any array sample with a contamination greater than 10%. We filter samples based on the call rate, which is a proxy for contamination and other errors, such as sample preparation errors. Note that most samples with a contamination rate greater than 10% will also not meet the call rate threshold.

We extract allele frequency information from the array VCF and convert it into the file format expected by BAFRegress. We then invoke BAFRegress with the following parameters:

Parameter	Value
task	"estimate"
freqfile	Allele frequency information for all sites, which was extracted from the single sample array VCF.

## Results

We estimated the contamination rate below 0.11 for all array samples. As the contamination rate increased, we did see a small decrease in the call rate (see Figure 3). Of the 312,945 array samples, 309,972 (99.0%) had an estimated contamination rate below 3.0% and 303,044 (96.8%) had a contamination rate less than 1%.



**Figure 3** -- Histogram of the array contamination rate estimates vs call rate. As the contamination rate increases, the call rate decreases.

# Short Read Whole Genome Sequencing (srWGS)

## Consistency across Genome Centers

The GCs use the same protocol for library construction (PCR Free Kapa HyperPrep), sequencer (NovaSeq 6000), software (DRAGEN v3.4.12), and software configuration. The software produces the metrics that are consumed by the sample QC processes. For more information about the sequencing processes used by the GCs, see previous work [\[5\]](#) and the NIH *All of Us* Research Program’s Return of Genetic Results FDA IDE (G200165).

## Single Sample QC

The list of specific QC processes for srWGS samples and an overview of the results can be found in [Table 2](#). Our WGS single sample QC uses the same sequencing process described previously [\[5\]](#) and in the NIH *All of Us* Research Program’s Return of Genetic Results FDA IDE (G200165). Most thresholds in our single sample QC process are identical to the clinical pipeline described previously [\[5\]](#), except for a higher contamination rate.

In some cases, we perform these tests twice for two reasons: 1) to confirm internal consistency between the GCs and the DRC and 2) to mark samples as passing (or failing) QC based on the research pipeline criteria. In this document, we are focused on downstream analytical QC processes after a sample has been sequenced, thus, there are some upstream processes not described here. The list of specific QC processes and an overview of the results can be found in [Table 2](#).

**Table 2 -- srWGS Single Sample QC processes**

QC process	Calculated at the DRC or GCs?	Passing criteria	Error modes addressed	v7 release results
Fingerprint concordance	Both	log-likelihood ratio > -3	-Sample swaps -Large amount of sample contamination	All srWGS samples are concordant with array samples.
Sex concordance	Both	Sex call is concordant with self-reported sex at birth. OR Self-reported sex at birth reported as “Other” or was not reported	-Sample swaps	All srWGS samples are concordant.  *Other refers to a participant self-reporting “Intersex”, “I prefer not to answer”, or “none of these fully describe me”
Cross-individual contamination rate	Both	< 0.03 (< 3%)	Sample contamination from another individual	All srWGS samples meet the threshold.  srWGS samples with corresponding arrays that have a contamination rate above

				10% were not released.
Coverage	GCs only	<p>≥ 30x mean coverage</p> <p>≥ 90% of bases at 20x coverage</p> <p>≥8e10 aligned Q30 Bases</p> <p>≥ 95% at 20x in regions of the 59 AoU Hereditary Disease Risk genes (AoUHDR) See <a href="#">Appendix G</a> for more information</p>	<p>-Sample preparation error</p> <p>-Poor sensitivity and precision of variant calling</p>	All srWGS samples meet the thresholds.

## Fingerprint Concordance

### Method

We filter variant calls to 113 sites (“fingerprint”) for both the array and srWGS SNP & Indel variants. We measure the concordance between the array and WGS data, using a log-likelihood ratio (fingerprint LOD) based on reads. We chose the threshold value, -3.0, to split a bimodal distribution (not shown). If the calls are not concordant (i.e., the fingerprint LOD does not meet the threshold), then there has likely been a sample processing error. A detailed description of fingerprint concordance is described in the Genome Analysis Toolkit documentation [\[6\]](#).

Note: \*One GC (Broad Institute) performed an internal check against a different fingerprint (Fluidigm SNP genotyping (SNPtype chemistry) using the 96.96 Dynamic Array), which did not use the same fingerprint sites as the array. The DRC treated these samples the same as from the other GCs and ran the array concordance as described in the main text of this document.

We call the fingerprint concordance tool “CheckFingerprint” using Picard (version 2.23.9) with the following parameters:

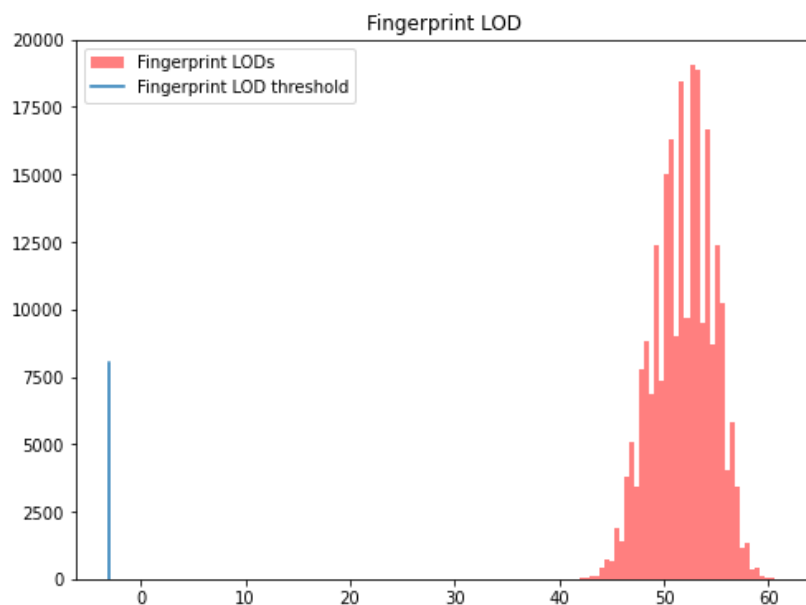
Parameter	Value
program name	“CheckFingerprint”
INPUT	The WGS cram to check concordance
REFERENCE_SEQUENCE	“gs://gcp-public-data--broad-references/hg38/v0/Homo_sapiens_assembly38.fasta”
GENOTYPES	VCF from corresponding array file
HAPLOTYPE_MAP	“gs://gcp-public-data--broad-references/hg38/v0/aou/fp/aou.fp.haplotype_database.txt”
IGNORE_READ_GROUPS	“true”

SAMPLE_ALIAS	Chipwell barcode from the header of the array file (array file passed in the GENOTYPES parameter)
--------------	---

Note: Quoted parameters are exact values, but quotes were not included in the actual call to the tool.

## Results

All samples in the v7 release passed the fingerprint concordance check based on arrays. As seen in [Figure 4](#), the passing samples exceeded the threshold. 1490 samples had a fingerprint LOD [\[6\]](#) less than 45 and the minimum fingerprint LOD was 13.



**Figure 4** -- Distribution of the Fingerprint LODs for srWGS v7 samples

## Sex Concordance

For srWGS data, we compared the computed sex from DRAGEN ([Appendix H](#)) and peddy [\[7\]](#) against the self-reported sex assigned at birth ([Appendix F](#)). If the two sources were not concordant, we assumed a potential sample swap, removed the sample, and investigated the source of the swap.

## Method

We compared variant and ploidy calls for chromosome X and Y against the self-reported sex assigned at birth for the sample. We check the sex ploidy call (e.g., XY or XX) from the DRAGEN pipeline (v 3.4.12) and use heterozygous chrX variant calls from peddy [\[7\]](#). If the concordance test fails against either of these calls, the sample fails QC and is not included in the release. If we do not have a “male” or “female” for the sex assigned at birth, because the participant reported it as “Intersex”, “I prefer not to answer”, “none of these fully describe me”, or skipped the question, we passed the sex concordance check for that sample, regardless of the

information from peddy and DRAGEN. The sex assigned at birth data from the CDR is described in [Appendix F](#).

DRAGEN invocations include a wide breadth of functionality, including ploidy calls (see [Appendix H](#) for the parameters).

The DRAGEN pipeline outputs a single sample VCF, which is primarily used in the clinical pipeline (for individual samples)[\[5\]](#), but we use it for our call to peddy. We call peddy with the following parameters:

Parameter	Value
vcf	Single sample VCF from DRAGEN (hard-filtered)
Pedigree file	We create this file dynamically based on the single sample and its sex call. Please note: This implies that we do not use pedigree information in our peddy call.

## Results

We do not include any srWGS samples that fail the sex concordance check in the released samples. It is important to note that some samples automatically passed this check solely because they did not answer “male” or “female” on the self-reported sex assigned at birth question (2.07% of srWGS samples). [Appendix F](#) has more details on this CDR question and the possible responses.

## Cross-Individual Contamination Rate

For all srWGS samples, we estimate the proportion of data coming from an individual other than the one being processed, referred to as the contamination rate.

## Method

We estimate the percent contamination from another individual by counting the number of reads at common homozygous alternate SNP sites. If there is a small amount of cross-individual contamination, we expect to see small numbers of reads supporting SNPs at these sites. We determine the percentage of the sample that may have come from a different individual using VerifyBamID2 [\[8\]](#), and the DRAGEN 3.4.12 pipeline. Contamination rate is a float value from 0.0 to 1.0, which represents 0 to 100%.

We use the following parameters for VerifyBamID2:

Parameter	Value
NumPC	“4”



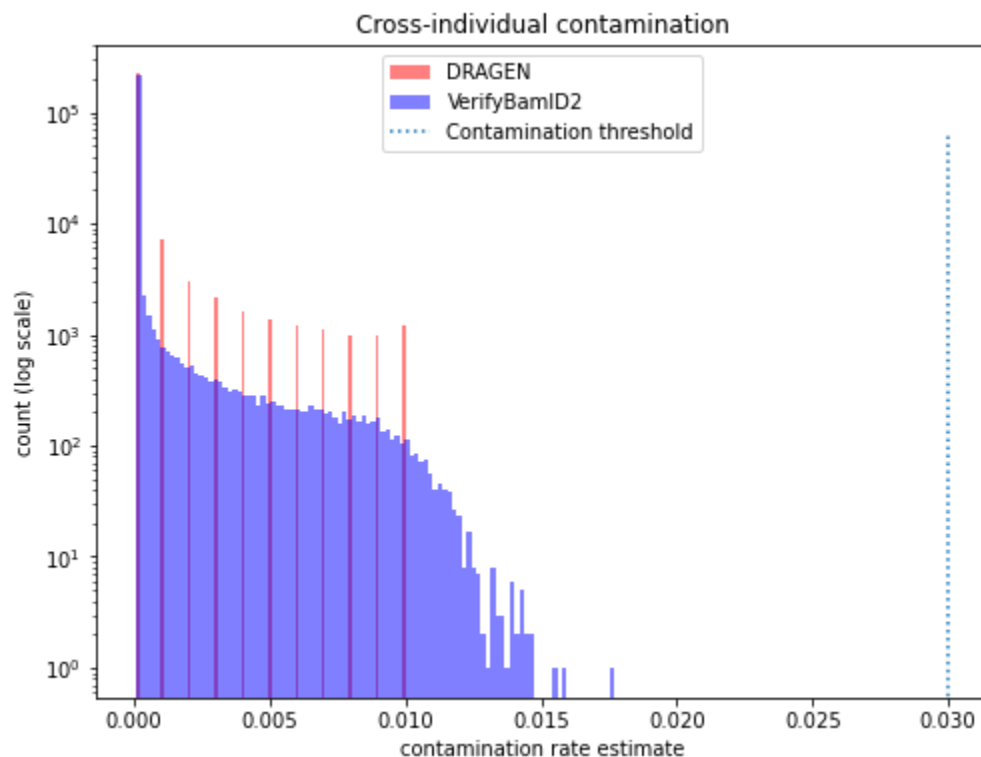
BamFile	WGS cram file
Reference	"gs://gcp-public-data--broad-references/hg38/v0/Homo_sapiens_assembly38.fasta"
UDPath	"gs://gcp-public-data--broad-references/hg38/v0/contamination-resources/1000g/1000g.phase3.100k.b38.vcf.gz.dat.UD"
BedPath	"gs://gcp-public-data--broad-references/hg38/v0/contamination-resources/1000g/1000g.phase3.100k.b38.vcf.gz.dat.bed"
MeanPath	"gs://gcp-public-data--broad-references/hg38/v0/contamination-resources/1000g/1000g.phase3.100k.b38.vcf.gz.dat.mu"
Verbose	specified

Please see [Appendix H](#) for the DRAGEN command line parameters, as the command line contains multiple functions, including calculating contamination.

## Results

The hard threshold for contamination was 0.03 for the research pipeline, higher than 0.01 for the clinical pipeline [\[5\]](#).

We did not include any samples with a contamination larger than 0.018 and only three samples greater than 0.015. [Figure 5](#) demonstrates the frequency of the contamination estimates for samples in the v7 release.



**Figure 5** -- srWGS contamination estimates from both sources (DRAGEN and VerifyBamID2). DRAGEN rounds the contamination estimate to three decimal places. Note the log scale of the counts (y-axis). Over 89.0% and 91.4% of srWGS samples had contamination estimates lower than  $1e-4$  by VerifyBamID2 and DRAGEN, respectively.

## Coverage


### Method

Coverage is defined as the number of reads covering the bases of the genome. Maintaining coverage is important for consistent statistical power and accurate variant calling. We apply several thresholds (summarized from the FDA IDE (G200165)):

- Mean coverage (threshold  $\geq 30x$ ) - This is the mean number of overlapping reads at every targeted base of the genome. Accuracy steadily decreases as mean coverage decreases, with a rapid decrease below 20x coverage, supporting a stringent threshold selection of a minimum of 30x.
- Genome coverage (threshold  $\geq 90\%$  at 20x) - Accuracy steadily decreases as the percent of bases with at least 20x coverage drops. Drop-off of performance is initially gradual, supporting a threshold of 90%.
- [All of Us Hereditary Disease Risk gene \(AoUHDR\)](#) coverage (threshold  $\geq 95\%$  at 20x) - For clinically relevant areas of the genome, we insist on higher mean coverage to ensure a higher calling accuracy. As we reduce the coverage in the AoUHDR region, the reduction in performance is slow initially but increases rapidly below 40%, showing that the threshold of 95% is conservative.
- Aligned Q30 bases (threshold  $\geq 8e10$ ) - All bases in the sequencing reads get a quality assignment, which is phred scaled (Q30  $\rightarrow$  probability of error is 0.001) [9]. As lower base quality counts increase, we see a reduction in accuracy with an inflection point starting around  $6e10$ .

## Results

As seen in [Figure 6](#), all srWGS samples exceed the thresholds that we set as part of the research pipeline. We had 281 (0.1%) samples with mean coverage greater than 70x.



**Figure 6** -- Coverage metrics for the v7 release srWGS samples. The orange line is the threshold for each metric. There are 281 samples (0.1%), with mean coverage greater than 70x, that are not included in the mean coverage (upper left) nor aligned q30 bases (lower right) plots. As expected, these samples were outliers in the number of aligned q30 bases (i.e., higher base count than samples with lower mean coverage).

## Short-read WGS SNP & Indel Joint Callset QC

The srWGS small variants are delivered as a joint callset and the QC steps in this section are performed on the joint callset, not individual samples [\[10\]](#). Please note that the QC steps described here apply during creation of the srWGS joint callset, after single sample QC. The joint callset is available in the RW and other auxiliary information (including lists of flagged samples) is available through the User Support Hub [\[1\]](#). The joint callset QC process is similar to that of gnomAD 3.1 [\[11\]](#), though not exactly the same. See a summary of the joint callset QC steps in [Table 3](#). Sample QC is performed before Variant QC.

We flag samples or variants as failing QC, rather than removing them from the callset, since we could not validate whether samples (especially population outliers) were problematic or were

just a part of a poorly-sampled ancestry. Flagged variants can also be a result of poorly-sampled ancestry.

**Table 3 -- srWGS SNP & Indel joint callset QC summary**

QC process	Sample or variant QC	Error modes addressed	v7 release results
Sample Hard Threshold Flag	sample	Extremely noisy samples	No samples flagged.
Sample Population Outlier Flag	sample	Noisy samples	551 samples flagged (0.2%). Based on regressing out the PCAs from callset metrics, such as snp_count.
Variant Hard Threshold Filters	variant	Artifacts that cannot be detected in a single sample	This has a simple implementation with high precision, which saves compute for downstream variant filtering. 59,496,403 were filtered 972,115,272 were not filtered
Allele-Specific VariantQualityScore Recalibration (AS-VQSR)	variant	Artifacts that cannot be detected in a single sample	See <a href="#">[12]</a> .
Sensitivity and Precision Evaluation	both	Poor variant detection	See <a href="#">Appendix I</a> for a list of samples.
<b>Auxiliary processes</b>			
Ancestry	sample	Flagging sample outliers and allows calculation of population level metrics, such as allele frequency (AF).	Error rate from holdout set (incl. Other): 0.046 Error rate from holdout set (not incl. Other): 0.002 Concordance vs self-reported: 0.915 See <a href="#">Appendix A</a> .  Number of independent, bi-allelic sites ("high-quality sites") used: 151159 See <a href="#">Appendix J</a> .
Relatedness and maximal independent set of samples	sample	Related samples, which confound analyses	19374 related pairs and 15376 samples in the maximal independent set.  See <a href="#">Appendix K</a> . This process produces a list of the sample pairs with kinship score, calculated by Hail <a href="#">[13]</a> . No samples are removed from the callset, but this allows researchers to easily remove a minimal set of samples to eliminate related samples in the callset.

## Sample Hard Threshold Flag

We flag srWGS individual samples based on these sample-level QC metrics. The flagged samples can be found in a list on the User Support Hub [\[1\]](#).

## Method

We initially flagged any samples with strong erroneous signals. We calculated all metrics using autosomal territory only. The criteria for being eliminated as “obviously erroneous”:

- number of SNPs: < 2.4M and > 5.0M
- number of variants not present in gnomAD 3.1: > 100k
- heterozygous to homozygous ratio (SNPs and Indel separately): > 3.3

## Results

We did not flag any samples for failing hard thresholds.

## Sample Population Outlier Flag

We flag srWGS individual samples based on the population outlier data. The flagged samples can be found in a list on the User Support Hub [\[1\]](#).

## Method

As part of ancestry prediction (see [Appendix A](#)), we regressed out sixteen principal component features computed and used the residuals to determine the outliers. We define outlier samples as being eight median absolute deviations (MADs) away from the median residual in any of the following metrics:

- i. number of deletions
- ii. number of insertions
- iii. number of SNPs
- iv. number of variants not present in gnomAD 3.1
- v. insertion : deletion ratio
- vi. transition : transversion (TiTv) ratio
- vii. heterozygous to homozygous ratio (SNPs and Indel separately)

## Results

We flagged 551 (0.2%) samples as outliers based on at least one of the above criteria (See [Table 4](#)). Plots of the first principal components against these eight metrics can be found in [Appendix L](#).

**Table 4 -- srWGS SNP & Indel population outlier sample counts**

Metric(s) considered	Flagged sample count
Indel heterozygous to homozygous ratio	307
Deletion count + Indel heterozygous to homozygous ratio + Insertion count + SNP count	74
Indel heterozygous to homozygous ratio + SNP	48

heterozygous to homozygous ratio	
Indel heterozygous to homozygous ratio + SNP count	41
Variants not present in gnomAD 3.1 count	28
Deletion count + Indel heterozygous to homozygous ratio + SNP count	26
SNP heterozygous to homozygous ratio	17
Ti/Tv ratio + Variants not present in gnomAD 3.1 count	5
Indel heterozygous to homozygous ratio + Variants not present in gnomAD 3.1 count	3
Ins/del ratio	1
Indel heterozygous to homozygous ratio + SNP count + SNP heterozygous to homozygous ratio	1

Total 

551
-----

## Variant Hard Threshold Filters

These site-level QC metrics for the srWGS SNP & Indel callset will flag variants, appearing as filtered in the site level filters of the VDS and VCF (`filters` in the VDS, `FILTER` in the VCF). These variants will still be included in cohorts, including in the Cohort builder.

### Method

If a variant does not meet the following criteria, it will be filtered:

- No high-quality genotype ( $GQ \geq 20$ ,  $DP \geq 10$ , and  $AB \geq 0.2$  for heterozygotes) called for the variant.
  - Allele Balance (AB) is calculated for each heterozygous variant as the number of bases supporting the least-represented allele over the total number of base observations. In other words,  $\min(AD) / DP$  for diploid GTs.
  - Filter field value: `NO_HQ_GENOTYPES`
- `ExcessHet` < 54.69
  - `ExcessHet` is a phred-scaled p-value. We cutoff of anything more extreme than a z-score of -4.5 (p-value of  $3.4e-06$ ), which phred-scaled is 54.69
  - Filter field value: `ExcessHet`
- `QUAL` score is too low (lower than 60 for SNPs; lower than 69 for Indels)
  - `QUAL` tells you how confident we are that there is some kind of variation at a given site. The variation may be present in one or more samples.

- Filter field value: LowQual

## Results

Unfiltered variants will have “.” or PASS in the site level filters fields in the srWGS joint callset SNP & Indel VCFs, VDS, and Hail MTs. Filtered variants will have the filter name in the site level filters of the VCF, VDS, or Hail MT (FILTER or filters). We recommend that researchers do not include variant sites that were filtered in their analyses. The variant counts can be found in [Table 5](#).

**Table 5 -- srWGS SNP & Indel variant hard threshold filter counts**

Filters	Numbers
None	972115272
'NO_HQ_GENOTYPES'	33526160
'NO_HQ_GENOTYPES', 'LowQual'	22245268
'LowQual'	3064830
'ExcessHet'	659051
'NO_HQ_GENOTYPES', 'ExcessHet'	1094

## Allele-Specific Variant Quality Score Recalibration (AS-VQSR)

These genotype-level QC metrics for the srWGS SNP & Indel callset will flag variants. The AS-VQSR tool scores genotypes, at some sites only some genotypes are filtered whereas at other sites all genotypes are filtered. We do not report the AS-VQSR scores in the srWGS SNP & Indel callset, we only report whether or not a genotype or variant is filtered.

A filtered genotype will appear as filtered in the genotype level filter (FT) in the VCF, VDS, and Hail MT. In the VDS, FT will contain True for PASS and False for FAIL. In the VCF or Hail MT, FT will contain PASS or FAIL. If all genotypes fail the AS-VQSR filtering at a variant site, the site will be filtered in the VDS filter field (filters) or the VCF/Hail MT filter field (FILTER). All variants will still be included in cohorts, including in the Cohort builder.

## Method

As part of the joint calling, we will filter variants with Allele-Specific Variant Quality Score Recalibration (AS-VQSR or VQSR) [\[12\]](#). This filtering technique uses machine learning to identify variants across samples that are likely artifacts. We used the following annotations as features for training:

- Variant Confidence/Quality by Depth (AS\_QD)
- Z-score From Wilcoxon rank sum test of Alt vs. Ref read mapping qualities (AS\_MQRankSum)

- Z-score from Wilcoxon rank sum test of Alt vs. Ref read position bias (AS\_ReadPosRankSum)
- Phred-scaled p-value using Fisher's exact test to detect strand bias (AS\_FS)he
- RMS Mapping Quality of reference vs alt reads (AS\_MQ) [SNPs only]
- Symmetric Odds Ratio of 2x2 contingency table to detect strand bias (AS\_SOR)

We used the default training sets as described in the GATK documentation [14] and Table 6. Training sets are flagged as true or training sites and assigned an initial prior likelihood score. Details of these parameters can be found in the GATK documentation [12,14], and the sites can be found as public resource downloads for the GATK [15].

**Table 6 – srWGS SNP and Indel AS-VQSR training and truth datasets**

Training Set Name	SNP or Indel	Truth	Training	Prior Likelihood	Description
Omni [16]	SNP	True	True	Q12 (93.69%)	This resource is a set of polymorphic SNP sites produced by the Omni genotyping array.
HapMap [17]	SNP	True	True	Q15 (96.84%)	This resource is a SNP callset that has been validated to a very high degree of confidence.
1000 Genomes [18]	SNP	False	True	Q10 (90%)	This resource is a set of high-confidence SNP sites produced by the 1000 Genomes Project.
Mills [19]	Indel	True	True	Q12 (93.69%)	This resource is an Indel callset that has been validated to a high degree of confidence.
Axiom [18]	Indel	False	True	Q10 (90%)	This resource is an Indel callset based on the Affymetrix Axiom array on 1000 Genomes Project samples.

## Sensitivity and Precision Evaluation

### Method

In the callset, we included four well-characterized control samples (four Genomes-in-a-Bottle samples (GiaB) [20] from HapMap [17] and Personal Genome Project; see Appendix I), which we can use to determine sensitivity and precision. The samples were sequenced with the same protocol as the *All of Us* samples. These samples do not appear in any user data (e.g., cohorts built using the RW).

We use the high confidence calling region, defined by GiaB v4.2.1, as the source of ground truth. In order to be called a true positive, a variant must match the chromosome, position,



reference allele, and alternate allele. In cases of sites with multiple alternate alleles, each alternate allele is considered separately.

## Results

Sensitivity and precision results can be seen in [Table 7](#).

**Table 7 -- Sensitivity and precision measurements for control samples using the *All of Us* sequencing protocol**

Variant type	Sample	Sensitivity	Precision
SNV	HG-001	0.995	>0.999
	HG-003	0.988	>0.999
	HG-004	0.988	>0.999
	HG-005	0.989	>0.999
Indel	HG-001	0.987	0.996
	HG-003	0.985	0.997
	HG-004	0.986	0.998
	HG-005	0.994	0.999

## Short Read Structural Variants (srWGS SVs)

Please note: As of June 17, 2024, we have released an additional set of v7 srWGS SVs, increasing the srWGS samples that have SV data to 40%. The QC steps for these 97,940 srWGS samples with SV calls is available as a [separate article](#).

The QC information for the original 11,390 srWGS SV samples is available in the [archived section](#) of the User Support Hub [\[1\]](#).

# Long Reads

There are 1,027 long read sequencing (lrWGS) samples in the v7 genomic data release. We applied a consistent sequencing protocol and performed QC checks for all samples. We also did quality checks on the joint SNP and Indel callset using information across samples. SNP and Indel variants and SVs are available for long read sequences in VCF format. In addition, long read sequences aligned to the grch38\_noalt and T2Tv2.0 references are available in BAM format [\[22\]](#). De novo assemblies are available in both FASTA format and graphical fragment assemblies (GFA) format. The data is described in the ‘How the *All of Us* Genomic data are organized’ article on the User Support Hub [\[1\]](#).

## Data generation for lrWGS

Only one sequencing center generated the long read whole genome sequencing (lrWGS) data with a single technology, so we did not perform any cross-sequencing center consistency checks.

Samples were selected for lrWGS sequencing if they had srWGS and Array data in this v7 release, are self-reported African American participants, and are unrelated. Please see [Appendix B](#) for the self-reported ancestry report for participants with long reads data.

HudsonAlpha was the single sequencing facility (“sequencing center”) commissioned for performing long reads sequencing for this cohort. The sequencing center used the Single Molecule Real Time (SMRT) sequencing technology [\[23\]](#) from Pacific Biosciences (PacBio). A SMRT cell is a chip conceptually similar to a flowcell in short read sequencing. Each SMRT cell is made of millions of zero-mode waveguides, where the molecules are trapped and sequenced. Each molecule is scanned multiple times. The scanned data goes through consensus-processing (Circular Consensus Sequencing or CCS), which combines multiple lower-accuracy readings of a single molecule to produce a high accuracy consensus sequence. The sequencing instrument outputs the CCS-processed reads for each SMRT cell as an unaligned BAM (uBAM). Please refer to the PacBio glossary [\[24\]](#) for a more thorough definition of terms.

The uBAM files containing data for one SMRT cell are then delivered to the DRC along with metadata (e.g. which samples are sequenced with this SMRT cell, barcodes used if applicable, etc). The DRC extracts the CCS reads with a quality score greater than Q20 (HiFi reads) from the uBAM.

## Single Sample QC for lrWGS

Samples can be sequenced across multiple SMRT cells, which are then demultiplexed and aggregated into BAMfiles for each sample. For each sample, we run QC on each individual demultiplexed SMRT cell and the aggregated sample bam file ([Table 8](#)). Three of the QC processes are performed on both the demultiplexed SMRT cell data and the aggregated sample data. As expected, no aggregated samples failed these three QC checks.

The HiFi reads of the demultiplexed SMRT cells are aligned to two references: “grch38\_noalt” and “T2Tv2.0”. grch38\_noalt corresponds to the GRCh38 reference with no alternate sequences [25,26]. T2Tv2.0 corresponds to the T2T-CHM13v2.0 reference, with the EBV contig added from the grch28\_noalt reference [27]. We check the grch38\_noalt BAM files before downstream processing, described in Table 8. In total, the DRC performed QC on 2,597 demultiplexed SMRT cell samples for a total of 1,027 aggregated samples.

Running three QC steps at both the demultiplexed SMRT cell level and aggregated sample level allows us to identify errors early in the process and to find quality issues at both levels. Please note that we do not use cross-sample information in the Single Sample QC process. Appendix M provides an overview of the processing steps for the IrWGS data, including processes after QC, such as variant calling and assemblies.

**Table 8 -- QC processes performed on single sample data in each SMRT cell and after aggregation**

QC Process	Aggregated sample or demultiplexed SMRT cell?	Passing criteria	Error modes addressed	V7 release results
Fingerprint concordance	Both	Log-likelihood ratio > 6	- Sample swaps - Large amounts of cross-individual contamination	All IrWGS samples are concordant with array samples.
Sex concordance	Both	Sex call is concordant with self-reported sex at birth. OR Self-reported sex at birth reported as “Other” or was not reported	- Sample swaps	All IrWGS samples are concordant.  *Other refers to a participant self-reporting “Intersex”, “I prefer not to answer”, or “none of these fully describe me”
Cross-individual contamination	Both	< 0.03 (<3%)	- Sample contamination from another individual	All IrWGS samples meet the threshold (Aggregated samples plotted in Figure 8).
Coverage	Aggregated sample	≥ 5x mean coverage	- Sample preparation errors - Poor sensitivity and precision of variant calling	All IrWGS samples passed this check.
Read length median	Aggregated sample	≥ 10,000 bp	- Shorter fragments significantly impacting variant calling performance	All IrWGS samples passed this check.
Outlier sample	Aggregated sample	Manual threshold	-Poor variant calling	Five total samples

filtering			performance	were outliers due to variant counts and are not included in the v7 lrWGS data release (See <a href="#">Figure 11</a> ).
-----------	--	--	-------------	---

## Fingerprint Concordance

### Method

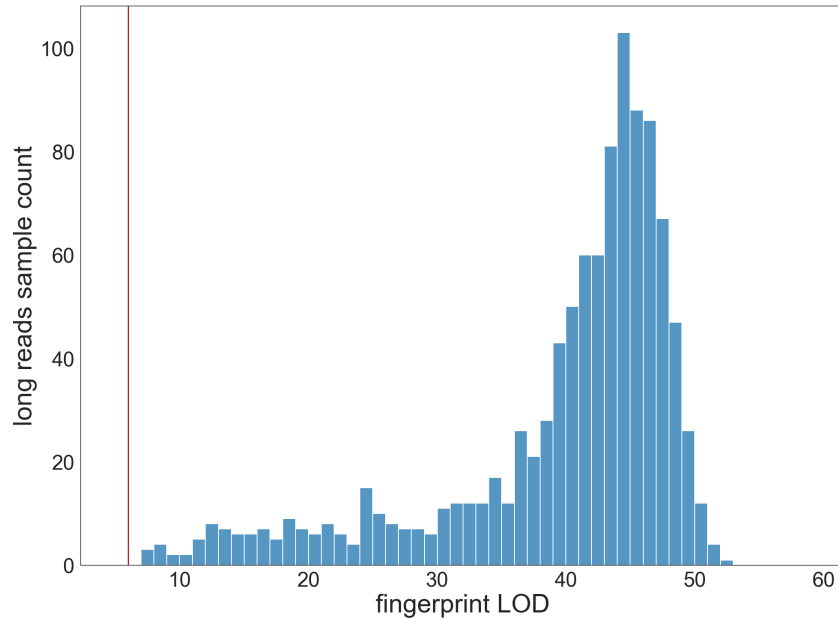
Each grch38\_noalt BAM is checked against a fingerprint VCF to verify their marked identity from the sequencing metadata. This is applied to both the SMRT cell demultiplexed reads for the sample and the aggregated sample reads. We use the same fingerprint VCFs that are used by the srWGS fingerprint verification pipeline. Please refer to the [Fingerprint Concordance](#) method of the srWGS SNP & Indel QC process, as we follow the same method for lrWGS data, with some engineering adaptations that do not change the algorithm. The HAPLOTYPE\_MAP file for lrWGS fingerprint concordance can be found at

“gs://gcp-public-data--broad-references/hg38\_noalt/v0/aou/fp/lr.aou.fp.haplotype\_database.no\_alt.txt”, which differs from the srWGS HAPLOTYPE\_MAP file only in the header section.

If a BAM failed fingerprint concordance, we launched a pipeline to find its identity by exhaustively testing to find a match against the other lrWGS fingerprint VCFs, whether planned or already sequenced. The true identity is the identity of the fingerprint VCF that returns the highest LOD that is above 6.0. We did not encounter any cases where a BAM file matched more than one other sample. If no matching VCF was found, then the lrWGS data was not released. The DRC also returns findings about swaps and their computationally-resolved identities back to the sequencing center for cross verification.

### Results

All lrWGS samples (including all corresponding demultiplexed SMRT cells) in the v7 release passed the fingerprint concordance check based on the corresponding array VCF ([Figure 7](#)).



**Figure 7** -- Fingerprint LOD histogram of the v7 IrWGS samples. All IrWGS v7 samples exceeded the fingerprint LOD threshold of 6.0.

## Sex Concordance

### Method

A simple sex concordance check is also performed on the grch38\_noalt version of each BAM. The tool `mosdepth` (v0.3.3) [28] is run to calculate coverage across the whole genome and over each chromosome. A formula is then used for inferring the sex ploidy of the sample

$$\text{Ploidy}_x = \text{round}( 2 * \text{cov}(\text{chrX}) / \text{cov}(\text{chr1}) )$$

$$\text{Ploidy}_y = \text{round}( 2 * \text{cov}(\text{chrY}) / \text{cov}(\text{chr1}) )$$

The self-reported sex assigned at birth from the CDR data for each sample is then checked for contradictions with the inferred sex chromosome ploidies. If we do not have a “male” or “female” for the sex assigned at birth for the sample from the CDR data, because the participant reported it as “Intersex”, “I prefer not to answer”, “none of these fully describe me”, or skipped the question, we passed the sex concordance check, regardless of the information from the inferred sex ploidy. The sex assigned at birth data from the CDR is described in [Appendix F](#).

### Results

We do not include any IrWGS samples that fail the sex concordance check in the v7 release samples, but one sample was included after a manual inspection. This sample, whose self-reported sex assigned at birth was male, received normalized X and Y coverage of 0.97 and 0.22. This happened in two separate SMRT cells for the sample. As a comparison, the mean normalized chrY coverage of female samples who passed the IrWGS sex concordance

check (samples whose self-reported sex at birth is female and the lrWGS calculated sex is female) is 0.05 with a 0.01 standard deviation. With rounding, the sample's chromosome X and Y ploidy was reported as 1 and 0, respectively; but because of the normalized 0.22 coverage on the Y chromosome, we manually marked the sample as passing and included it in the release.

## Cross-individual Contamination Rate

### Method

VerifyBamID2 (version 2.0.1) was adapted into a pipeline for estimating cross-individual contamination for lrWGS data. VerifyBamID2 was originally designed for short read sequencing. To make the process scalable, we converted the grch38\_noalt BAM to a pileup format at selected sites, where VerifyBamID2 genotypes the input, in parallel per-chromosome.

We also evaluated the accuracy of VerifyBamID2 around the 3% contamination cutoff for lrWGS data to determine if the tool would erroneously pass or fail samples. We did this through an *in silico* mixture of samples, simulating different contamination scenarios and at different levels:

- Cross contamination from a sample from a different population and of opposite sex.
- Cross contamination from a sample from a different population and of the same sex.
- Cross contamination from a sample from the same population of different sex.
- Cross contamination from within a family, i.e. parent-child contamination.

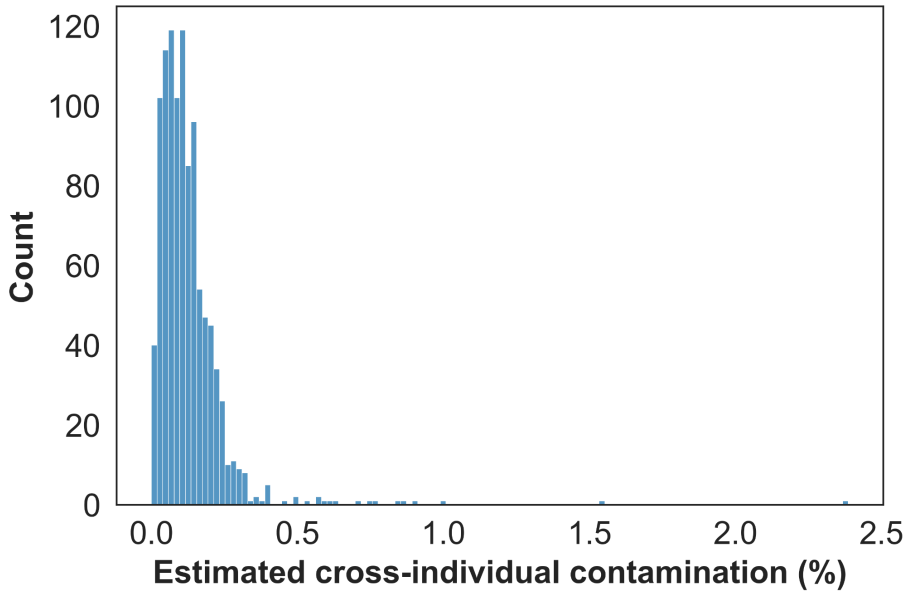
We did not have publicly accessible long reads data for assessing the case where the contaminant is from the same population and the same sex. Given that the sites used by VerifyBamID2 for estimation are all autosomal sites, we don't believe this case will have any effect. All *in silico* mixed BAMs have coverage ~8X to emulate the production coverage.

We tested six levels of contamination (3%, 9%, 17%, 33%, and 50%). At 3%, 9%, and 17%, the error between VerifyBamID2 and our *in silico* mixture was never over 10% of the testing contamination level (eg, error was < 0.3% when testing an *in silico* mixture of 3%). At higher tested contamination levels (33% and 50%), the error stayed within 20%. Note that if contamination were to be this high, fingerprint verification would have failed the sample.

We observed from this experiment that for unrelated samples, VerifyBamID2's estimations are in line with the expected contamination level. For related samples, VerifyBamID2 tends to significantly underestimate the contamination level. This does not affect the v7 lrWGS release because all samples are unrelated.

### Results

We did not include any lrWGS samples with a cross-individual contamination rate higher than 2.5% ([Figure 8](#)). This is below the error (2.7%) we observed in the *in silico* testing of VerifyBamID2 at 3% contamination, which indicates that all released samples have a cross-contamination rate below 3%.



**Figure 8** -- Estimated cross-individual contamination for the IrWGS samples, calculated by VerifyBamID2. All samples met the 3% threshold, even when accounting for possible errors in VerifyBamID2.

## Coverage


### Method

Coverage is defined as the number of reads covering the bases of the genome. Maintaining coverage is important for consistent statistical power and accurate variant calling. We applied a coverage threshold of  $\geq 5x$  mean coverage for the IrWGS samples in order to remove any samples that did not have enough coverage for statistical power and accurate variant calling.

### Results

We did not release any samples that did not meet the mean coverage threshold. The frequency of mean coverage can be seen in [Figure 9](#).





**Figure 9** -- Coverage histogram of the lrWGS samples in the v7 release. Note that a mean coverage of 5x was the minimum to be included in the release.

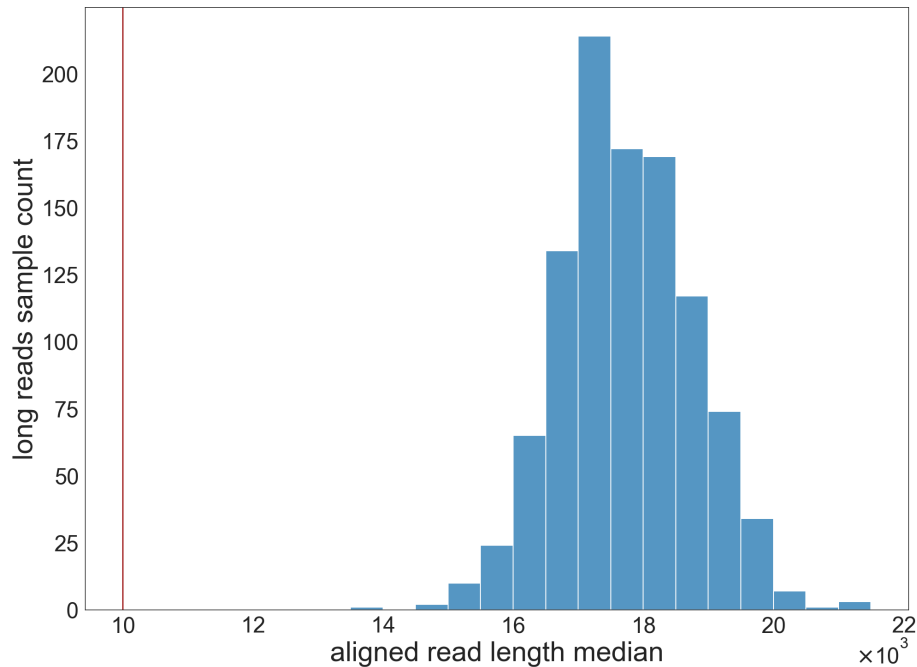
## Read Length Median

### Method

We calculated the read length median to determine if any samples had shorter fragments that would significantly impact the variant calling performance. The threshold read length median was  $\geq 10,000$  base pairs and all lrWGS samples passed this check.

### Results

We did not release any samples that did not meet the read length median threshold. A distribution of the read length median can be seen in [Figure 10](#).



**Figure 10** -- Read length median histogram of the IrWGS samples in the v7 release. Note that a read length median of 10,000 base pairs was the minimum to be included in the release.

## Outlier Sample Filtering

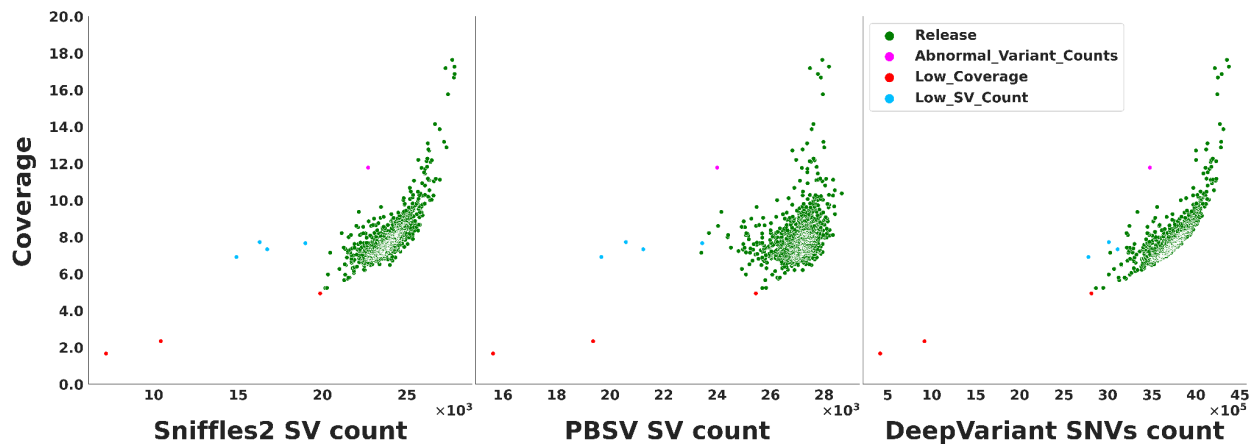
### Method

Variant calling was performed on single sample BAM files for long reads data. SNP and Indel single sample callsets were generated using the PEPPER-Margin-DeepVariant pipeline [29] (version 0.4.1 for PEPPER and version 1.3.0 for DeepVariant). We filtered out events with a QUAL score less than 40. We called SVs with Sniffles2 [30] and PBSV [21].

We used two independent criteria when detecting outlier samples using their single sample VCFs: abnormal SV counts and SNP and Indel counts, given their coverage. We found outliers by plotting the variant counts versus the coverage and manually evaluating the distribution across the entire cohort, seen in [Figure 11](#).

### Results

We removed four samples from the v7 release due to their low SV counts and we removed one sample from the release due to the combined SV and small variant count being an outlier. We do not include these samples in the v7 IrWGS release.



**Figure 11** -- We compared the variant count from long read samples to the sample coverage and manually identified outliers. The Sniffles2 and PBSV plots are SV counts and the DeepVariant plot represents SNP and Indel counts. Each dot represents one sample. Green dots (1027) are included in the v7 Release. Cyan dots (4) are excluded because they appear to be outliers when judged by their SV counts. Magenta dots (1) are excluded because they appear to be outliers when judged by their SV and SNP and Indel counts. Red dots (3) are excluded due to low coverage.

## Joint Callset QC for the IrWGS SNP/Indel callsets

Joint callset QC is performed on the joint SNP and Indel callset from the long reads data ([Table 9](#)). The single sample VCFs were joined using GLNexus (version 1.4.1) [\[31\]](#), parallelized per chromosome. Final joint-called SNP and Indel callsets were converted to Hail MatrixTable (Hail MT) format for analysis. For the Hail MT joint callset QC relatedness and sample population outlier analyses, we used high quality sites that can be called accurately ([Appendix J](#)). If a sample was an outlier, we did not include it in the v7 release, as opposed to the srWGS joint callset QC, which flagged but did not remove outlier samples.

**Table 9** -- IrWGS joint callset QC steps

QC process	Error modes addressed	V7 release results
Relatedness	- Sample swaps - Related samples, which confound analyses	No sample pairs had kinship scores > 0.1, indicating that no samples were related.
Sample Population Outlier	- Noisy samples	All outlier samples were removed from the v7 release and are not included in the IrWGS joint callset

## Relatedness

### Method

We used the Hail `pc_relate` function to determine the kinship score of pairs, as we did with the short reads data in [Appendix K](#). We ran `pc_relate` on the Hail MT based on the IrWGS SNP/Indel joint callset.

### Results

We found that no IrWGS sample pairs had kinship scores  $> 0.1$ . This indicated that no IrWGS samples were related to a 2nd degree or closer, as we expected.

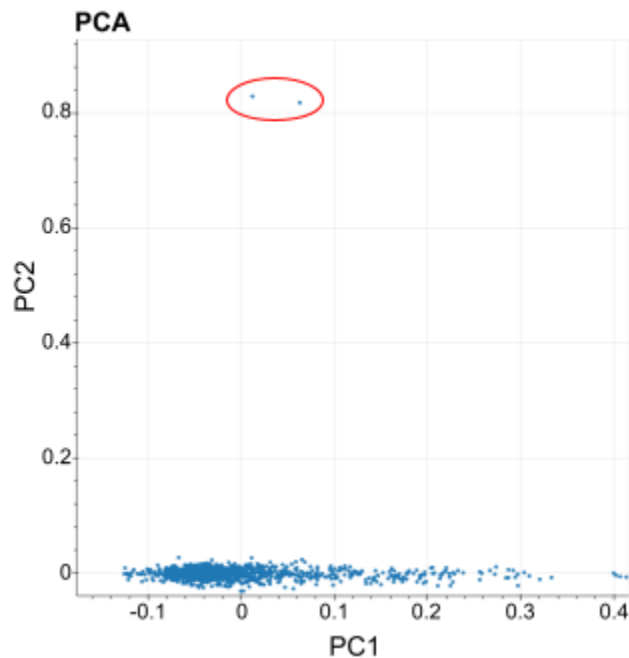
## Sample Population Outlier

### Method

We used the high quality sites from the IrWGS SNP/Indel joint callset ([Appendix J](#)) to do a population analysis and find population outliers. Following a similar method to the [srWGS SNP & Indel sample population outlier flag](#) analysis, we regressed out sixteen principal component features for the joint IrWGS genotype matrix in Hail.

### Results

We manually inspected samples in PC1 and PC2 and discovered two outlier samples in the callset, which were excluded from the release ([Figure 12](#)).



**Figure 12** -- Principal component analysis of the population structure of the IrWGS joint callset in order to determine outlier samples. The two outliers were identified manually and removed from the v7 IrWGS release.

# Known Issues

The issues below apply to the v7 release genomic data (arrays, srWGS, srWGS SVs, long reads, and auxiliary data). We have provided suggested actions for researchers to workaround the issues and provided remediation plans when necessary. Sample lists relevant to these issues can be found in the User Support Hub [\[1\]](#).

## Known Issue #1: Small subset of samples missing corresponding CDR data

**Update as of April 18, 2024.** *The 20 array samples have been removed from the array Hail MT and PLINK files, released as version 7.1 on April 18, 2024.*

Six srWGS samples and 20 Array samples in this release are missing their corresponding CDR data. The affected participants are consented to appear in the genomic data.

Affects:

- srWGS SNP & Indel samples: VDS, VCF, PLINK, and Hail MT formats
- Array samples: VCFs, PLINK, and Hail MT

Suggested action:

- If you are not using CDR data (e.g., surveys, EHR), then no action.
- Otherwise, remove samples without corresponding CDR data. We will provide the lists of srWGS and array samples without corresponding data in the CDR.

Remediation:

- We will provide two lists (srWGS and array) of the affected samples through the CDR.

## Known Issue #2: 11 samples were affected by a sample swap incident

**Update as of April 18, 2024.** *The 11 array samples have been removed from the array Hail MT and PLINK files, released as version 7.1 on April 18, 2024.*

We identified a sample swap incident from the GCs which affects 11 samples with genomic data, one srWGS sample and 11 array samples in the v7 release. This was due to an internal issue that we have now remediated.

Affects:

- srWGS SNP & Indel samples: VDS, VCF, PLINK, and Hail MT formats
- Array samples: VCFs, PLINK, and Hail MT

Suggested action:

- Remove samples that are affected by the sample swap from your analysis. We provide two list files containing the research IDs of the affected samples, one file for srWGS and one file for array.

Remediation:

- We have fixed these sample swaps for future releases.

## Known issue #3: Array samples (N=416) from previous release are missing in this release

A total of 416 array samples from the previous release are not included in the newest v7 release. This happened for multiple reasons:

1. For this release (v7 C2022Q4R9), we have reprocessed the arrays (details of the new approach are found in [Appendix E](#)). As a result, some array samples passed QC in the previous release, but did not pass QC checks in this release.
2. The participant withdrew consent between releases.
3. The sample was not reprocessed in time for this release. These samples will appear in a future release.

At this time, we cannot provide a breakdown of the counts for each of the above reasons.

Affects:

- Array samples: VCFs, PLINK, and Hail MT.

Suggested action:

- If you are not using any array data from the previous release (C2022Q2R2), then no action is necessary.
- If your cohort includes one of the samples that is missing you should upgrade your cohort to include only samples in the v7 set.

Remediation:

- We have released a list of missing array samples on the User Support Hub [\[1\]](#).
- We will continue to re-process the array samples (ETA 2023). All new array samples will use the new process as described in [Appendix E](#).

## Known Issue #4: Single array sample missing from Array Hail MT and PLINK files

**Update as of April 18, 2024.** *The sample has been included in the array Hail MT and PLINK files, released as version 7.1 on April 18, 2024.*

For sample 3518297, we only have a corresponding array in VCF. This sample was not included in the array Hail MT and PLINK files. This was due to an internal issue in synchronization between our srWGS and array sample lists.

Affects:

- Array Hail MT and PLINK files

Suggested action:

- No action. For array analyses, we recommend that you proceed with this one sample missing. We believe that the effort to add one sample (0.0003%) to the MT or PLINK files is not worth the expense.

Remediation:

- This will be fixed in the next release.

## Known Issue #5: Larger than expected changes in ancestry predictions from previous release

We have found that the predicted ancestry for 7744 participants changed from the previous release (7.9 % of the v6 C2022Q2R2 data release). Most of the changes were to and from the “Other” classification (7519 participants, 7.6% of the v6 release). This was due to a change in the SNP set used as the features in the PCA for prediction. We used chr20 and chr21 in v6 and the autosomal exome in the current v7 release (see [Appendix A](#)). The ancestry changes may affect your analysis if you migrate from the previous release to this release. In future releases, we do not expect large ancestry changes (see Remediation below). The VAT uses these computed ancestries to generate *All of Us* population (gvs\_\*\_\*) annotations. The genomic data in the public Data Browser are also dependent on the ancestry predictions for populating population information about variants.

### Affects:

- Ancestry predictions, if you are migrating from v6 ancestry to v7.
- Variant Annotation Table (VAT)
- Public Data Browser

### Suggested action:

- If you are not using ancestry predictions from the previous v6 release, then no action
- If you are migrating your analysis from the previous v6 release to the v7 release and you used ancestry predictions, we recommend that you rerun downstream analyses that are affected by computed ancestry.

### Remediation:

- We are planning to continue to use the current (v7) SNP set as features for ancestry prediction and thus, we do not expect large ancestry changes in future releases. In an internal test, we used the new feature set with the previously released v6 data and compared the ancestry predictions to the current v7 data with the new feature set. We found that only 1890 samples (1.9%) changed ancestry assignment. Note that of these samples, all samples switched to/from the “Other” classification.

## Known issue #6: Ancestry prediction has higher error rates for Middle Eastern ancestry

A paucity of labeled Middle Eastern samples reduced the performance of the random forest classifier. This caused the confidence to dip when predicting ancestry for Middle Eastern samples, which caused a larger proportion of these samples relative to other computed ancestries, to be classified as Other (“oth”). The VAT uses these computed ancestries to generate *All of Us* population (gvs\_mid\_\* and gvs\_oth\_\*) annotations. The genomic data in the public Data Browser are also dependent on the ancestry predictions for populating population information about variants.

See [Table A.2](#) for details of the ancestry prediction performance

### Affects:



- Ancestry predictions
- Variant Annotation Table (VAT)
- Public Data Browser

Suggested Action:

- When limiting cohorts to samples with computed ancestry of Middle Eastern (“mid”), use the ancestry predictions that do not include “other”. In other words, use the “ancestry\_pred” column, instead of “ancestry\_pred\_other”.

Remediation:

- We have completed investigating other approaches and have minimized this error in the data. However, we do not have enough samples of middle eastern ancestry in our training data to fully remediate this issue. We will be adding this information in the FAQ section for all future releases.

## Known Issue #7: VDS issue with GT

In the srWGS SNP & Indel VDS, we have precalculated the genotypes (GT) from the local genotype field (LGT). If you filter any participants from the VDS, you will need to recalculate the GT field. Examples of how to do this are found in example notebooks in the Researcher Workbench.

Affects:

- srWGS SNP & Indel VDS

Suggested Action:

- If you are not using the VDS, no action.
- If you filter participants from the VDS, make sure to recalculate the GT fields.

Remediation:

- We have provided documentation in Researcher Workbench showing how to recalculate GT fields from the VDS.
- In future releases, we will also provide a FAQ question and eliminate this Known Issue.

## Known Issue #8: AS\_VQSLOD is incorrect in the VDS and dropped in callset data with all participants

There is inconsistency in the AS\_VQSLOD annotation in the srWGS SNP & Indel VDS. This annotation indicates information regarding the allele-specific filtering model and we generally do not recommend that researchers use it for their analysis. Instead, we recommend using the filters annotations. In the VDS, AS\_VQSLOD is included but incorrect. In the reduced SNP & Indel callset VCF, Hail MT, and PLINK files, AS\_VQSLOD is dropped. AS\_VQSLOD is correct in the cohort builder.

Affects:

- srWGS SNP & Indel samples: VDS

Suggested action:

- Ignore (or drop) AS\_VQSLOD from any analyses involving the VDS.

- Update your analyses to not use this annotation

Remediation:

- We will provide the correct AS\_VQSLOD in future releases.

## Known Issue #9: Smaller callset ChrM VCFs are empty

We provide the srWGS SNP and Indel callset in VCF format over limited genomic regions, sharded by chromosome. The chrM VCFs are empty and do not contain any data. We cross-checked the complete callset VDS and verified that there are no calls on chrM.

Affects:

- srWGS SNP & Indel samples: VCF format

Suggested action:

- No action.
- If you are using the srWGS SNP & Indel smaller callsets in VCF format, do not include the chrM.vcf.bz file in your analysis.

Remediation:

- We will not release these empty files in the next release.

## Known issue #10: srWGS SNP & Indel VDS and VCFs from the Cohort Browser will have extraneous INFO field (AS\_YNG)

The srWGS joint callset includes AS\_YNG (an INFO field) in the VDS, which should be ignored by researchers. We have removed AS\_YNG from all other genomic srWGS SNP & Indel VCFs and Hail MTs, but it will appear when creating a VCF from the Cohort Browser

Affects:

- srWGS SNP & Indel VCF files created from the Cohort Browser
- srWGS SNP & Indel variants: VDS

Suggested Action:

- Do not include the AS\_YNG field in any analyses when creating a VCF from the Cohort Browser.
- Ignore (or drop) AS\_YNG from any analyses involving the VDS.

Remediation:

- We will update the VDS to remove the AS\_YNG field (ETA 2024)
- We will update the Cohort Browser data to remove the AS\_YNG field (ETA 2023)

## Known issue #11: srWGS SNP & Indel variant calls on chromosome Y need additional filtering

We see variants with heterozygous calls in chromosome Y, which cannot be correct germline calls. After manual review, we believe that regions of chromosome Y are prone to misalignment artifacts (low mappability). This will cause heterozygous calls in chrY that are likely artifacts. We have not investigated whether these are somatic mutations.

Affects:

- srWGS SNP & Indel variants: VDS, VCF, Hail MT formats

Suggested Action:

- If you do not use variant calls on chrY, then no action.
- Otherwise, we recommend that you use AD, GQ, and GT to filter variants on chromosome Y.

Remediation (ETA 2023):

- We will provide a set of regions (via a BED file) that researchers can use to mask regions of the genome with poor calling accuracy for chromosome Y. It is not currently available with the v7 release.

## Known Issue #12: QUAL information has been removed for srWGS SNP & Indel variants

The srWGS SNP & Indel variants are now released in VDS format (see the VDS article on the User Support Hub [\[1\]](#)). In the VDS format, the actual QUALApprox annotation is not included. This will affect other srWGS short variant files, such as the smaller callsets (eg, exome). The filters used for srWGS SNP & Indel variants are the same and correct, but the annotation is not included.

Affects:

- srWGS SNP & Indel variants: VDS, VCF, Hail MT, and PLINK formats

Suggested action:

- Use the filter field to determine the quality of variants
- If this annotation is important to your (current or future) analyses, please contact the User Support Team [\[1\]](#). We do not plan to remediate this unless we hear from researchers.

Remediation:

- This will not be remediated, this known issue will move to a FAQ for all future releases.

## Known Issue #13: srWGS callset using new convention for genotype filtering flag

The srWGS SNP & Indel variants are now released in VDS format (see the VDS article on the User Support Hub [\[1\]](#)). In order to reduce the genotype metadata stored and reduce the size of the srWGS SNP & Indel variant dataset, the genotype filtering reason is no longer available. Genotype filtering, which is in the VCF and the VDS as the FT annotation, is reported as PASS, FAIL, or “.”. Treat “.” as PASS. In previous AoU releases, we reported more filtering information, including annotations about VQSR filtering.

Affects:

- srWGS SNP & Indel variants: VDS, VCF, PLINK, and Hail MT formats

Suggested action:

- Update your analyses to check the genotype filter (FT) for “FAIL” instead of “low\_VQSLOD\_INDEL” or “low\_VQSLOD\_SNP” to determine whether or not to use a srWGS SNP & Indel variant in your analysis.

Remediation:

- This will not be remediated, this known issue will move to a FAQ for all future releases.

## Known Issue #14: Data processing issue affecting array data

**Update as of April 18, 2024.** We identified a sample swap incident that affects 63 samples with array data in the v7 release.

Affects:

- Array samples: IDATs, VCFs, PLINK bed, and Hail MT

Suggested action:

- Update to the new 7.1 merged array Hail MT or PLINK bed files.
- Remove affected single samples IDAT or VCFs from analysis. We provided a list file containing the research IDs of the affected samples.

Remediation:

- We generated a new array merged Hail MT and merged PLINK bed dataset, released April 18, 2024.

## Known Issue #15: Array and srWGS data with bone marrow transplant history

**Update as of April 18, 2024.** The DRC has identified 1430 participants with a history of an allogeneic bone marrow transplant or a bone marrow transplant of an unknown type. This affects 1430 array samples and 12 srWGS samples.

Affects:

- Array samples: IDATs, VCFs, PLINK bed, and Hail MT formats
- srWGS samples: CRAMs, VDS, VCFs, PLINK bed, and Hail MT formats

Suggested action:

- Update to the new 7.1 merged array Hail MT or PLINK bed files.
- Remove affected samples from your analysis. We provide two list files containing the research IDs of the affected samples, one file for srWGS and one file for array.

Remediation:

- We generated a new array merged Hail MT and merged PLINK bed dataset, released April 18, 2024.

# FAQ

1. Why do you fail samples based on contamination rate for srWGS, but not for array samples?

srWGS analyses (e.g., mosaicism) rely on other signals, such as read counts, which are affected by contamination. Low rates of contamination do not affect array calls and problematic levels of contamination will be reflected in the array call rate.

2. Did you remove samples from participants with bone marrow transplants?

Yes, we removed both array and srWGS samples associated with participants that have received bone marrow transplants from allogeneic transplantation (transplantation from another person), according to the corresponding electronic health record (EHR) and survey responses provided by participants (Overall Health). We did not remove samples who received bone marrow transplants from autologous transplantation (transplantation from themselves).

3. Are all samples in the srWGS joint callset sourced from blood?

Yes. Although the program does have saliva srWGS samples, we did not include these samples in the v7 release. Once we identify any batch effects between saliva and blood samples (ETA 2023), we will reassess the inclusion of saliva samples in the joint srWGS callset. If we decide that the batch effects will have minimal impact, we will include saliva samples in the srWGS joint callsets in 2024.

**Update as of April 18, 2024.** The DRC has identified 1430 participants in the v7 data release with a history of allogeneic transplantation (transplantation from another person) or a bone marrow transplant of an unknown type. Please see [Known Issue #15](#) for more information.

# References

- [1] **All Of Us User Support Hub** <https://aousupporthelp.zendesk.com/hc/en-us>
- [2] Illumina GenCall Data Analysis Software. (n.d.). Retrieved October 21, 2021, from [https://www.illumina.com/Documents/products/technotes/technote\\_gencall\\_data\\_analysis\\_software.pdf](https://www.illumina.com/Documents/products/technotes/technote_gencall_data_analysis_software.pdf)
- [3] CollectArraysVariantCallingMetrics (Picard), Retrieved October 21, 2021, from <https://gatk.broadinstitute.org/hc/en-us/articles/360037593871-CollectArraysVariantCallingMetrics-Picard->
- [4] G. Jun et al., **Detecting and Estimating Contamination of Human DNA Samples in Sequencing and Array-Based Genotype Data**, American journal of human genetics doi:10.1016/j.ajhg.2012.09.004 (volume 91 issue 5 pp.839 - 848)
- [5] E Venner, D Muzny, et al., **Whole-genome sequencing as an investigational device for return of hereditary disease risk and pharmacogenomic results as part of the All of Us Research Program**, *Genome Medicine* (2022). <https://doi.org/10.1186/s13073-022-01031-z>
- [6] **Detecting sample swaps with Picard tools – GATK**. (n.d.). Retrieved October 21, 2021, from <https://gatk.broadinstitute.org/hc/en-us/articles/360041696232-Detecting-sample-swaps-with-Picard-tools>
- [7] Pedersen and Quinlan, **Who's Who? Detecting and Resolving Sample Anomalies in Human DNA Sequencing Studies with Peddy** The American Journal of Human Genetics (2017) <http://dx.doi.org/10.1016/j.ajhg.2017.01.017>
- [8] Zhang F, et al. **Ancestry-agnostic estimation of DNA sample contamination from sequence reads**. *Genome Research* (2020). <https://doi.org/10.1101/gr.246934.118>
- [9] **Phred-scaled quality scores – GATK**. (n.d.). Retrieved January 31, 2022, from <https://gatk.broadinstitute.org/hc/en-us/articles/360035531872-Phred-scaled-quality-scores>.
- [10] Van der Auwera GA & O'Connor BD. (2020). **Genomics in the Cloud: Using Docker, GATK, and WDL in Terra (1st Edition)**. O'Reilly Media. P.400
- [11] **gnomAD v3.1 New Content, Methods, Annotations, and Data ....** (n.d.). Retrieved February 1, 2022, from <https://gnomad.broadinstitute.org/news/2020-10-gnomad-v3-1-new-content-methods-annotations-and-data-availability>.
- [12] Van der Auwera GA & O'Connor BD. (2020). **Genomics in the Cloud: Using Docker, GATK, and WDL in Terra (1st Edition)**. O'Reilly Media. P.166
- [13] **Relatedness - Hail**. (n.d.). Retrieved October 21, 2021, from [https://hail.is/docs/0.2/methods/relatedness.html#hail.methods.pc\\_relate](https://hail.is/docs/0.2/methods/relatedness.html#hail.methods.pc_relate).
- [14] **Which training sets arguments should I use for running VQSR ....** (n.d.). Retrieved February 1, 2022, from <https://gatk.broadinstitute.org/hc/en-us/articles/4402736812443-Which-training-sets-arguments-should-I-use-for-running-VQSR->.
- [15] **Resource bundle – GATK**. (n.d.). Retrieved February 1, 2022, from <https://gatk.broadinstitute.org/hc/en-us/articles/360035890811-Resource-bundle>.
- [16] **The Omni Family of Microarrays**. (n.d.). Retrieved February 16, 2022, from [https://www.illumina.com/Documents/products/datasheets/datasheet\\_gwas\\_roadmap.pdf](https://www.illumina.com/Documents/products/datasheets/datasheet_gwas_roadmap.pdf).

- [17] International HapMap Consortium. **The International HapMap Project**. Nature. 2003 Dec 18;426(6968):789-96. doi: 10.1038/nature02168. PMID: 14685227.
- [18] The 1000 Genomes Project Consortium, **A global reference for human genetic variation**, Nature 526, 68-74 (01 October 2015) doi:10.1038/nature15393
- [19] Mills R.E. et al. **An initial map of insertion and deletion (INDEL) variation in the human genome**. Genome Res. 2006;16:1182–1190. doi:10.1101/gr.4565806.
- [20] Krusche, P., Trigg, L., Boutros, P.C. et al. **Best practices for benchmarking germline small-variant calls in human genomes**. Nat Biotechnol 37, 555–560 (2019).  
<https://doi.org/10.1038/s41587-019-0054-x>
- [21] **PacBio structural variant calling and analysis tools (PBSV)**, Retrieved March 3, 2023, from <https://github.com/PacificBiosciences/pbsv>.
- [22] Li, H. and Handsaker, B. et al. “The Sequence Alignment/Map format and SAMtools.” Bioinformatics, 25 (2009): 2078–2079, <https://doi.org/10.1093/bioinformatics/btp352>
- [23] Wenger, A.M., Peluso, P., Rowell, W.J. et al. Accurate circular consensus long-read sequencing improves variant detection and assembly of a human genome. Nat Biotechnol 37, 1155–1162 (2019). <https://doi.org/10.1038/s41587-019-0217-9>
- [24] **Pacbio glossary**, Retrieved March 8, 2023, from <https://www.pacb.com/wp-content/uploads/2015/09/Pacific-Biosciences-Glossary-of-Terms.pdf>
- [25] Li, Heng. “Which Human Reference Genome to Use?” Heng Li's Blog, 13 Nov. 2017, <https://lh3.github.io/>. Accessed 2 Mar. 2023.
- [26] Schneider, Valerie A et al. “Evaluation of GRCh38 and de novo haploid genome assemblies demonstrates the enduring quality of the reference assembly.” Genome research vol. 27,5 (2017): 849-864. doi:10.1101/gr.213611.116
- [27] Rhie A, Nurk S, Cechova M, Hoyt SJ, Taylor DJ, et al. [The complete sequence of a human Y chromosome](#). bioRxiv, 2022.
- [28] Pedersen, B.S. and Quinlan, A.R. “Mosdepth: quick coverage calculation for genomes and exomes”, Bioinformatics, 34(2018):867–868 <https://doi.org/10.1093/bioinformatics/btx699>
- [29] Shafin, K., Pesout, T., Chang, PC. et al. “Haplotype-aware variant calling with PEPPER-Margin-DeepVariant enables high accuracy in nanopore long-reads.” Nat Methods 18, 1322–1332 (2021). <https://doi.org/10.1038/s41592-021-01299-w>
- [30] **Sniffles2 (PBSV)**, Retrieved March 3, 2023, from <https://github.com/fritzsedlazeck/Sniffles>
- [31] Yun, T., et al. “Accurate, scalable cohort variant calls using DeepVariant and GLnexus” Bioinformatics, 36 (2020): 5582–5589, <https://doi.org/10.1093/bioinformatics/btaa1081>
- [32] Karczewski, K.J., Francioli, L.C., Tiao, G. et al. **The mutational constraint spectrum quantified from variation in 141,456 humans**. Nature 581, 434–443 (2020).  
<https://doi.org/10.1038/s41586-020-2308-7>
- [33] M'Charek, A. **The Human Genome Diversity Project: An Ethnography of Scientific Practice** (Cambridge Studies in Society and the Life Sciences). Cambridge: Cambridge University Press. (2005) doi:10.1017/CBO9780511489167
- [34] Ho, TK . **Random Decision Forests**. Proceedings of the 3rd International Conference on Document Analysis and Recognition, Montreal, QC, 14–16 August 1995. pp. 278–282.
- [35] [Scikit-learn: Machine Learning in Python](#), Pedregosa et al., Journal of Machine Learning Research 12, pp. 2825-2830, (2011).
- [36] **Downloads | gnomAD**. (n.d.). Retrieved February 1, 2021, from

<https://gnomad.broadinstitute.org/downloads#v3-hgdp-1kg>.

[37] Frankish A, Diekhans M, Jungreis I, et al. **GENCODE 2021**, *Nucleic Acids Research*, Volume 49, Issue D1, 8 January 2021, Pages D916–D923, <https://doi.org/10.1093/nar/gkaa1087>

[38] **Genetics - Hail**. (n.d.). Retrieved October 21, 2021, from [https://hail.is/docs/0.2/methods/genetics.html#hail.methods.hwe\\_normalized\\_pca](https://hail.is/docs/0.2/methods/genetics.html#hail.methods.hwe_normalized_pca).

[39] Laurie CC, Doheny KF, et al. **Quality control and quality assurance in genotypic data for genome-wide association studies**. *Genet Epidemiol*. 2010 Sep;34(6):591–602. doi: 10.1002/gepi.20516. PMID: 20718045; PMCID: PMC3061487.

[40] Green RC, Berg JS, Grody WW, Kalia SS, Korf BR, Martin CL, et al. **ACMG recommendations for reporting of incidental findings in clinical exome and genome sequencing**. *Genet Med*. 15:565–574. (2013)

[41] Erdős, P. **On cliques in graphs**, *Israel Journal of Mathematics*, 4 (4): 233–234, (1966), doi:10.1007/BF02771637, MR 0205874, S2CID 121993028

[42] **Long Reads Pipeline**, Retrieved March 7, 2023 from <https://github.com/broadinstitute/long-read-pipelines>




## Appendix A: Ancestry

We computed categorical ancestry for all of the srWGS SNP & Indel samples in *All of Us* and made these available to researchers. These predictions are also the basis for population allele frequency calculations in the Variant Annotation Table (e.g. gvs\_afr\_ac) and data in the Genomic Variants section of the public Data Browser. We used the high-quality set of sites (HQ sites), described in [Appendix J](#), to determine an ancestry label for each sample. The ancestry categories are based on the same labels used in gnomAD [\[32\]](#), Human Genome Diversity Project [\[33\]](#), and 1000 Genomes [\[18\]](#):

- African (afr)
- Latino/Native American/Ad Mixed American (amr)
- East Asian (eas)
- Middle Eastern (mid)
- European (eur) -- Composed of Finnish (FIN) and Non-Finnish European (NFE)
- Other (oth) -- not belonging to one of the other ancestries or is an admixture.
- South Asian (sas)

We trained a random forest classifier [\[34,35\]](#) on a training set of the HGDP and 1kg samples variants on the autosomal exome, obtained from gnomAD [\[36\]](#). This exome was derived from the exon regions of all autosomal, basic, protein-coding transcripts in GENCODE v42 [\[37\]](#). We generated the first 16 principal components (PCs) of the training sample genotypes (using the hwe\_normalized\_pca in Hail [\[38\]](#)) at the high-quality variant sites (see [Appendix J](#)) for use as the feature vector for each training sample. We used the truth labels from the sample metadata, which can be found alongside the VCFs. Note that we do not train the classifier on the samples labeled as “Other.” We use the label probabilities (“confidence”) of the classifier on the other ancestries to determine ancestry of “Other”.

To determine the ancestry of *All of Us* samples, we project the *All of Us* samples into the PCA space of the training data and apply the classifier (see [Figure A.1](#)). Since we do not have truth labels, we can not determine the accuracy of our *All of Us* predictions. As a proxy for the accuracy of our *All of Us* predictions, we look at the concordance between the survey results and the predicted ancestry. The ancestry predictions can be found in [Table A.1](#).



**Figure A.1** -- Ancestry predictions for the *All of Us* srWGS samples plotted on the first two principal components (PC1 on x-axis and PC2 on the y-axis) of the genotype calls.

**Table A.1** -- Breakdown of the computed ancestries in *All Of Us* srWGS data

Computed Ancestry (sorted by percentage)	Count (percentage)
European	123072 (50.2%)
African	53944 (22.0%)
Latino/Admixed American	40838 (16.6%)
Other	19289 (7.9%)

East Asian	5381 (2.2%)
South Asian	2342 (1.0%)
Middle Eastern	528 (0.2%)

Total: 245394 (100.0%)

We evaluated the performance of the ancestry predictions using two different test datasets:

1. A holdout set of training samples. We tested performance with and without the “Other” ancestry
  - a. Error rate (incl Other): 0.050
    - i. See [Table A.2](#)
    - ii. Please see [Known Issue #6](#), since the error rate is higher for Middle Eastern (mid) ancestry. Our classifier conflates Middle Eastern and Other.
  - b. Error rate (not incl Other): 0.002
    - i. See [Table A.3](#)

**Table A.2 -- Error rate (incl. Other) on labeled training data using holdout set**

	Predicted						
Actual	AFR	AMR	EAS	EUR	MID	OTH	SAS
AFR	198	0	0	0	0	2	0
AMR	0	50	0	0	0	0	0
EAS	0	0	199	0	0	1	0
EUR	0	0	0	198	0	2	0
MID	0	0	0	0	49	1	0
OTH	0	2	2	3	25	7	8
SAS	0	0	0	0	0	2	198

**Table A.3 -- Error rate (not incl. Other) on labeled training data using holdout set**

	Predicted					
Actual	AFR	AMR	EAS	EUR	MID	SAS
AFR	199	1	0	0	0	0
AMR	0	50	0	0	0	0
EAS	0	0	199	0	0	1

EUR	0	0	0	200	0	0
MID	0	0	0	0	50	0
SAS	0	0	0	0	0	200

2. We measured the concordance of the ancestry predictions against the self-reported race and ethnicity of the *All of Us* samples. The concordance should be worse than the holdout HGDP samples, but this is expected. Self-reported race and ethnicity does not correspond to the populations listed above (race and ethnicities are social constructs and the ancestry predictions are computed from the genotypes). We expect these to be correlated, but not corresponding. This concordance evaluation is meant to discover large errors in our ancestry predictions.

“Concordant” labeling between HGDP/1kg populations and *All of Us* race/ethnicities:

1. African (AFR) → Black
2. Latino/Ad Mixed American (AMR) → Hispanic
3. East Asian (EAS) → Asian
4. Finnish (FIN) → White
5. Middle Eastern (MID) → MENA
6. Non-Finnish European (NFE) → White
7. Other (OTH) → Other (do not include skipped)
8. South Asian (SAS) → Asian

We do not include any samples where the self-reported race/ethnicity is “Skip”, “Prefer not to answer”, or was not filled in. If a participant selected that their race/ethnicity was not a possible selection (“NoneOfThese”), we counted them as “Other”.

Based on the procedure above, the concordance between self-reported race/ethnicity and the ancestry predictions: 0.898.

## Appendix B: Self-reported race/ethnicity

As seen in [Table B.1](#), the race/ethnicity breakdown of the genomic data is similar to all participants *All of Us* CDR release C2022Q4R9. Samples with “Skip” responses include participants that answered “prefer not to answer”, entered blank text, or did not respond to the survey question.

Please see [Known Issue #1](#), as six srWGS samples and 10 array samples are missing CDR data. All other array, srWGS, and lrWGS samples have corresponding survey data ([Appendix C](#)).

**Table B.1 -- Self-reported Race/Ethnicity breakdown of the genomic data**

Self-reported Race/Ethnicity	Survey response counts (%)	Array counts (%)	srWGS counts (%)	lrWGS counts (%)
Asian	13838 (3.3%)	9605 (3.1%)	7422 (3%)	–
Asian, White	1894 (0.5%)	1284 (0.4%)	992 (0.4%)	–
Black	77069 (18.6%)	62514 (20.0%)	50064 (20.4%)	911 (88.7%)
Black, White	2390 (0.6%)	1743 (0.6%)	1351 (0.6%)	54 (4.8%)
Hispanic	64672 (15.6%)	52904 (16.9%)	41938 (17.1%)	–
Hispanic, White	6512 (1.6%)	4718 (1.5%)	3682 (1.5%)	–
Other	15294 (3.7%)	11264 (3.6%)	8709 (3.5%)	49 (4.8%)
Skip	91422 (2.2%)	6872 (2.2%)	5387 (2.2%)	13 (1.3%)
White	222646 (53.8%)	162021 (51.8%)	125843 (51.3%)	–
<b>Total</b>	<b>413457 (100.0%)</b>	<b>312925 (100.0%)</b>	<b>245388 (100.0%)</b>	<b>1027 (100.0%)</b>

## Appendix C: Data type availability with genomic data

Please see [Known Issue #1](#), as six srWGS samples and 10 array samples are missing CDR data. All other array, srWGS, and lrWGS samples have corresponding survey data. The srWGS samples are a subset of the array data. The srWGS SV samples are a subset of the srWGS SNP & Indel data (all srWGS samples with SNP and Indel data have SV data). The long reads samples are a subset of the srWGS SNP & Indel and array data.

Additionally, array ([Table C.1](#)), srWGS SNP & Indel ([Table C.2](#)), and lrWGS ([Table C.3](#)) data have other corresponding non-genomic data. This can be one or more of the following:

- Electronic Health Records (EHR)
- Physical Measurements (PM)
- Participant Provided Information (PPI/surveys)
- Fitbit (FB)

Descriptions of the non-genomic data can be found on the [All of Us Data Sources](#) page.

**Table C.1 -- Array overlap with non-genomic data types**

Data Combination	Description	Participant Count
Array	any Array data	312925
Array and PPI	any Array AND any PPI	312925
Array and PPI and PM	any Array AND any PPI AND any PM	303064
Array and EHR	any Array AND any EHR	255052
Array and PPI and EHR	any Array AND any PPI AND any EHR	255052
Array and PPI and EHR and PM	any Array AND any EHR AND any PM AND any PPI	254203
Array and Fitbit	any Array AND any Fitbit	11763
Array and PPI and Fitbit	any Array AND any PPI AND Fitbit	11763
Array and PPI and PM and Fitbit	any Array AND any PPI AND any PM AND any Fitbit	10442
Array and Fitbit and PPI and EHR	any Array AND any Fitbit AND and PPI AND any EHR	8867
Array and PPI and EHR and PM and Fitbit	any Array AND any EHR AND and PM AND any PPI AND any Fitbit	8778

**Table C.2 -- srWGS SNP & Indel overlap with non-genomic data types**

Data Combination	Description	Participant Count
srWGS	any srWGS data	245388
srWGS and PPI	any srWGS AND any PPI	245388

srWGS and PPI and PM	any srWGS AND any PPI AND any PM	245149
srWGS and EHR	any srWGS AND any EHR	206173
srWGS and PPI and EHR	any srWGS AND any PPI AND any EHR	206173
srWGS and PPI and EHR and PM	any srWGS AND any EHR AND any PM AND any PPI	206109
srWGS and Fitbit	any srWGS AND any Fitbit	8812
srWGS and PPI and Fitbit	any srWGS AND any PPI AND Fitbit	8812
srWGS and PPI and PM and Fitbit	any srWGS AND any PPI AND any PM AND any Fitbit	8798
srWGS and Fitbit and PPI and EHR	any srWGS AND any Fitbit AND and PPI AND any EHR	7445
srWGS and PPI and EHR and PM and Fitbit	any srWGS AND any EHR AND and PM AND any PPI AND any Fitbit	7444

**Table C.3 -- lrWGS overlap with non-genomic data types**

<b>Data Combination</b>	<b>Description</b>	<b>Participant Count</b>
lrWGS	any lrWGS data	1027
lrWGS and PPI	any lrWGS AND any PPI	1027
lrWGS and PPI and PM	any lrWGS AND any PPI AND any PM	1027
lrWGS and EHR	any lrWGS AND any EHR	985
lrWGS and PPI and EHR	any lrWGS AND any PPI AND any EHR	985
lrWGS and PPI and EHR and PM	any lrWGS AND any EHR AND any PM AND any PPI	985
lrWGS and Fitbit	any lrWGS AND any Fitbit	38
lrWGS and PPI and Fitbit	any lrWGS AND any PPI AND Fitbit	38
lrWGS and PPI and PM and Fitbit	any lrWGS AND any PPI AND any PM AND any Fitbit	38
lrWGS and Fitbit and PPI and EHR	any lrWGS AND any Fitbit AND and PPI AND any EHR	37
lrWGS and PPI and EHR and PM and Fitbit	any lrWGS AND any EHR AND and PM AND any PPI AND any Fitbit	37

## Appendix D: Genome Centers and Data and Research Center

Below is the listing of the three Genome Centers (GCs), the Data and Research Center (DRC), and the Biobank.

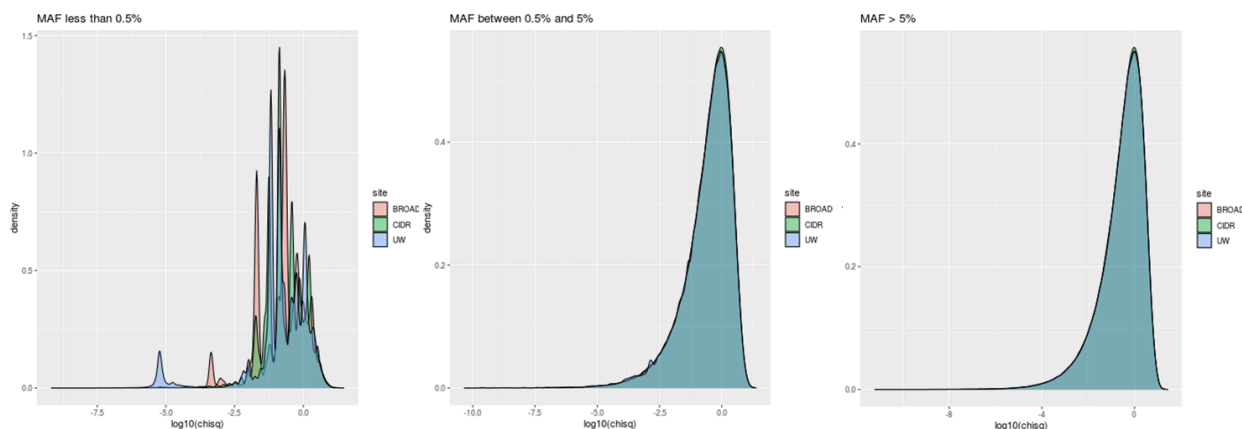
Role	Principal Investigator(s)
Genome Center	Richard Gibbs - Baylor College of Medicine, Johns Hopkins University Eric A. Boerwinkle - Baylor College of Medicine, Johns Hopkins University Kimberly F. Doheny - Baylor College of Medicine, Johns Hopkins University Stacey Gabriel - Broad Institute Gail Jarvik - Northwest Genomics Center at the University of Washington Evan Eichler - Northwest Genomics Center at the University of Washington
Data and Research Center	Paul Harris - Vanderbilt University Medical Center Dan M. Roden - Vanderbilt University Medical Center Melissa Basford - Vanderbilt University Medical Center Anthony Philippakis - Broad Institute David Glazer - Verily Life Sciences
Biobank	Stephen Norman Thibodeau - Mayo Clinic



# Appendix E: Array processing overview

See [Figure E.2](#) for an overview of the array genotyping process for the *All of Us* Research Program. The three GCs used identical array products, scanners, resource files, and genotype calling software. The GCs used the Illumina Global Diversity Array (GDA) (<https://www.illumina.com/products/by-type/microarray-kits/infinium-global-diversity.html>).

For the v7 data release (C2022Q4R9), new cluster definition files (.egt) was created at Johns Hopkins using raw data from 12,983 samples from all 3 genotyping centers (3,782-Broad, 4,342-Johns Hopkins, 4,859-UW) in order to reduce batch effects. Manual review and editing of cluster boundaries was performed for 67,812 assays including all X, MT and Y SNPs, rare variant calls with “new hets” detected by z-call (new hets > 2, total hets >=4, and MAF <=0.0025) GEM trait SNPs, fingerprint sites for array concordance to WGS datasets and all assays within the bed file regions for health-related return of results. 11,916 assays were dropped based on manual review and 75,237 assays were dropped based on call rate <99% and/or cluster separation <0.4. 681 trios were examined for mendelian segregation errors, 15 SNPs were dropped due to >1 mendelian error. A homogeneous subset of 7,511 samples was defined using PCA and MCD (minimum covariance determinant method). Using this homogeneous sample subset, HWE and sex differences in allele frequency were evaluated. 4,005 SNPs were dropped due to Hardy Weinberg equilibrium p-value less than  $10^{-4}$  and 258 SNPs were dropped due to a sex difference in allele frequency of >0.2. \Batch effects were evaluated by comparing allele frequencies between genotyping centers within the homogenous sample subset. Chi-square statistics were Broad 0.73, Johns Hopkins 0.74, UW 0.74.



**Figure E.1** Comparisons shown in [Figure E.1](#) broken out into MAF bins.

	Different Genotypes		Missing_1		Missing_2		Same	
	original	reclustered	original	reclustered	original	reclustered	original	reclustered
CIDR_Broad	0.0045%	0.0003%	0.4283%	0.0337%	0.0533%	0.0135%	99.51%	99.95%
CIDR_UW	0.0053%	0.0004%	0.3671%	0.0337%	0.1767%	0.0484%	99.45%	99.92%
Broad_UW	0.0032%	0.0001%	0.0498%	0.0135%	0.2346%	0.0484%	99.71%	99.94%

**Figure E.2** Data for the program control sample HG001 was compared to evaluate the performance of the new cluster file. When comparing data between the 3 genotyping centers, missing data rates were decreased and concordance rates were increased.

Array product details:

- Bead pool file: GDA-8v1-0\_D1 .bpm
- EGT cluster file: GDA-8v1-1\_A1\_AoUUpdated.08.17.21\_ClusterFile.egt
- gentrain v.3
- reference hg19 (Note: We liftover to hg38 before publishing array data in the RW. The IDAT files are raw files and thus have no reference.)
- gencall cut-off 0.15
- 1,814,226 assays
  - 1,767,452 SNVs
  - 36,839 indels
  - 9,934 IntensityOnly (probes intended only for Copy Number Variant (CNV) calling)

Chemistry: Illumina Infinium LCG using automated protocol

Liquid handling robotics: Various platforms across the genome centers

Scanners: Illumina iSCANs with Automated Array Loader

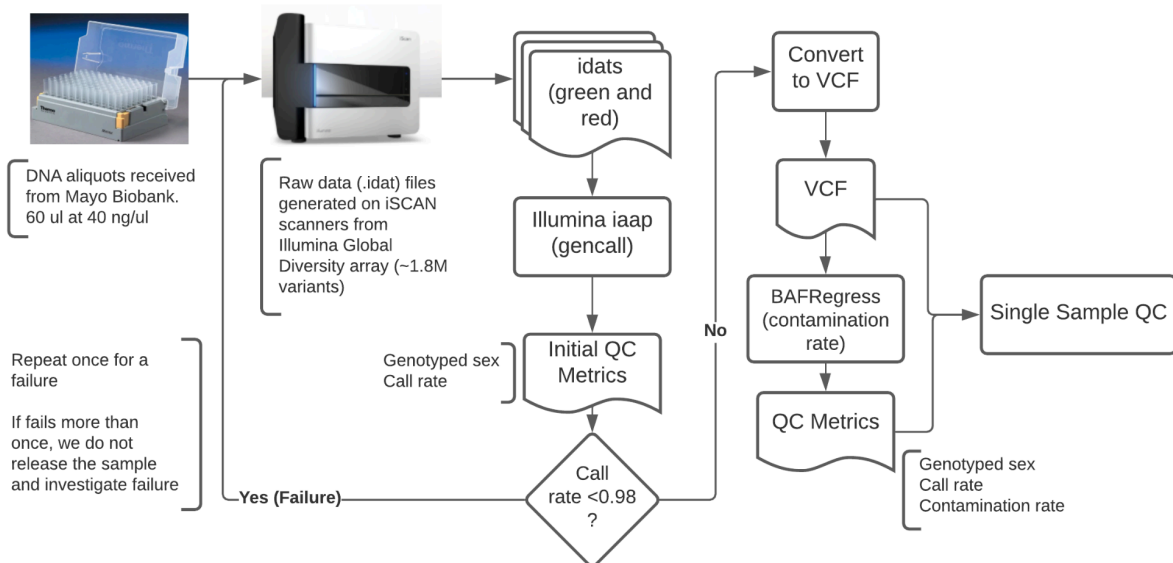
Software:

- Illumina IAAP Version:  
iaap-cli-linux-x64-1.1.0-sha.80d7e5b3d9c1fd9c2e99b472a90652fd3848bbc7.tar.gz
  - IAAP converts raw data (.idat files – 2 per sample) into a single .gtc file per sample using the .bpm file (defines strand, probes sequences, and illumicode address) and the .egt file (defines the relationship between intensities and genotype calls)
- Picard-2.26.4
  - Picard tool, GTCtoVCF, converts the .gtc file into a vcf file.
- BAFRegress version 0.9.3 [\[4\]](#)
  - BAFRegress measures the within species DNA sample contamination using B allele frequency data from Illumina genotyping arrays using a regression model

Quality Control:

Each genome center ran the GDA array under Clinical Laboratory Improvement Amendments (CLIA) compliant protocols. We generated .gtc files and uploaded metrics to in-house Laboratory Information Management Systems (LIMS) systems for quality control review. At batch level (each set of 96 well plates run together in the laboratory at one time), each GC included positive control samples, which were required to have > 98% call rate and >99% concordance to existing data, in order to approve release of the batch of data. At the sample level, the call rate and sex are the key quality control determinants [\[39\]](#). Contamination is also measured using BAFRegress [\[4\]](#) and reported out as metadata. Any sample with a call rate below 98% is repeated one time in the laboratory. Genotyped sex is determined by plotting normalized X versus normalized Y intensity values for a batch of samples [\[39\]](#). Any sample

discordant with 'sex assigned at birth' [reported by an All of Us participant](#) is flagged for further detailed review. If multiple sex discordant samples are clustered on an array or on a 96 well plate, the entire array or plate will have data production repeated. Samples identified with sex chromosome aneuploidies are also reported back as metadata (XXX, XXY, XYY, etc). A final processing status of "PASS," "FAIL" or "ABANDON" is determined before release of data to the DRC. An array sample will PASS if the call rate is > 98% and the genotyped sex and sex assigned at birth are concordant. If we do not have a "male" or "female" for the sex assigned at birth, because the participant reported it as "Intersex", "I prefer not to answer", "none of these fully describe me", or skipped the question, the array sample is marked as PASS. The sex assigned at birth data from the CDR is described in [Appendix F](#). An array sample will FAIL if the genotyped sex and the sex assigned at birth are discordant or if the call rate is less than 98% on the first run of the sample. An array sample will have the status ABANDON if the call rate is less than 98% after at least 2 attempts at the GC.



**Figure E.3** -- Overview of the array processing pipeline.

## Appendix F: Self-reported sex assigned at birth

See [Table F.1](#) for the counts and percentages of participant responses to “What was your biological sex assigned at birth?” in the Basics survey (based on *All of Us* CDR release C2022Q4R9). The CDR code for this question is sex\_at\_birth. These participant responses are used for the participant self-reported sex at birth information used in sex concordance checks.

Please see [Known Issue #1](#), as six srWGS samples and 10 array samples are missing CDR data. All other array, srWGS, and lrWGS samples have corresponding survey data ([Appendix C](#)).

**Table F.1 -- v7 release participants response breakdown to sex assigned at birth question**

V7 release	Array		srWGS SNP & Indel		lrWGS	
Sex assigned at birth responses	counts	percent	counts	percent	counts	percent
Female	185210	59.19	145563	59.32	712	69.33
Male	121173	38.72	94756	38.61	294	28.63
Other responses*	6542	2.09	5069	2.07	21	2.04
Total	312925		245388		1027	

Percentages may not add to 100 due to rounding.

\*The \*Other responses count includes any or no response for sex\_at\_birth. The available options in the CDR are “I prefer not to answer”, “None of these fully describe me”, “Intersex”, “No matching concept”, and “PMI: Skip”. “No matching concept” and “PMI: Skip” are separate counts both referring to no response for sex\_at\_birth. These are separate because participants in “No matching concept” did select a gender option for this survey question. The terms used here are the Concept Names as they appear in the CDR.

## Appendix G: *All of Us* Hereditary Disease Risk genes

The following gene symbols are in the *All of Us* Hereditary Disease Risk (AoUHDR) genes. We have additional srWGS QC criteria in the regions covered by these genes, described in [Table 2](#) of the main text. In the v7 callset, the AoUHDR genes are the same as the American College of Medical Genetics and Genomics' list of 59 genes where incidental findings should be reported (ACMG59) [\[40\]](#). The AoUHDR gene list may change in future releases.

ACTA2, ACTC1, APC, APOB, ATP7B, BMPR1A, BRCA1, BRCA2, CACNA1S, COL3A1, DSC2, DSG2, DSP, FBN1, GLA, KCNH2, KCNQ1, LDLR, LMNA, MEN1, MLH1, MSH2, MSH6, MUTYH, MYBPC3, MYH11, MYH7, MYL2, MYL3, NF2, OTC, PCSK9, PKP2, PMS2, PRKAG2, PTEN, RB1, RET, RYR1, RYR2, SCN5A, SDHAF2, SDHB, SDHC, SDHD, SMAD3, SMAD4, STK11, TGFBR1, TGFBR2, TMEM43, TNNT2, TP53, TPM1, TSC1, TSC2, VHL, and WT1

# Appendix H: DRAGEN invocation parameters

[Table H.1](#) summarizes the parameters used by the GCs to generate GVCFs, contamination estimates, and sex ploidy calls from the DRAGEN for srWGS data.

**Table H.1 DRAGEN 3.4.12 parameters run at all GCs**

Parameter	Parameter Value	Description
-f	n/a	Overwrite if output exists
-r	<hg38-ref-dir>	The reference to use
--fastq-list	<path-to>/fastq_list.csv	A list of fastq files to use as input for this sample
--fastq-list-sample-id	<sampleID>	The sample ID to use for naming this sample
--output-directory	<output-dir>	The location of the final output files
--intermediate-results-dir	<int-results-dir>	The location to write intermediate outputs
--output-file-prefix	[CenterID]_[Biobankid_Sampleid]_[LocalID:optional]_[Rev#]	Standardized naming prefix for each output file
--enable-variant-caller	TRUE	Turn on variant call outputs
--enable-duplicate-marking	TRUE	Mark duplicate reads during alignment
--enable-map-align	TRUE	Produce an alignment from unaligned read input
--enable-map-align-output	TRUE	Store the output of the alignment
--output-format	CRAM	Store the alignment as a CRAM file
--vc-hard-filter	DRAGENHardQUAL:all:QUAL<5.0;LowDepth:all:DP<=1'	This parameter setting changes the threshold on the quality to 5.
--vc-frd-max-effective-depth	40	Setting this parameter puts a limit on the penalty value that is applied for variant calls that deviate from the expected 50% allele fraction for heterozygous variants.
--qc-cross-cont-vcf	<path-to/SNP_NCBI_GRCh38.vcf>	Marker sites to use for contamination estimation
--qc-coverage-region-1	<path-to/wgs_coverage_regions.bed>	Regions to use for coverage analysis (whole genome)
--qc-coverage-reports-1	cov_report	The type of reports requested for qc-coverage-region-1
--qc-coverage-region-2	<path-to/HDRR_regions.bed>	Regions to use for coverage analysis (HDR reportable regions)
--qc-coverage-reports-2	cov_report	The type of reports requested for qc-coverage-region-2
--qc-coverage-region-3	<path-to/PGx_regions.bed>	Regions to use for coverage analysis

		(PGx reportable regions)
--qc-coverage-reports-3	cov_report	The type of reports requested for qc-coverage-region-3

# Appendix I: Samples used in the Sensitivity and Precision Evaluation

In order to calculate the sensitivity and precision of the srWGS SNP and Indel joint callset, we included four well-characterized samples in the v7 callset ([Table I.1](#)). We sequenced the NIST reference materials (DNA samples) from Genome in a Bottle (GiaB) and performed variant calling as described in the main text. We used the corresponding published set of variant calls for each sample as the ground truth in our sensitivity and precision calculations [\[20\]](#).

Please note that the control samples do not appear in the data released to researchers.

**Table I.1 -- Samples used in sensitivity and precision evaluation**

Control Sample	Ground Truth	Genome Center	GVCF origin	Notes
HG-001	GiaB	BI	DRAGEN 3.4.12	NA12878
HG-003	GiaB	UW	DRAGEN 3.4.12	Ashkenazi Trio NA24149 - Father
HG-004	GiaB	BI	DRAGEN 3.4.12	Ashkenazi Trio NA24143 - Mother
HG-005	GiaB	BI	DRAGEN 3.4.12	Han ancestry NA24631- Son

Genome Center:

BI -- Broad Institute

UW -- University of Washington



## Appendix J: High quality site determination (srWGS)

In order to do relatedness and ancestry checks, we identified a corpus of sites that can be called accurately in both our ancestry training set (HGDP+1KG) and our target data (*All of Us* srWGS callset). We used a similar methodology that gnomAD used to determine high-quality sites [\[11\]](#):

1. Autosomal, bi-allelic single nucleotide variants (SNVs) only
2. Allele frequency > 0.1%
3. Call rate > 99%
4. LD-pruned with a cutoff of  $r^2 = 0.1$

Our aim was to assemble a set of independent sites where we can be confident of the accuracy.

We identified 150229 high-quality (HQ) sites in the v7 callset. These were HQ sites in both the HGDP+1kg training VCF and the *All of Us* v7 callset. A sites-only VCF of the HQ sites is available in the RW (access required).

## Appendix K: Relatedness (srWGS)

We used the Hail `pc_relate` function to determine the kinship score to report any pairs with a kinship score over 0.1. This analysis was done with the srWGS SNP and Indel data and the lrWGS SNP and Indel data. The kinship score is half of the fraction of the genetic material shared (ranges from 0.0 - 0.5).

- Parent-child or siblings: 0.25
- Identical twins: 0.5




Please see the [Hail `pc\_relate` function \[13\]](#) documentation for more information, including interpretation.

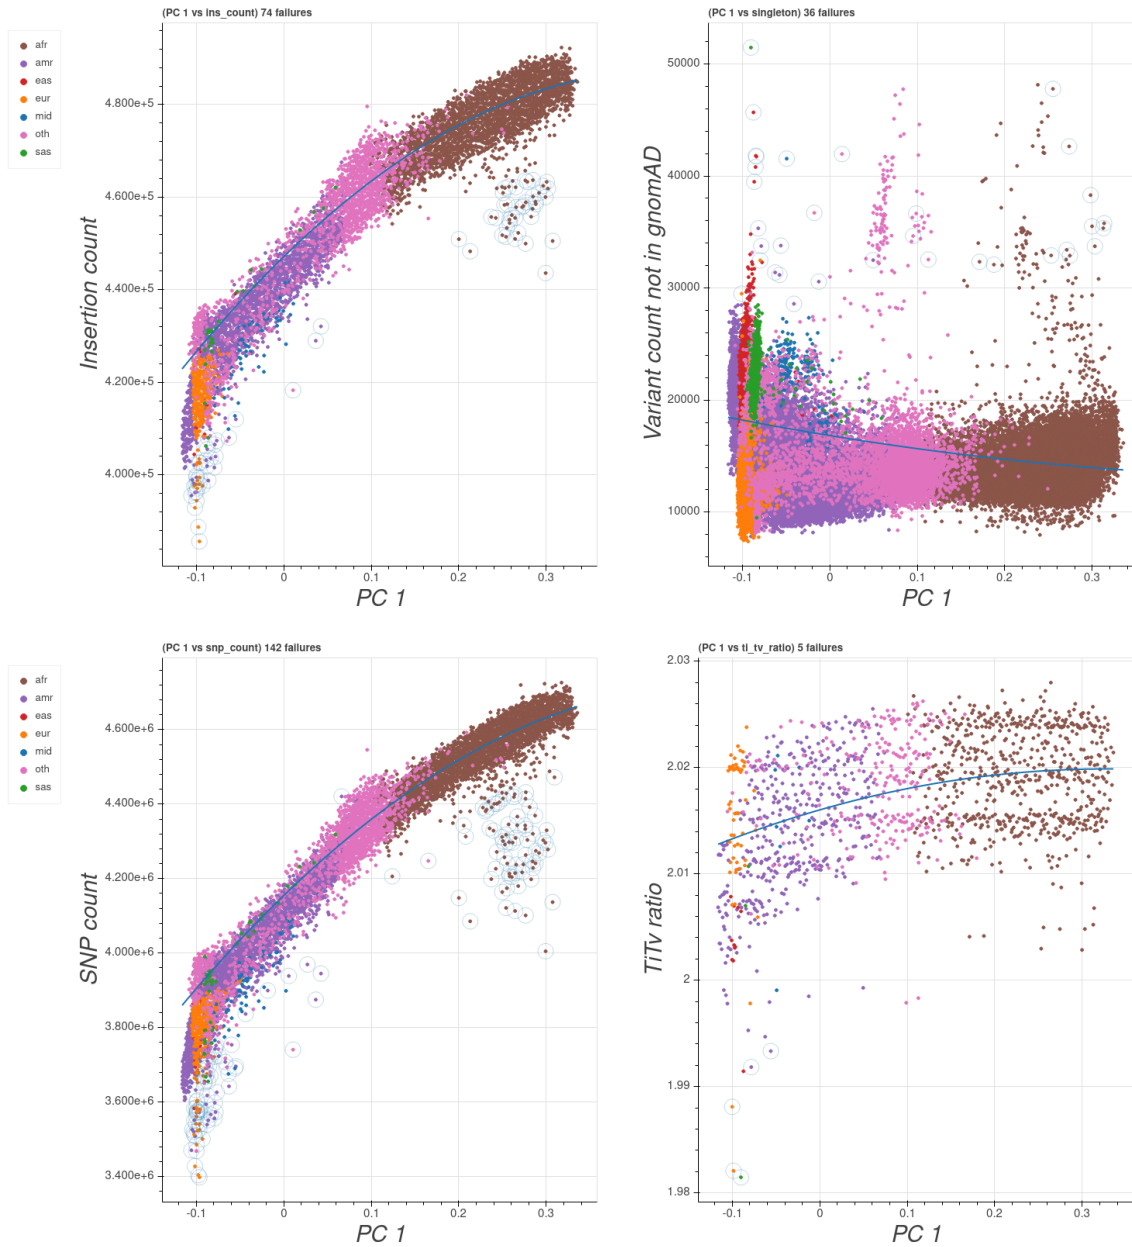
We will determine the [maximal independent set \[41\]](#) for related samples to minimize the number of samples that would need pruning. Using the HQ sites identified in [Appendix J](#), researchers can remove first and second degree relatives.

We estimated 19,374 related pairs and 15,376 samples in the maximal independent set for kinship scores above 0.1. The sample pairs, with kinship score, and the set are available in the RW (access required).

## Appendix L: Plots of the first principal component against population outlier QC metrics

[Figure L.1](#) contains the plots of the first principal component against metrics used for determining [sample population outliers](#) in srWGS sample QC. Note that we use sixteen principal components for determining which samples should be flagged for being outliers in a metric. The blue line shows the linear regression fit in the first dimension (residuals are calculated as the distance from this hyperplane). The failure count over these plots will sum higher than the 551 flagged samples, since samples can get flagged for multiple criteria. Please see the next page for [Figure L.1](#).





**Figure L.1** -- Sample population outlier plots for eight metrics (see [Population Outlier Flagging](#)). Each metric (y-axis) is plotted against the first (of sixteen) principal components (x-axis). Outliers are identified by regressing out the principal components and determining if the residual is over 8 MADs from the sample population.

# Appendix M: Long Read Workflow Diagrams

The following figures summarize the workflows utilized to process the AoU Long Read samples. A standard sample will be processed in the following sequence of workflows:

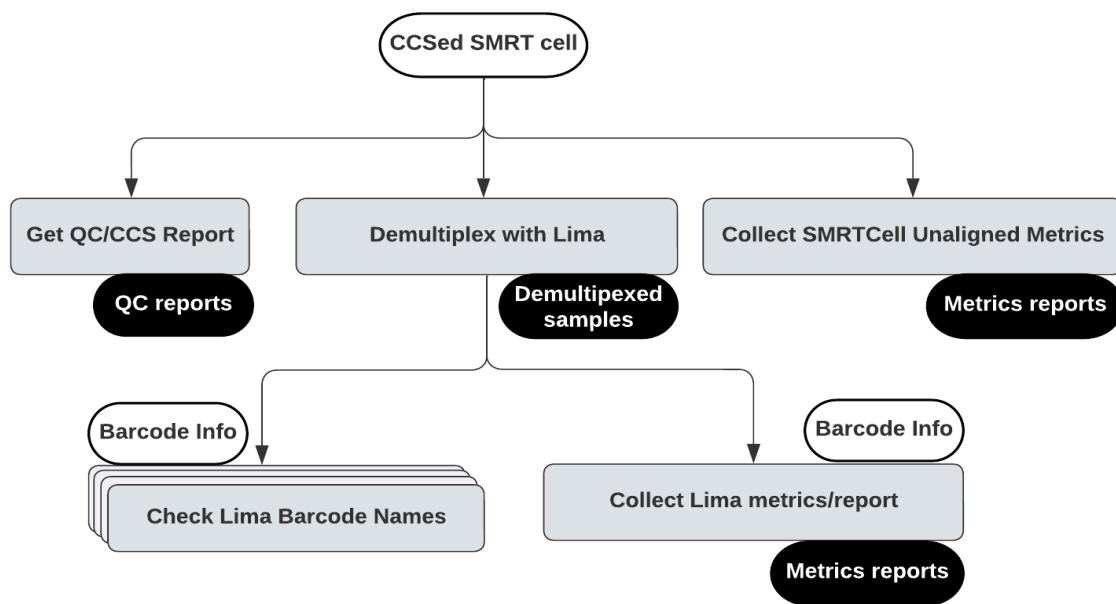
- 1) Preprocessing CCSed SMRT cells (HiFi reads extraction, demultiplex if applicable) [Figure N.1]

[Figure N.1]

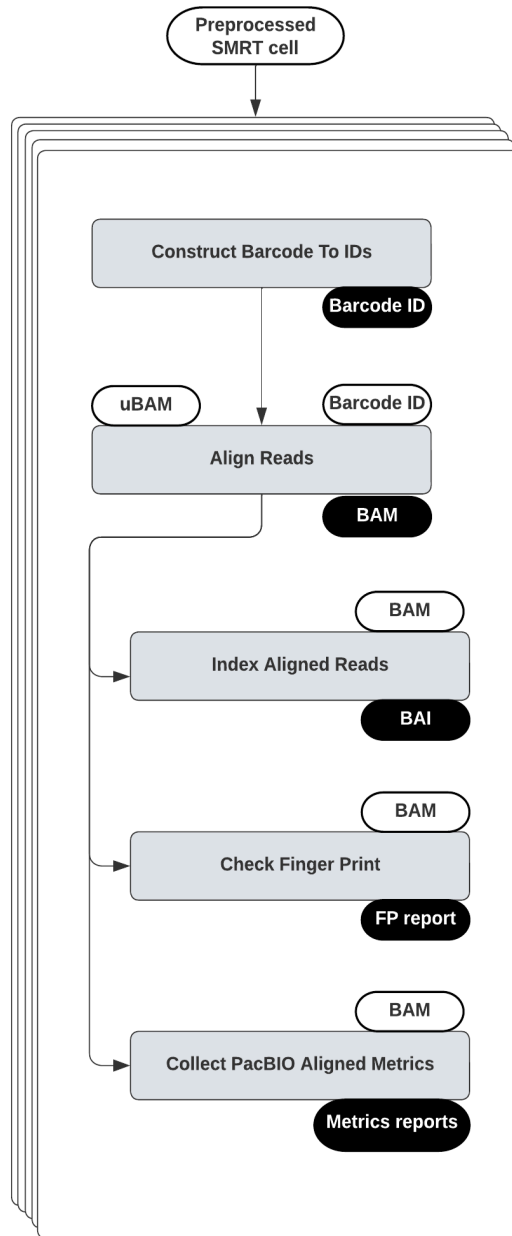
- 2) Post-processing SMRT cells (alignment, QC controls at atomic unit level) [Figure N.2]

- 3) Whole Genome Variant Calling (data aggregation, de novo assembly, and variant calling) [Figures N.3 and N.4]

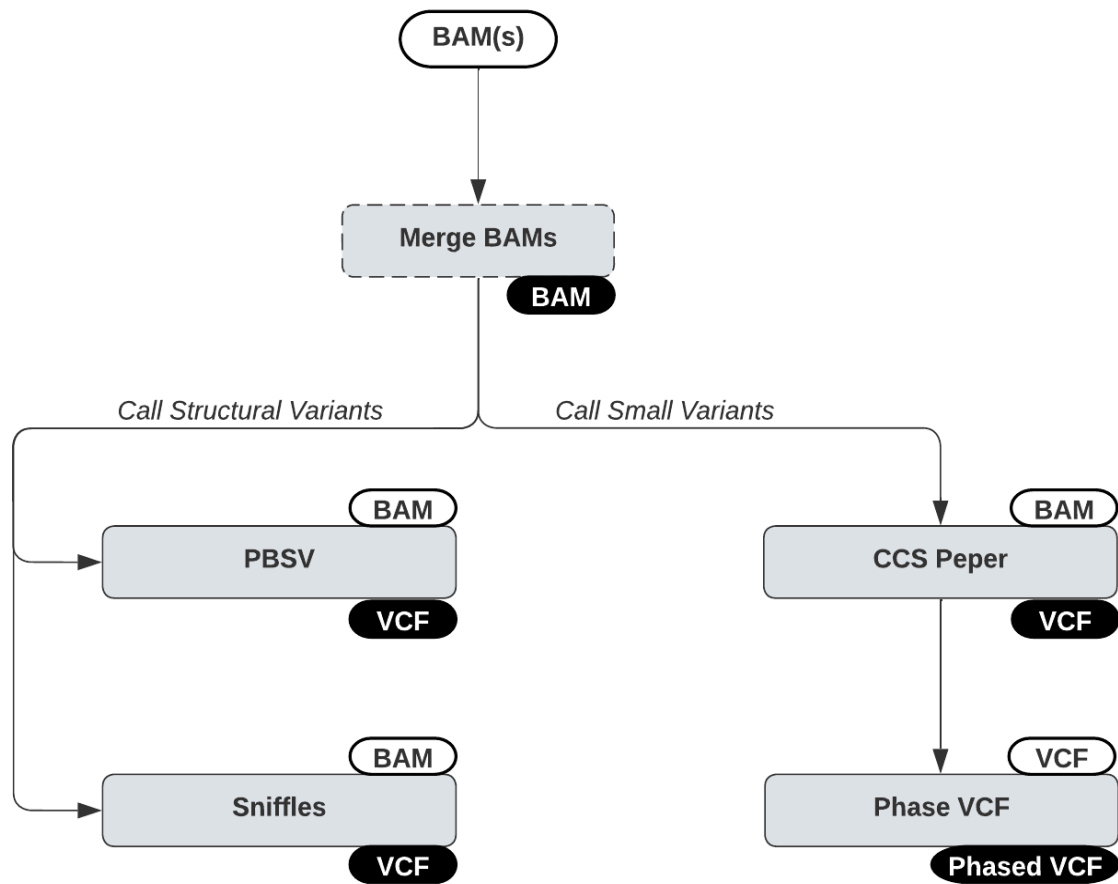
The workflows are available as Workflow Definition Language (WDL) files in a public github repository [\[42\]](#).



**Figure M.1** Preprocessing CCSed SMRT cell. A workflow for preprocessing barcoded (potentially multiplexed) SMRTCell. The cell is assumed to be CCSed on-instrument, and the whole data folder is mirrored onto a cloud bucket.

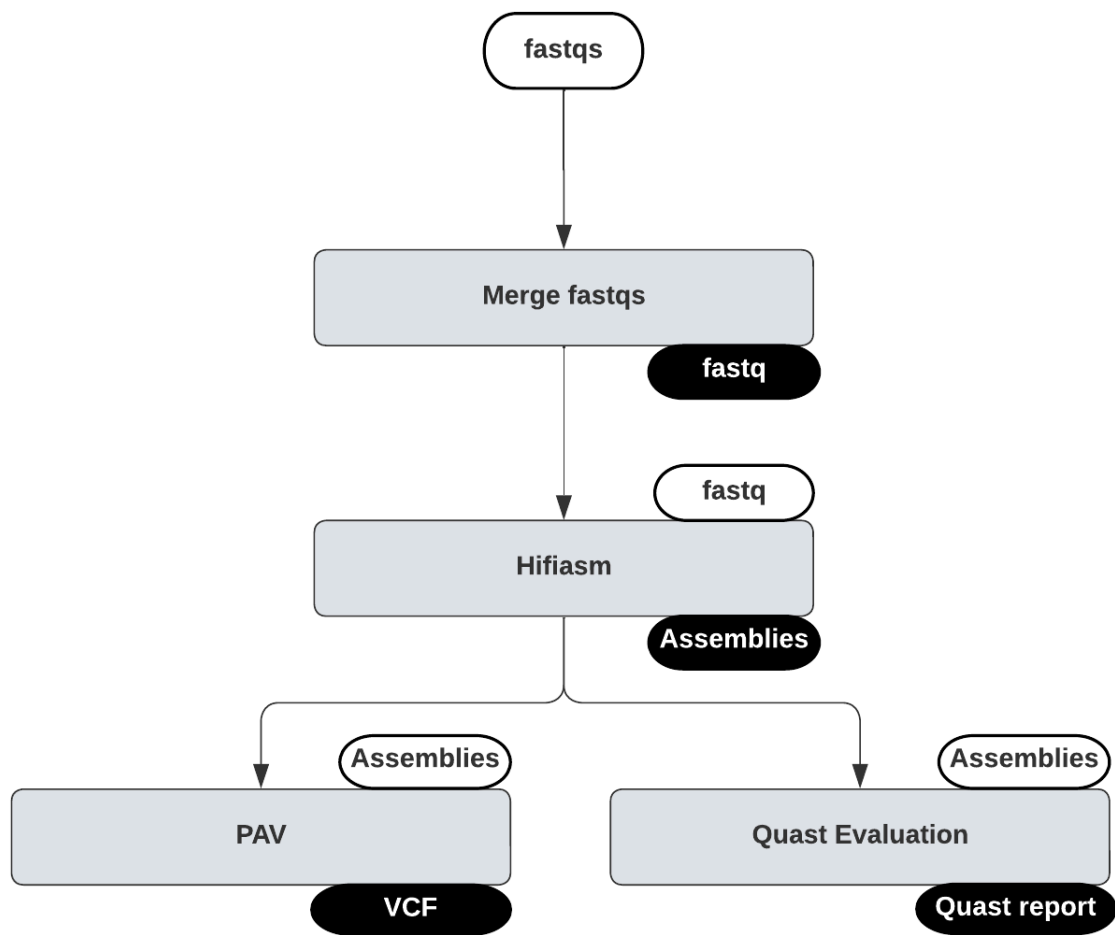


**Figure M.2** Post-processing demultiplexed SMRTCell (CCS already performed). The various IDs are assumed to be in-phase with the barcode names. Alignment to the two references (grch38\_noalt and T2Tv2.0) are done independently and only the grch38\_noalt version goes through the three sample SMRT cell (pre-sample-aggregation) QC processes.



**Figure M.3** Single sample whole genome variant calling. The workflow merges information from each sample's SMRT cell into a single BAM prior to variant calling. We run the entire workflow separately for each reference.





**Figure M.4** HiFi FASTQ files for a sample are merged before sent for de novo assembly with Hifiasm. The haplotype-resolved assembly is then evaluated with QUAST. This assembly will then be used by PAV for calling variants.

# Appendix N: IrWGS analysis versions and parameters

**Table N.1 – IrWGS pipeline software versions and parameters**

Software	Version used	Functionality	Invocation parameters
extracthifi	1.0.0	Extracting HiFi reads from CCS bam.	extracthifi <unaligned.ccs.bam> <unaligned.hifi.bam>
lima	2.6.0	SMRT cell demultiplexing.	lima --dump-removed --split-bam-named --hifi-preset SYMMETRIC <unaligned.hifi.bam> Sequel_16_barcode_v3.fasta <demux.bam>
pbmm2	1.4.0	HiFi reads alignment.	pbmm2 align <unaligned.hifi.bam><reference.fasta> --preset CCS --sample <sample_name> --strip --sort --unmapped
samtools	1.13	BAM aggregation and conversion to FASTQ.	<u>Aggregation</u> samtools merge -p -c --no-PG -@ 2 --write-index -o <agg.bam> <input.bam>[,<input.bam>, ...]  <u>Conversion</u> samtools fastq <input.bam> <output.fastq>
Hifiasm	0.16.1	<i>de novo</i> assembly.	<u>Primary and alt assembly</u> hifiasm -o <output_prefix> -t <cpu_cores_to_use> -primary <input.fastq>[,<input.fastq>, ...]  <u>Haplotype resolved assembly</u> hifiasm -o <output_prefix> -t <cpu_cores_to_use> <input.fastq>[,<input.fastq>, ...]
pbsv	2.6.0=h9ee0642_0	Single sample SV calling per chromosome. After this step, chromosomes are merged.	pbsv discover --tandem-repeats <trf.bed> <aligned.bam> <output.svsig.gz>  pbsv call -ccs <reference.fasta> <input.svsig.gz> <output.vcf>
sniffles2	2.0.6	Single sample SV calling	sniffles -i <input.bam> --minsvlen 50 --sample-id <sample_id> --vcf <output.vcf> --snf <output.snf>

pav (the tool)	Branch aou with hash fa43453 in repo <a href="https://github.com/EichlerLab/pav">https://github.com/EichlerLab/pav</a>	The specific pav docker that we ran	
pav (WDL pipeline)		Single sample SV and small variant calling from assembly	pav pipeline at <a href="https://github.com/broadinstitute/pav-wdl/tree/sh_more_resources_pet">https://github.com/broadinstitute/pav-wdl/tree/sh_more_resources_pet</a> It is currently in development. We ran the pipeline in the state that is documented in the git commit hash 5558ebdbd0be3f2eb722b10774a1e305a20fa569
Pepper	Docker kishwars/pepper_deepvariant:r0.4.1	Prepping BAM for small variant calling	run_pepper_margin_deepvariant \ call_variant \ -b <input.bam> \ -f <reference_fasta> \ -s <sample_name> \ -o <output_dir> \ -p <output_prefix> \ --phased_output \ --ccs
DeepVariant	1.3.0	Single sample SNP and Indel variant calling	/opt/deepvariant/bin/run_deepvariant \ --model_type=PACBIO \ --ref=<reference_fasta> \  --reads=<pepper.prepared.input.bam> \ --output_vcf=<output.vcf.gz> \  --output_gvcf=<output.g.vcf.gz> \ --use_hp_information
Margin	Docker kishwars/pepper_deepvariant:r0.4.1	Single sample SNP and Indel phasing	margin phase \ <input.haplotagged.bam> \ <reference_fasta> \ <unphased_vcf> \  /opt/margin_dir/params/misc/allParams.phase_vcf.json \ -M \ -o <output_dir>/<prefix>