

TECHNICAL NOTE

CNVpytor: a tool for copy number variation detection and analysis from read depth and allele imbalance in whole-genome sequencing

Milovan Suvakov¹, Arijit Panda¹, Colin Diesh², Ian Holmes² and Alexej Abyzov^{1,*}

¹Department of Quantitative Health Sciences, Center for Individualized Medicine, Mayo Clinic, Rochester, MN 55905, USA and ²Department of Bioengineering, University of California, Berkeley, CA 94720, USA

*Correspondence address. Alexej Abyzov, Department of Quantitative Health Sciences, Center for Individualized Medicine, Mayo Clinic, 200 1st street SW, Rochester, MN 55905, USA. E-mail: abyzov.alexej@mayo.edu

Abstract

Background: Detecting copy number variations (CNVs) and copy number alterations (CNAs) based on whole-genome sequencing data is important for personalized genomics and treatment. CNVnator is one of the most popular tools for CNV/CNA discovery and analysis based on read depth. **Findings:** Herein, we present an extension of CNVnator developed in Python—CNVpytor. CNVpytor inherits the reimplemented core engine of its predecessor and extends visualization, modularization, performance, and functionality. Additionally, CNVpytor uses B-allele frequency likelihood information from single-nucleotide polymorphisms and small indels data as additional evidence for CNVs/CNAs and as primary information for copy number–neutral losses of heterozygosity. **Conclusions:** CNVpytor is significantly faster than CNVnator—particularly for parsing alignment files (2–20 times faster)—and has (20–50 times) smaller intermediate files. CNV calls can be filtered using several criteria, annotated, and merged over multiple samples. Modular architecture allows it to be used in shared and cloud environments such as Google Colab and Jupyter notebook. Data can be exported into JBrowse, while a lightweight plugin version of CNVpytor for JBrowse enables nearly instant and GUI-assisted analysis of CNVs by any user. CNVpytor release and the source code are available on GitHub at <https://github.com/abyzovlab/CNVpytor> under the MIT license.

Keywords: copy number variations; copy number alternations; whole-genome sequencing; Python

Introduction

The continuous reduction of cost has enabled whole-genome sequencing (WGS) to be widely used in different research projects and clinical applications. Consequently, many approaches for processing, analyzing, and visualizing WGS data have been developed and are being improved. Detection and analysis of copy number variations (CNVs) based on WGS data is one of them. Research directions related to cancer genomics, single-cell sequencing, and somatic mosaicism create huge amounts of data and demands for processing on the cloud that require further

improvements in CNV callers, moving to parallel processing, better compression, modular architecture, and new statistical methods.

CNVnator is a method for CNV analysis based on read depth (RD) of aligned reads. It has been determined to have high sensitivity (86–96%), low false-discovery rate (3–20%), and high genotyping accuracy (93–95%) for germline CNVs in a wide range of sizes from a few hundred base pairs to chromosome size events [1–5]. Since its development a decade ago, the tool has been widely used in different scientific areas by researchers around

Received: 16 May 2021; Revised: 21 August 2021; Accepted: 22 October 2021

© The Author(s) 2021. Published by Oxford University Press GigaScience. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

the world for detection of CNVs in a variety of species with different genome sizes: bacteria [6], fungi [7], plants [8–10], insects [11], fish [12], birds [13], mammals [14–17], and humans [2, 18–20]. It has been used to discover somatic variations in cancer and disease studies [21] and to find mosaic variants in human cells [22]. Although CNVnator was developed to detect germline CNVs, it is well suited to discover copy number alteration (CNAs) present in a relatively high (>50%) fraction of cells, such as somatic alteration found in cancers. It was not, however, designed for nor capable of aiding analysis of copy number-neutral changes.

Here we describe CNVpytor, a Python extension of CNVnator. CNVpytor inherits the reimplemented core engine of CNVnator and extends visualization, modularization, performance, and functionality. Along with RD data, it enables consideration of allele frequency of single-nucleotide polymorphisms (SNP) and small indels as an additional source of information for the analysis of CNV/CNA and copy number-neutral variations. Along with RD data, this information can be used for genotyping genomic regions and visualization.

Results

Analysis of RD signal

CNVpytor inherits the RD analysis approach developed in CNVnator [1]. Briefly, it consists of the following steps: reading alignment file and extracting RD signal, binning RD signal, correcting the signal for GC bias, segmenting the signal using the mean-shift technique, and calling CNVs (Fig. 1). RD signal can be parsed from BAM, SAM, or CRAM alignment files and is counted in 100-bp intervals, resulting in a small footprint of intermediate .pytor files in HDF5 format (Table 1). Because of using the pysam [23] library for parsing, this step (the most time-consuming one) is parallelized and can be conducted very efficiently, particularly in comparison with the older tool (Table 1). The binning step integrates RD over larger bins that are limited to multiples of initially stored 100-bp bins. Bin size can be adjusted by user depending on application. The recommended minimal bin size, which depends on sequencing coverage, is provided in Supplementary Table S1. Next, technical biases in the read depth signal that are correlated with GC content (so-called GC biases) are removed using a GC correction procedure. For the human reference genomes GRCh37 and GRCh38, per bin GC content is pre-calculated and supplied as a resource with the CNVpytor package. For other genomes, GC content can be calculated during runtime from a provided FASTA file or precalculated and added to the CNVpytor resource for future use. Once information about read coverage (and variants, see below) is extracted from an alignment file, the following analysis steps take place (i.e., read input and write output) with the same file. As a result, histograms for each processed bin size and information about CNV calls, including coordinates, different statistics, and P-values, are all stored in the same .pytor file and can be extracted into Excel (TSV file) or a VCF file.

Analysis of variant data

A novel feature of CNVpytor is the analysis of information from SNPs and small indels imported from a VCF file. An imbalance in the number of haplotypes can be measured using allele frequencies traditionally referred to as B-allele frequency (BAF) [24–27]. The main advantage of using BAF compared to RD is that BAF

values do not require normalization and are distributed around 0.5 by binomial distribution for heterozygous variants (HETs). Additionally, BAF is complementary to RD signal, as it changes for copy number-neutral events such as loss of heterozygosity. However, BAF dispersion can be measured incorrectly owing to systematic misalignment particularly in repeat regions, incomplete reference genome, or site-specific noise in sequencing data. To mitigate this issue, we filtered out HETs in the fraction of genome that is inaccessible to short read technologies, as defined by the strict mask from the 1000 Genomes Project [28]. Such filtering removes almost all HETs with outlier values of BAF, while values for the retained variants closely follow binomial distribution (Fig. 1c). To integrate BAF information within bins, CNVpytor calculates the likelihood function that describes an imbalance between haplotypes (see Methods). Currently, BAF information is used when genotyping a specific genomic region where, along with estimated copy number, the output contains the average BAF level and 2 independent P-values calculated from RD and BAF signal. Variant data can also be plotted in parallel with RD signal (Fig. 2). Same as for RD signal, binned information calculated from variants is stored in and can be extracted from the .pytor file.

Running CNVpytor

CNVpytor is to be run in a series of steps (Fig. 3). For enhanced flexibility, RD and BAF processing workflows proceed in parallel. In this way each workflow can be run at different times or even on different computers. For example, data parsing steps can be run on a cloud where data (i.e., alignment files) are accessible, resulting in <25 Mb .pytor files that then can be copied to a local computer/cluster where the remaining calculation steps will be performed. If necessary, a user can run additional calculations (e.g., conduct processing with different bin size) using the same .pytor file in the future, allowing for further flexibility in data analysis.

Routine processing steps can be followed by CNV visualization and analysis in the Viewer session, which can be interactive or hands-off. Implementation of interactive mode is inspired by a Linux shell with tab completion and with a help page similar to the man pages. In this mode, a user can instantly make various visualizations, preview and filter CNV calls, annotate calls, create joint calls across multiple samples, and genotype specified regions. The viewer does not save results into the .pytor file, and outputs are printed and plotted on the screen or exported to an output file(s). Hands-off mode executes user-written script(s) with CNVpytor commands. Such scripts can be used as part of the processing pipeline where, e.g., images of signals around called CNVs are generated and stored for possible future inspection. Through the viewer interface, it is possible to directly access Python and run code. This allows user to access some standard features of underlying libraries, e.g., matplotlib library can be used to customize plots.

CNVpytor can be used as the Python module. All functionalities, like reading and editing CNVpytor data files, and all calculation steps and visualizations can be performed by calling functions or classes. This way CNVpytor can be easily integrated in different platforms and computing environments; e.g., CNVpytor can be run from Jupyter Notebook on a local machine or in cloud services, e.g., Google Colab. CNVpytor is also integrated into OmniTier's "Compstor Novos" variant calling workflow [29].

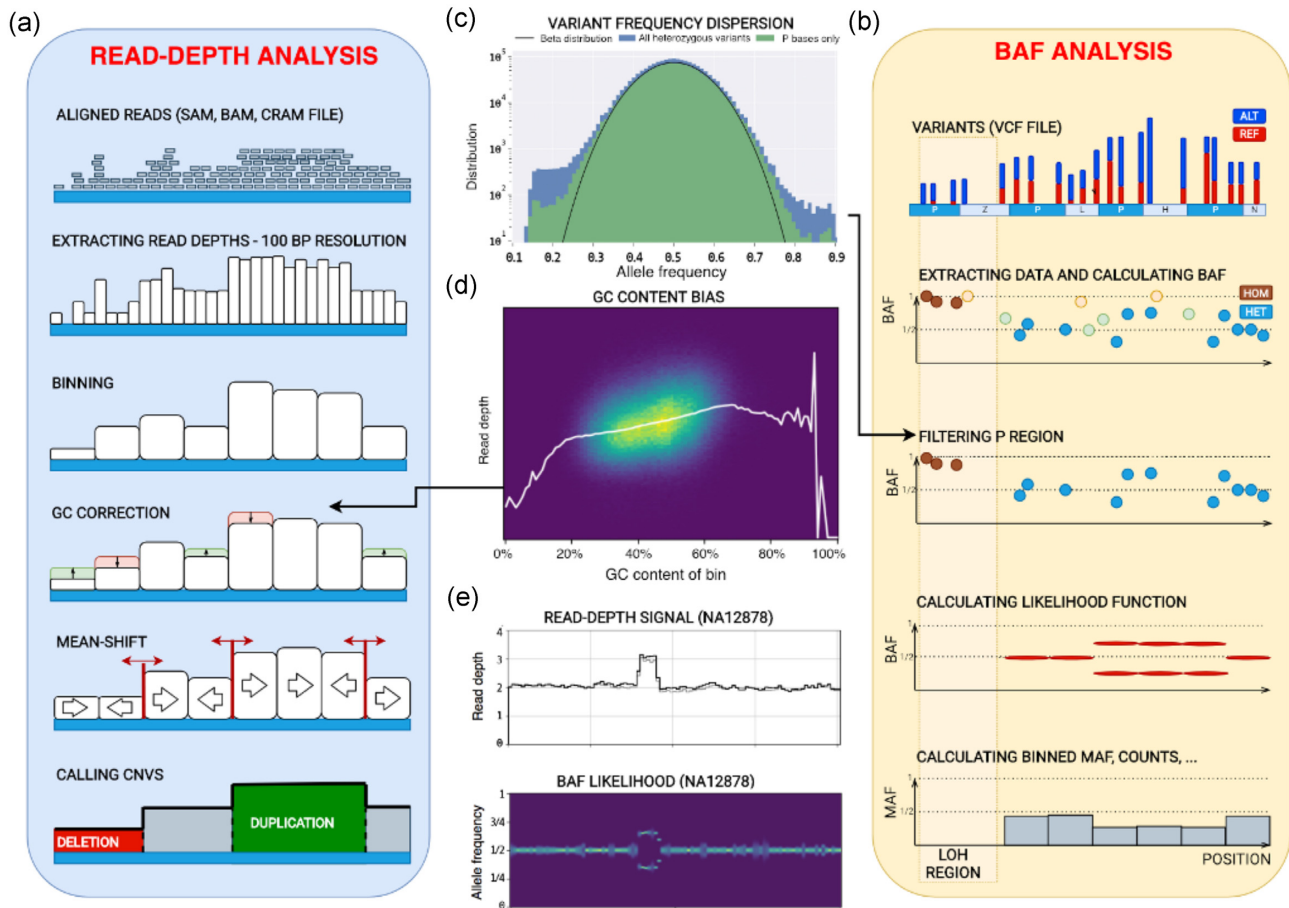


Figure 1: Schematics of core algorithm and data processing steps. (a) Read depth analysis steps include parsing alignment file, calculating and storing read depths in 100-bp intervals, binning using user-specified bin size, correcting RD for GC bias, segmenting by mean-shift, and calling CNVs. (b) B-allele frequency (BAF) analysis steps include reading variant file, storing the data about SNPs and small indels, filtering variants using strict mask, calculating BAF for heterozygous variants (HETs), and calculating likelihood function for bins. For CNVs, BAF signal splits away from value 0.5 expected for HETs. (c) Distribution of the variant allele frequency for all variants and variants within strict mask as defined by the 1000 Genomes Project. Black line shows fit by Gaussian distribution. (d) An example of RD depending on GC within bin. Statistics of RD signal within bins of the same percentage of GC content is used to correct for GC bias in the signal. White line represents average RD level for bins with given GC content. (e) An example of RD and BAF signals for a germline duplication in NA12878 sample (raw RD signal is in grey, GC-corrected RD signal is in black, brighter color of BAF likelihood corresponds to higher values of the likelihood).

Table 1: Efficiency of parsing alignment file on modern computers in relation to sequencing coverage and engaged number of CPU cores

| Sequencing coverage of the human genome | Parsing time | | | | | File size | | |
|---|--------------|----------|---------|----------|---------------|---------------|-------------|-----------------------|
| | CNVnator | CNVpytor | | | CNVnator (Gb) | CNVpytor (Mb) | | |
| | | 4 Cores | 8 Cores | 23 Cores | | RD parsing | BAF parsing | With 1 and 10 kb bins |
| 5× | 10 min | 5 min | 3 min | <2 min | 1 | 18 | 20 | 250 |
| 30× | 1 h | 28 min | 18 min | 10 min | 1.5 | 19 | 20 | 250 |
| 100× | 3.3 h | 1.5 h | 1 h | 33 min | 2 | 20 | 20 | 250 |

Data visualization and result curation

Visualization of multiple tracks/signals can be done interactively by mouse and by typing relevant commands, as well as by running scripts with CNVpytor commands provided as inputs to CNVpytor. CNVpytor has extended visualization capabilities with multiple novel features as compared to CNVnator. For each sample (i.e., input file) multiple data tracks such as RD signal, BAF of SNPs, and binned BAF likelihood can be displayed in an adjustable grid layout as specified by a user. Specifically, multi-

ple regions across multiple samples can be plotted in parallel, facilitating comparison across samples and different genomic loci (Fig. 2). To get a global view, a user can visualize an entire genome in a linear or circular fashion (Fig. 4, Supplementary Figs S1 and S2). Such a view can be useful in judging the quality of samples and in visually checking for aneuploidies.

Some additional features include GC-bias curve plot and 2D histogram (Fig. 1d), allele frequency distribution per region (Fig. 1c), and comparing RD distributions between 2 regions (Sup-

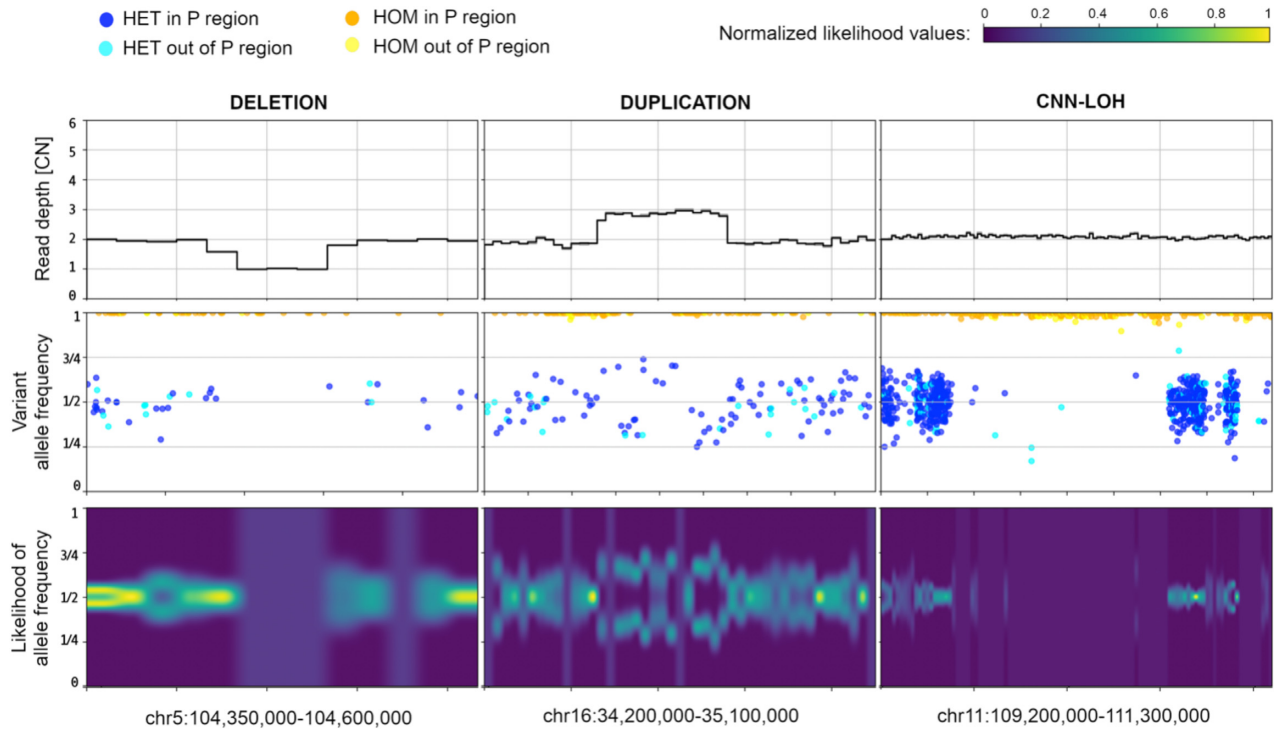


Figure 2: BAF signal corroborates and complements RD signal. Example of CNVpytor region plots produced for deletion (left), duplication (middle), and CNN-LOH (right) for NA12878 sample. Within the coordinates of heterozygous deletion, there is a 50% decrease in RD signal and a loss in heterozygosity in BAF signal (i.e., no heterozygous SNPs in the region). Duplication of 1 haplotype results in the increase of RD signal by 50% and in a split in VAF distribution of SNPs and a split in BAF likelihood function. In the CNN-LOH region, few reliable heterozygous SNPs are detected while RD signal does not change. Likelihood function values are normalized to the maximum value across the range.

plementary Fig. S3). Figure resolution, layout grid, colors, marker size, titles, and plotting style are adjustable by the user. Overall CNVpytor has more functionalities than other available software (Supplementary Table S2).

Integration with JBrowse

CNVpytor also has implemented functionality to export data into formats that can be embedded into JBrowse, a web-based genome browser used to visualize multiple related data tracks (Fig. S4). The export enables users to utilize JBrowse capabilities to visualize, compare, and cross-reference CNV calls with other data types (such as ChIP-seq, RNA-seq, ATAC-seq, etc.) and annotations across genome. Exported data provide 3 resolutions of RD and BAF tracks (1, 10, and 100 kb bins) while the appropriate resolution is chosen automatically by JBrowse depending on the size of the visualized genomic region. Multiple .pytor files can be exported at once.

Alternatively, a user can utilize a lightweight CNVpytor plugin for JBrowse. The plugin takes information about coverage from a relatively small (as compared to BAM) VCF file and on the fly performs the read depth and BAF estimation, segmentation, and calling. For read depth analysis, the plugin fetches the information from the DP field in the VCF file and uses it as a proxy for actual coverage. Since for large bin sizes such an estimate corresponds well to the actual value (Fig. S5), the plugin enables quick and easy review of large copy number changes in a genome. For BAF analysis, the plugin conducts analyses the same way as a stand-alone application. All temporary values are stored in the browser cache for fast and interactive visual-

ization of a genomic segment. As well as improving responsiveness by eliminating the network lag of a client-server application, this ensures that no information about a personal genome is transferred to external servers. Once the analyses are complete, the results are instantly visualized using JBrowse's native capabilities. Usage cases of the plugin are: 1) quick and visual cross-referencing of copy number profiles between multiple samples and in relation to other data types, and 2) a review of a personal genome(s) for large CNVs in a simple user-friendly environment.

Conclusion

Development of new, maintenance, and improvement of existing bioinformatics tools are driven by changing data types, demands for newer and user-verifiable analyses, necessity for processing larger datasets, and the evolving nature of computational infrastructures and platforms. CNVpytor brings the functionality of its predecessor CNVnator to a new level and significantly expands it. CNVpytor is faster and virtually effortless to install, requires minimum space for storage, enables analysis of BAF for call confirmation and genotyping, provides users with instant and extended visualization and convenient functionality for result curation (including merging over multiple samples), and is equipped for integration with other tools. The utilized method is suitable to segment RD signal in the case of mosaic or somatic cancer sample CNAs, and the alteration will be called if its cell frequency is >50% (Supplementary Fig. S6). A more accurate approach to discover cancer and mosaic CNAs is under development. The prototype of the somatic CNA caller

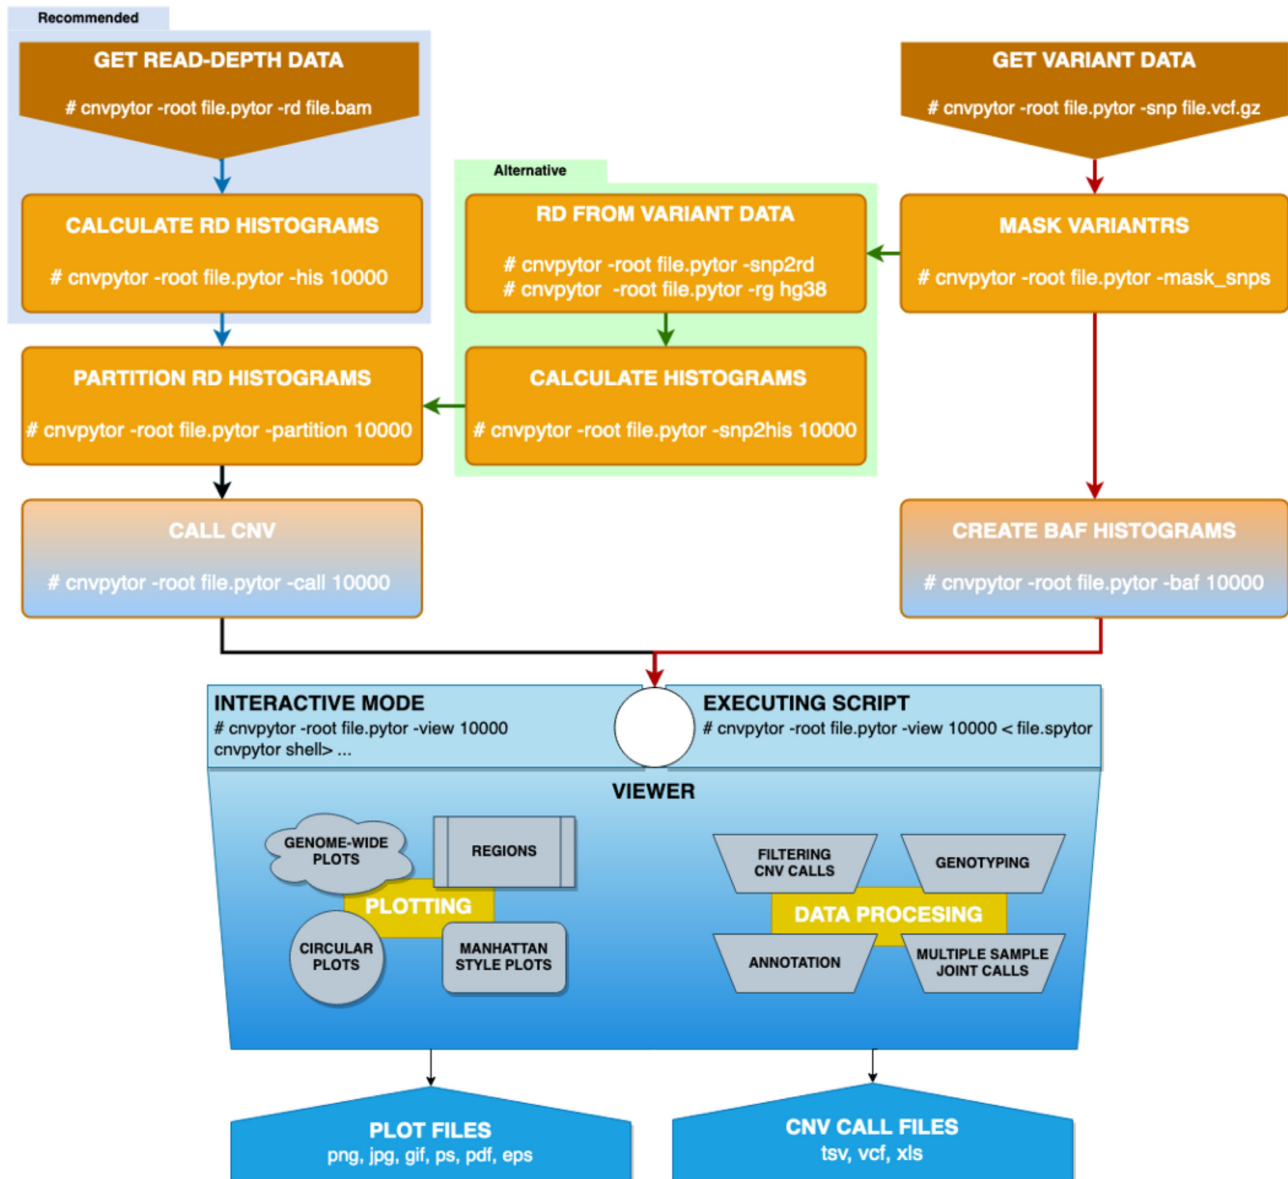


Figure 3: CNVpytor workflow and steps used in data processing. Left: Reading RD data from alignment file, creating histograms, segmentation, calling CNVs. Right: Reading SNP and indel data from VCF file, filtering variants using strict mask, calculating histograms and likelihood function. Middle: Alternatively, if alignment file is not available, RD signals can be calculated from variant data (green arrows). Visualization using both RD and/or BAF data can be done from an interactive command line interface or automatically by running script file.

is functional, available in the current version, and documented on the CNVpytor GitHub page. Provided good genome amplification and high sequencing coverage in a single cell, detection of CNAs in the cell becomes like detecting germline variation, i.e., every somatic variant will be present in 1 of 2 haplotypes. Advancement of single-cell amplification approaches [30] renders CNVpytor applicable to analysis of single-cell data. A lightweight plugin for JBrowse enables on-the-fly visualization and analysis convenient for wide categories of users. Overall, CNVpytor establishes a framework for discovering and analyzing copy number changes from WGS data either by an individual researcher or clinician or in a collaborative and shared environment.

Materials and Methods

RD analysis

Calculations for RD binning, mean-shift algorithm, partitioning, and calling CNVs are explained in detail in the CNVnator article [1]. The only difference in CNVpytor implementation is how information about GC content is obtained. For the 2 versions GRCh37 (hg19) and GRCh38 of the human reference genome, information about GC and AT content for each 100-bp bin is provided as resource data within the CNVpytor package. This way the user does not need to have reference genome FASTA for GC correction. In spite of slightly different implementation, called regions by CNVnator and CNVpytor overlap by >99%. Addition-

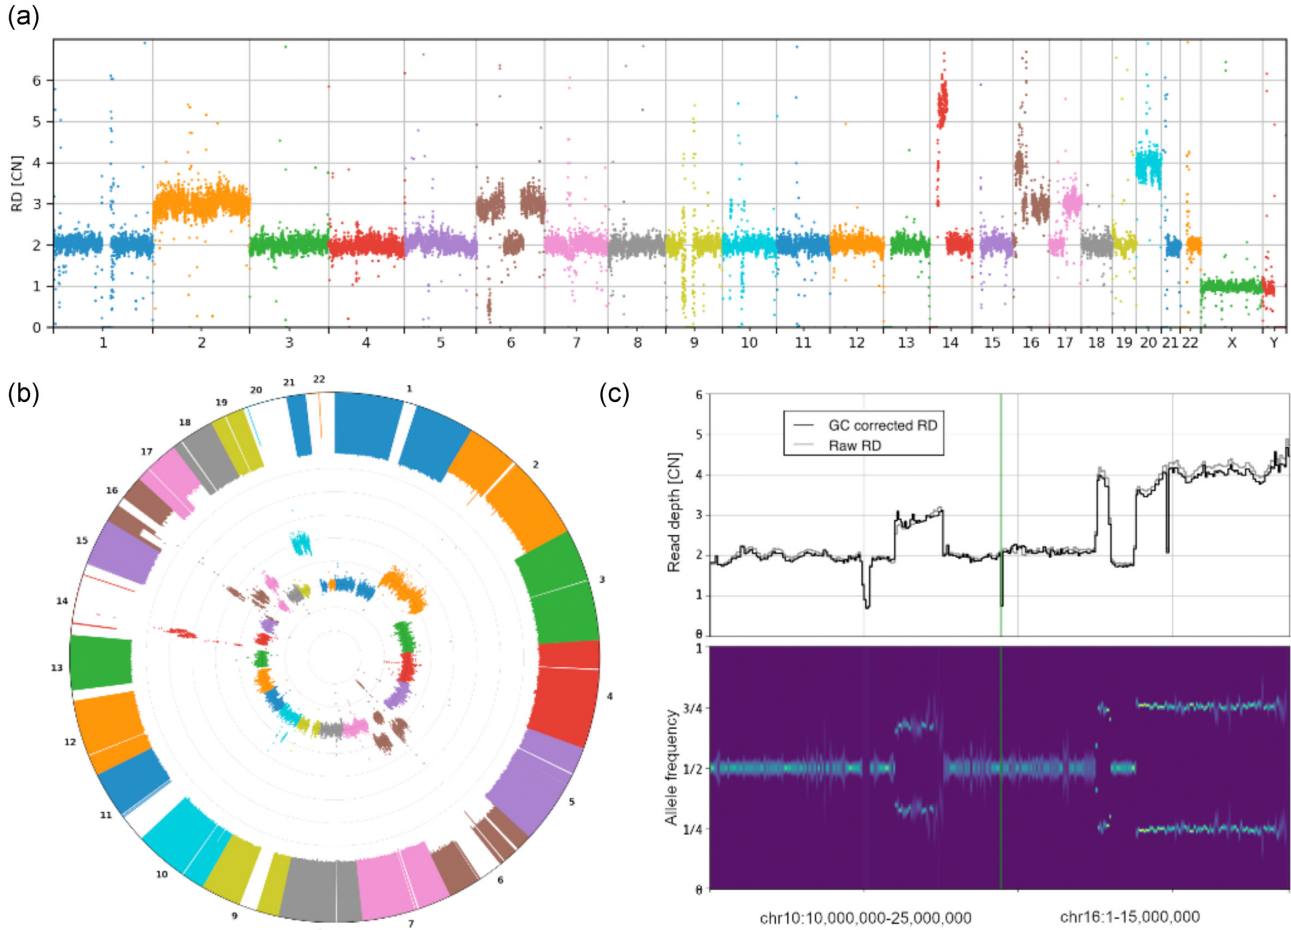


Figure 4: Novel features of visualization using data for HepG2 cell line. Genome-wide visualizations can be useful in some cases, including cancer studies and clinical applications such as screening for chromosomal abnormalities. They also can be useful for quick insight in sequencing or single-cell amplification quality. To demonstrate genome-wide plot types, HepG2 immortal cell line sample is used. (a) Manhattan and (b) circular style plot of RD signal. Large CNVs and chromosome copy number changes are apparent. One can also judge the dispersion of the RD signal. The inner circle in (b) shows RD signal while the outer one shows minor allele frequency (MAF). Regions where we can see loss of MAF signal with normal RD, e.g., chromosome 14 or 22, are chromosomal CNV-LOH. (c) Examples of smaller CNVs not apparent in the global view. Each CNV is ~15 Mb. The displayed regions with multiple CNVs (that are possibly complex events) were hardly visible on the circular plot.

ally, CNVpytor can calculate read depth from coverage of imported variants (Fig. 3). This is an approximate but rather precise solution to cases when the alignment file is not available (Supplementary Fig. S5).

Variant data

CNVpytor imports information about SNPs and single-letter indels from the variant (VCF) file. All other variants are ignored. For each variant the following data are stored in the CNVpytor file: chromosome, position, reference base, alternative base, reference count (ref_i), alternative count (alt_i), quality, and genotype (0/1 or 1/1).

In the normal case with 1 copy of each haplotype, heterozygous SNPs are visible in half of the reads coming from that haplotype. In other words, B-allele frequency distributes around 1/2. Contrarily, in the regions with constitutional duplication of 1 haplotype, heterozygous SNPs are expected with a frequency equal to 2/3 or 1/3 depending on whether they are located on a duplicated haplotype. This split from value 1/2 can be visible in a plot of BAFs vs position of variants as shown in the right panel of Fig. 1. Similarly, for homozygous deletion complete loss

of heterozygous SNPs is expected. In the case of somatic sub-clonal CNA (e.g., frequently observed in cancer genomes), the ratio between haplotypes can be an arbitrary number depending on cell frequencies with the CNA. Consequently, the split in BAF plot varies from 0 through 1. Measuring the level of this split can provide useful information about type of CNV. Moreover, copy number-neutral loss of heterozygosity (CNN-LOH) can be detected this way.

For each stored variant, we can calculate 2 frequencies defined in following way:

- B-allele frequency (BAF): $BAF = \frac{alt_i}{ref_i + alt_i}$
- Minor allele frequency (MAF): $MAF = \min\left(\frac{alt_i}{ref_i + alt_i}, \frac{ref_i}{ref_i + alt_i}\right) = \min(BAF, 1 - BAF)$

One of the characteristics of next-generation sequencing is that some bases are not accessible for variant discovery using short reads, owing to the repetitive nature of the human genome. In the 1000 Genomes Project, a genome mask is created to tabulate bases for variant discovery. There are ~74% of bases marked passed (P), which corresponds to ~77% of non-N bases (1000 Genomes Project Consortium [28]). We use that mask to fil-

ter out variants called in non-P regions. This way we eliminate ~22% of variants, but there is a benefit because with fewer false-positive heterozygous SNPs the statistics is improved (Fig. 1c) and this improves the quality of further calculations. The same way as for GC content, information about strict mask P regions for 2 versions of the human reference genome is stored in resource files that are part of the CNVpytor package.

Calculating BAF likelihood function

The ratio between reads coming from one or another haplotype is distributed following a binomial distribution. If counts are known, then one can calculate the likelihood function for that ratio:

$$L(p_i | alt_i, ref_i) = B(x; alt_i + 1, ref_i + 1) = \frac{1}{B(alt_i, ref_i)} p_i^{alt_i} (1 - p_i)^{ref_i},$$

where p_i is allele frequency for variant i and $B(alt_i, ref_i)$ is the normalization constant.

By multiplying likelihood functions of individual variants in a bin, one can obtain likelihood for each bin. This is true only if the real value of that ratio does not change within that bin. However, we are using non-phased variant counts, which means that there is a 50% chance that variant is coming from one or another haplotype. The frequency of a fixed haplotype is sometimes described by either a BAF or a $1 - BAF$ distribution.

In that case, we have to use the symmetrized beta distribution for likelihood:

$$\begin{aligned} L(p_i | alt_i, ref_i) &= \frac{B(x; alt_i + 1, ref_i + 1) + B(x; ref_i + 1, alt_i + 1)}{2} \\ &= \frac{p_i^{alt_i} (1 - p_i)^{ref_i} + p_i^{ref_i} (1 - p_i)^{alt_i}}{2B(alt_i, ref_i)} \end{aligned}$$

The likelihood function for each bin is calculated as a product of individual likelihood functions of variants within that bin:

$$\begin{aligned} L(p_b) &\sim \prod_{i \in \text{bin}(b)} L(p_i | alt_i, ref_i) \sim \prod_{i \in \text{bin}(b)} \\ &\left(p_i^{alt_i} (1 - p_i)^{ref_i} + p_i^{ref_i} (1 - p_i)^{alt_i} \right), \end{aligned}$$

where p_b is the allele frequency for bin b . To calculate likelihood functions, we use discretization. Interval $[0, 0.5]$ is discretized using some resolution (default is 101 point) and for each the point function is calculated by multiplying values of symmetrized beta distribution for each variant. The position of maximum likelihood represents the most probable BAF value in a particular bin. Along with the likelihood function average values of variant BAF and MAF are calculated per bin and stored in the CNVpytor file, together with counts of homozygous and heterozygous variants.

Filtering CNV calls

For each CNV call the following values are calculated: (i) event type: “deletion” or “duplication”; (ii) coordinates in the reference genome; (iii) CNV size; (iv) RD normalized to 1; (v) e-val1: P-value calculated using i-test statistics between RD difference in the region and global (i.e., across whole genome) mean; (vi) e-val2: P-value from the probability of RD values within the region to be in the tails of a Gaussian distribution of binned RD; (vii) e-val3: same as e-val1 but without first and last bin; (viii) e-val4: same as e-val2 but without first and last bin; (ix) q0: fraction of reads

mapped with zero quality within call region; (x) pN: fraction of N bases (i.e., unassembled reference genome) within call region; (xi) dG: distance to nearest gap in reference genome.

There are 5 parameters in viewer mode used for filtering calls: CNV size, e-val1, q0, pN, and dG. Those parameters will define which calls CNVpytor will plot or print out. When calls are printed or exported, CNVpytor optionally can generate graphical file(s) with a plot of the CNV call region containing user-specified tracks.

Annotating CNV calls

To annotate called regions, we use Ensembl REST API (overlap/region resource). It is an optional step that requires web connection and is executed when calls are previewed by the user or exported to an output file. The annotation is added in an additional column in the output and contains a string with gene names, Ensembl gene IDs, and information about the position of genes relative to CNV (i.e., inside, covering, or intersecting left/right breakpoints of the CNV region).

Genotyping

The copy number of a provided genomic region is calculated as a mean RD within the region divided by mean autosomal RD scaled by 2. To achieve better precision, first and last bin content are weighted by the fraction of overlap with the provided region. Optionally CNVpytor can provide additional values: (i) e-value from the probability of RD values within the region to be in the tails of a Gaussian distribution of binned RD (analogous to e-val2); (ii) q0: fraction of reads mapped with q0 quality within call region; (iii) pN: fraction of reference genome gaps (Ns) within call region; (iv) BAF level estimated using maximum likelihood method; (v) number of homozygous variants within the region; (vi) number of heterozygous variants within the region; (vii) P-value based on BAF signal.

Merging calls over multiple regions

To make a joint call set for multiple samples, CNVpytor proceeds in the following way:

1. Filter calls using user-defined ranges for size, p-val, q0, pN, and dG;
2. Sort all calls for all samples by start coordinate;
3. Select first call in that list that is not already processed and select calls from other samples with reciprocal overlap >50%;
4. For selected calls, calculate genotypes within the region of intersection and, optionally, annotate with overlapping genes.

If specified, for each joint CNV call CNVpytor will create a graphical file with a plot of the call region containing user-specified tracks.

Data format and compression

For data storage and compression, we used HDF5 file format and h5py Python library. Additional compression is obtained by storing RD signal using 100-bp bins. The same bin size is used for storing reference genome AT, GC, and N content. Data organization within the .pytor file is implemented in an IO module, which can be used to open and read different datasets from an

external application. The Python library xlwt is used to generate spreadsheet files compatible with Microsoft Excel.

Visualizations

The Matplotlib [31] Python library is used for creating and storing visualizations. Different plotting styles are available within matplotlib. Derived installed libraries can be used, as well as a variety of file formats for storing graphical data.

Module dependences

CNVpytor depends on several widely used Python packages: requests 2.0 or higher, gnureadline, pathlib 1.0 or higher, pysam 0.15 or higher, numpy 1.16 or higher, scipy 1.1 or higher, matplotlib 2.2 or higher, h5py 2.9 or higher, xlwt 1.3 or higher. All dependences are available through pip installer, which makes installation of CNVpytor straightforward.

Data Availability

An archival copy of the code and links to data used to create figures are available via the GigaScience database, GigaDB [32].

Availability of Supporting Source Code and Requirements

Project name: CNVpytor

Project home page: <https://github.com/abyzovlab/CNVpytor>

Operating systems: Platform independent

Programming language: Python

Other requirements: requests ≥ 2.0 , gnureadline, pathlib ≥ 1.0 , pysam ≥ 0.15 , numpy ≥ 1.16 , scipy ≥ 1.1 , matplotlib ≥ 2.2 , h5py ≥ 2.9 , xlswriter ≥ 1.3 , pathlib ≥ 1.0

License: MIT License

RRID:SCR_021627

bio.tools ID: cnvpytor

Additional Files

Supplementary Figure 1: Genome-wide plot for K562 cell line: normalized read depth (top), B-allele frequency of individual SNPs (middle), and BAF likelihood function (bottom). Bin size is 100 kb.

Supplementary Figure 2: Circular plot for K562 cell line. Inner circle corresponds to read depth; outer, to binned MAF. Bin size is 100 kb.

Supplementary Figure 3: Comparison of read depth statistics between 2 regions.

Supplementary Figure 4: JBrowse export example. (a, b) CNVpytor-produced read depth and binned BAF data for a glioblastoma cancer sample for chromosome 1 deletion are visible. (c) JBrowse view of the same data. Same color-coding schema is followed here.

Supplementary Figure 5: Comparison between read depth signal parsed from alignment file and variant file for 3 samples: RD Manhattan plot comparison for NA12878 sample (a), K562 sample (b), and HepG2 (c); distribution of differences in copy number within bins for same samples (d–f). Bin size is 10 kb.

Supplementary Figure 6: CNVpytor application on polyp subclonal CNA [33]. Raw RD signal is in grey, GC-corrected RD signal is in black, segmentation is in red, and CNA calls in green.

Supplementary Table 1: Recommended minimal bin size for given coverage to ensure relative deviation of RD signal $<10\%$

for 150-bp reads. For 100-bp reads one can use 33% smaller bin size.

Supplementary Table 2: Comparison between CNVpytor features with other similar tools.

Abbreviations

API: application programming interface; BAF: B-allele frequency; bp: base pairs; CNA: copy number alteration; CNV: copy number variation; CPU: central processing unit; Gb: gigabase pairs; GUI: graphical user interface; HETs: heterozygous variants; kb: kilobase pairs; Mb: megabase pairs; RD: read depth; REST: representational state transfer; SNP: single-nucleotide polymorphism; WGS: whole-genome sequencing.

Competing Interests

The authors declare that they have no competing interests.

Funding

This study is supported by National Cancer Institute grant U24CA220242 and funds from the Center for Individualized Medicine at Mayo Clinic.

Authors' Contributions

A.A. conceived and supervised this study. M.S. designed and developed CNVpytor software. A.P. and M.S. tested the software. A.P. and C.D. developed the plugin for JBrowse. I.H. co-supervised the development of the plugin for JBrowse. M.S., A.P., and A.A. wrote the manuscript. All authors read and approved the final manuscript.

Acknowledgements

We thank Abyzov lab members for useful discussions and suggestions. We are grateful to users on the GitHub for useful suggestions and bug reports.

References

1. Abyzov A, Urban AE, Snyder M, et al. CNVnator: an approach to discover, genotype, and characterize typical and atypical CNVs from family and population genome sequencing. *Genome Res* 2011;21(6):974–84.
2. Mills RE, Walter K, Stewart C, et al. Mapping copy number variation by population-scale genome sequencing. *Nature* 2011;470(7332):59–65.
3. Duan J, Zhang JG, Deng HW, et al. Comparative studies of copy number variation detection methods for next-generation sequencing technologies. *PLoS One* 2013;8(3):e59128.
4. Legault MA, Girard S, Lemieux-Perreault LP, et al. Comparison of sequencing based CNV discovery methods using monozygotic twin quartets. *PLoS One* 2015;10(3):e0122287.
5. Trost B, Walker S, Wang Z, et al. A comprehensive workflow for read depth-based identification of copy-number variation from whole-genome sequence data. *Am J Hum Genet* 2018;102(1):142–55.
6. Coll F, Preston M, Guerra-Assunção JA, et al. PolyTB: a genomic variation map for *Mycobacterium tuberculosis*. *Tuberculosis* 2014;94(3):346–54.

7. Cabañes FJ, Sanseverino W, Castellá G, et al. Rapid genome resequencing of an atoxigenic strain of *Aspergillus carbonarius*. *Sci Rep* 2015;5(1):9086.
8. Fuentes RR, Chebotarov D, Duitama J, et al. Structural variants in 3000 rice genomes. *Genome Res* 2019;29(5):870–80.
9. Gordon SP, Priest H, Des Marais DL, et al. Genome diversity in *Brachypodium distachyon*: deep sequencing of highly diverse inbred lines. *Plant J* 2014;79(3):361–74.
10. Wallace JG, Bradbury PJ, Zhang N, et al. Association mapping across numerous traits reveals patterns of functional variation in maize. *PLoS Genet* 2014;10(12):e1004845.
11. Choi JY, Bubnell JE, Aquadro CF. Population genomics of infectious and integrated *Wolbachia pipiensis* genomes in *Drosophila ananassae*. *Genome Biol Evol* 2015;7(8):2362–82.
12. Chain FJ, Feulner PG, Panchal M, et al. Extensive copy-number variation of young genes across stickleback populations. *PLoS Genet* 2014;10(12):e1004830.
13. Yi G, Qu L, Liu J, et al. Genome-wide patterns of copy number variation in the diversified chicken genomes using next-generation sequencing. *BMC Genomics* 2014;15(1):962.
14. Hermesen R, de Ligt J, Spee W, et al. Genomic landscape of rat strain and substrain variation. *BMC Genomics* 2015;16(1):357.
15. Wang H, Wang C, Yang K, et al. Genome wide distributions and functional characterization of copy number variations between Chinese and Western pigs. *PLoS One* 2015;10(7):e0131522.
16. Gokcumen O, Tischler V, Tica J, et al. Primate genome architecture influences structural variation mechanisms and functional consequences. *Proc Natl Acad Sci U S A* 2013;110(39):15764–9.
17. Pezer Ž, Harr B, Teschke M, et al. Divergence patterns of genic copy number variation in natural populations of the house mouse (*Mus musculus domesticus*) reveal three conserved genes with major population-specific expansions. *Genome Res* 2015;25(8):1114–24.
18. Abel HJ, Larson DE, Regier AA, et al. Mapping and characterization of structural variation in 17,795 human genomes. *Nature* 2020;583(7814):83–89.
19. Sudmant PH, Rausch T, Gardner EJ, et al. An integrated map of structural variation in 2,504 human genomes. *Nature* 2015;526(7571):75–81.
20. Nagasaki M, Yasuda J, Katsuoka F, et al. Rare variant discovery by deep whole-genome sequencing of 1,070 Japanese individuals. *Nat Commun* 2015;6(1):8018.
21. Han L, Zhao X, Benton ML, et al. Functional annotation of rare structural variation in the human brain. *Nat Commun* 2020;11(1):2990.
22. Guo H, Duyzend MH, Coe BP, et al. Genome sequencing identifies multiple deleterious variants in autism patients with more severe phenotypes. *Genet Med* 2019;21(7):1611–20.
23. Gilman P, Janzou S, Guittet D, et al. PySAM (Python Wrapper for System Advisor Model “SAM”). 2019. Golden, CO, USA: National Renewable Energy Lab.
24. Peiffer DA, Le JM, Steemers FJ, et al. High-resolution genomic profiling of chromosomal aberrations using Infinium whole-genome genotyping. *Genome Res* 2006;16(9):1136–48.
25. Loh PR, Genovese G, Handsaker RE, et al. Insights into clonal haematopoiesis from 8,342 mosaic chromosomal alterations. *Nature* 2018;559(7714):350–5.
26. Boeva V, Popova T, Bleakley K, et al. Control-FREEC: a tool for assessing copy number and allelic content using next-generation sequencing data. *Bioinformatics* 2012;28(3):423–5.
27. Zhu M, Need AC, Han Y, et al. Using ERDS to infer copy-number variants in high-coverage genomes. *Am J Hum Genet* 2012;91(3):408–21.
28. Auton A, Brooks LD, Durbin RM, et al. A global reference for human genetic variation. *Nature* 2015;526(7571):68–74.
29. OmniTier. CompStor Novos. 2020; <https://www.omnitier.com/compstor-novos>.
30. Gonzalez-Pena V, Natarajan S, Xia Y, et al. Accurate genomic variant detection in single cells with primary template-directed amplification. *Proc Natl Acad Sci U S A* 2021;118(24):e2024176118.
31. Hunter JD. Matplotlib: A 2D graphics environment. *Comput Sci Eng* 2007;9(3):90–5.
32. Suvakov M, Panda A, Diesh C, et al. Supporting data for “CNVpytor: a tool for copy number variation detection and analysis from read depth and allele imbalance in whole-genome sequencing.” *GigaScience Database* 2021; <http://doi.org/10.5524/100938>.
33. Kim M, Druliner BR, Vasmatzis N, et al. Inferring modes of evolution from colorectal cancer with residual polyp of origin. *Oncotarget* 2017;9(6):6780–92.