

cerebra: A tool for fast and accurate summarizing of variant calling format (VCF) files

Lincoln Harris¹, Rohan Vanheusden¹, Olga Botvinnik¹, and Spyros Darmanis¹

¹ Chan Zuckerberg Biohub, San Francisco, CA, United States

DOI: [10.21105/joss.02432](https://doi.org/10.21105/joss.02432)

Software

- [Review](#) ↗
- [Repository](#) ↗
- [Archive](#) ↗

Editor: [Lorena Pantano](#) ↗

Reviewers:

- [@betteridiot](#)
- [@afrubin](#)

Submitted: 25 June 2020

Published: 01 October 2020

License

Authors of papers retain copyright and release the work under a Creative Commons Attribution 4.0 International License ([CC BY 4.0](#)).

Motivation

A single “typo” in the genome can have profound consequences on an organism’s biology. Identifying the protein changes that accompany these genomic typos is a fundamental challenge in bioinformatics. Tools exist for identifying genomic variants and predicting their associated polypeptide-level changes, however, wrangling genomic variant calls and protein variant predictions across thousands of samples remains a substantial challenge. *cerebra* addresses this need by offering a fast and accurate framework for summarizing genomic variant calls and protein variant predictions across many samples.

To find variants in the genome, researchers often begin with a [DNA-sequencing](#) (DNA-seq) or [RNA-sequencing](#) (RNA-seq) experiment on their samples of interest. Sequencing is followed by alignment of reads to the reference genome with tools like [STAR](#) or [BWA](#), followed by variant calling with tools like [GATK HaplotypeCaller](#) or [freebayes](#) (Dobin et al., 2013; Garrison & Marth, 2012; Li & Durbin, 2009; Poplin et al., 2018). Variant callers produce tab delimited text files in the [variant calling format](#) (VCF) for each processed sample. VCF files encode: i) the genomic position, ii) reference vs. observed DNA sequence, and iii) quality associated with each observed variant. Shown below are the first 4 lines of a sample VCF file. Note that only a single record is displayed, and that the record line has been artificially wrapped.

```
##fileformat=VCFv4.2
##source=HaplotypeCaller
#CHROM POS ID REF ALT QUAL FILTER INFO FORMAT
chr1 631391 . C T 72.28 . AC=2;AF=1.00;AN=2;DP=2;
ExcessHet=3.0103;FS=0.000;MLEAC=1;MLEAF=0.500;MQ=NaN;
QD=25.36;SOR=2.303 GT:AD:DP:GQ:PL 1/1:0,2:2:6:84,6,0
```

Current methods for variant calling are incredibly powerful and robust, however, a single sequencing run can generate as many as 10^8 unique VCF records. Only a small portion of these VCF records are likely to be relevant to the researcher. In addition, variant callers report only the genomic location of the variant, and not the effect the variant has on the translated protein sequence. To address the unmet need for high-throughput VCF summary tools, we introduce *cerebra*, a python package that provides fast and accurate protein variant summarizing of VCF files.

Functionality

cerebra comprises three modules: i) `germline-filter` is intended for working with cancer data and removes variants that are common between tumor and normal samples, ii) `cou`

nt-variants reports total number of variants in each sample, and iii) find-peptide-variants reports the likely protein variants in each sample. Here we use *variant* to refer to single nucleotide polymorphisms (SNPs) and short insertions/deletions. *cerebra* is not capable of reporting larger structural variants such as copy number variations and chromosomal rearrangements.

A data structure crucial to *cerebra* is the *genome interval tree*, which matches RNA transcripts and polypeptides to each feature in the genome (Figure 1). *Interval trees* are self-balancing binary search trees that store numeric intervals and can quickly retrieve every such interval that overlaps a given query interval (*see also*). Given n nodes, interval trees have theoretical average-case $O(\log n)$ and worst-case $O(n)$ time complexity for search operations, making them tractable for genome-scale operations (Cormen, Leiserson, Rivest, & Stein, 2009, *see also*). Tree construction proceeds at $O(n \log n)$ time complexity, making construction rather than search the bottleneck for most VCF sets (Alekseyenko & Lee, 2007). The *genome interval tree* is constructed with a reference genome sequence (FASTA format, often with a .fa extension), and a genome annotation (gene transfer format, GTF, .gtf extension). We rely on the *ncls* python library for fast interval tree construction and lookup operations (Alekseyenko & Lee, 2007).

In order to analyze multiple VCF records at once, we use *multiprocessing* with the Python *pathos* library module (McKerns, Strand, Sullivan, Fang, & Aivazis, 2011). We extract relevant information – including genomic interval, observed base, and read coverage – from each variant record. In the *germline-filter* module variants are compared to one another and filtered out if found to be identical. In *count-variants* variants are simply matched to whichever gene they came from. In *find-peptide-variants* variants are queried against our *genome interval tree* – if a matching interval is found we convert the DNA-level variant to a protein variant. Finally, protein variants across all VCF files are reported in tabular format.

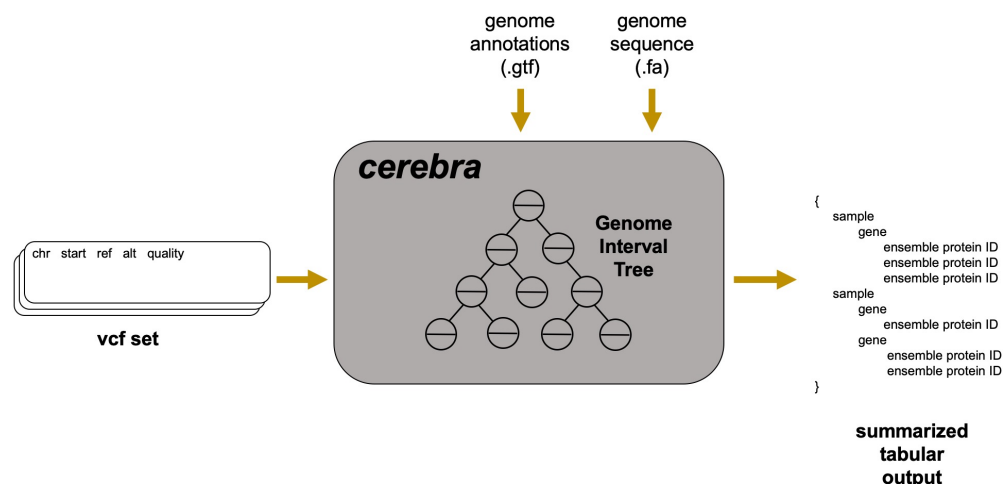


Figure 1: Workflow describing the *find-peptide-variants* module. We construct a genome interval tree from a genome annotation (.gtf) and a reference genome sequence (.fa), then process VCF files in parallel to create a single tabular output file (CSV or JSON).

germline-filter

Variant calling is often applied to the study of cancer. If the research project is centered around a “tumor vs. normal” question, then *germline-filter* is the proper starting point. This

module removes germline variants that are common between tumor and normal samples, and thus excludes variants unlikely to be pathogenic for the cancer under study. The user provides a very simple metadata file (see [USAGE.md](#)) that indicates which tumor samples correspond to which normal samples. Using the [vcfpy](#) library we quickly identify shared variants across tumor/normal matched VCF files, then write new VCFs that contain only the unique variants (Holtgrewe & Beule, 2016). These steps are performed by a [subprocess pool](#) so that we can process multiple discrete chunks of input at the same time.

The output of `germline-filter` is a set of trimmed-down VCF files, which will be used for the next two steps. If you do not have access to “normal” samples then proceed directly to `count-variants` or `find-peptide-variants`.

`count-variants`

The `count-variants` module reports the raw variant counts for every gene across every sample. We first create a *genome interval tree* from the reference GTF, then read in a VCF file and convert it to a [vcfpy](#) object, then processes VCF records in [parallel](#). Each variant is matched to its corresponding gene, and gene-wise counts are stored in [shared memory](#).

If working with cancer samples, the user has the option to limit the reported variants to those also found in Wellcome Sanger Institute’s [COSMIC](#) database (Tate et al., 2018). While certainly not exhaustive, this database contains an extensive list of known human variants. This option is designed to limit the search space to known and potentially actionable targets.

`count-variants` produces two output files, one containing raw variant counts and one containing COSMIC filtered variant counts for every gene in the genome.

`find-peptide-variants`

The `find-peptide-variants` module reports the protein variants potentially associated with each genomic variant. First we load the reference GTF, then construct an index (`.fai`) of the genome fasta file with [pyfaidx](#) to enable fast random memory access (Shirley, Zhaorong, Pedersen, & Wheelan, 2015). We then create a *genome interval tree* that will be used to quickly match genomic coordinates from VCF records to protein variants. The user again has the option to limit the search space to variants found in the COSMIC database. VCF records are read in simultaneously; individual records are converted to *GenomePosition* objects to keep track of their genomic intervals and observed DNA bases. *GenomePositions* are then queried against the *genome interval tree*. If an overlapping interval is found, we retrieve the protein variant from this node of the *genome interval tree*. Protein variants are converted to [ENSEMBL](#) protein IDs, in accordance with the [HGVS](#) sequence variant nomenclature (Dunnen et al., 2016; Yates et al., 2019). The output is a hierarchically ordered text file (CSV or JSON) that reports the the ENSEMBL protein ID and the gene associated with each variant, for each experimental sample.

Variant callers are known to produce a great deal of false positives, especially when applied to single-cell RNA-seq data (Enge et al., 2017). To address this concern, we have included the coverage option. If indicated this option will report counts for both variant and wildtype reads at all variant loci. We reasoned that variants with a high degree of read support are less likely to be false positives. This option is designed to give the user more confidence in individual variant calls.

We should emphasize that `find-peptide-variants` does not *definitively* report protein variants but rather the *likely* set of protein variants. Definitively reporting protein variants from RNA-seq requires knowledge of alternate splicing – this represents an open problem in the field (Huang & Sanguinetti, 2017). For example, if a read picks up a variant in exon 2 of a given gene, we can report each of the potential isoforms of that gene that contain exon

2, but we **cannot** infer which of those particular isoforms are actually present in our sample (see Figure 2). For the example shown in Figure 2 we would translate and report $t1$ and $t3$ as both of these contain exon 2. It is possible the sample does not actually express both of these isoforms, however, determining the isoform landscape of a sample from RNA-seq is outside the scope of this project.

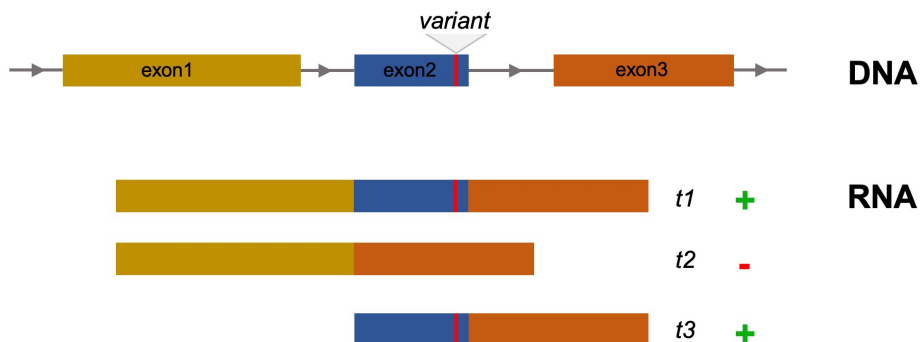


Figure 2: For a given mutational event, *cerebra* reports ALL potentially affected isoforms.

To assess performance of `find-peptide-variants` we obtained VCFs from a single-cell RNA-seq study conducted on lung adenocarcinoma patient samples (Maynard et al., 2020). These VCFs were produced with STAR (alignment) and GATK HaplotypeCaller (variant calling), and are on the order of megabytes, typical of a single-cell RNA-seq experiment. *cerebra* was run on standard hardware (MacBook Pro, 2.5GHz quad-core processor, 16 GB RAM). As show in Figure 3 *cerebra* processed a set of 100 VCF files in approximately 34 minutes.




Figure 3: cerebra processes 100 VCF files (~400 Mb in total) in ~34 minutes.

The first 10 or so minutes of cerebra `find-peptide-variants` do not involve any VCF processing, instead, this time is attributed to the *genome interval tree* construction phase. After the tree is built, files are processed in a near-linear manner. Also of note is that cerebra's search operations take advantage of multiprocessing. cerebra should scale better to high-memory machines than single-threaded tools, though it has been designed to run on standard hardware.

Conclusions

RNA/DNA sequencing paired with fast and accurate summarizing of variants is often crucial to understanding the biology of an experimental system. We present a tool that can be used to quickly summarize the variant calls contained within a large set of VCF files. As sequencing costs continue to drop, large-scale variant calling will become more accessible, and summary tools like cerebra will become essential for drawing meaningful conclusions in a reasonable time frame. Our tool offers the advantages of parallel processing and a single, easy-to-interpret output file (CSV or JSON).

cerebra is already enabling research, see (Maynard et al., 2020), a study that examines the tumor microenvironment of late-stage drug-resistant carcinomas with single-cell RNA-sequencing. Understanding the mutational landscape of individual tumors was essential to this study, and given the sheer volume of VCF records, would not have been possible without cerebra. We hope that cerebra can provide an easy-to-use framework for future studies in the same vein.

Acknowledgments

Funding for this work was provided by the [Chan Zuckerberg Biohub](#). The authors would like to thank Ashley Maynard, Angela Pisco and Daniel Le for helpful discussions and support.

Correspondence

Please contact ljharris018@gmail.com

Code

cerebra is written in Python 3. Code and detailed installation instructions can be found at <https://github.com/czbiohub/cerebra>. In addition, cerebra can be found on [PyPi](#) and [Dockerhub](#).

References

- Alekseyenko, A., & Lee, C. (2007). Nested Containment List (NCList): a new algorithm for accelerating interval query of genome alignment and interval databases. *Bioinformatics*. doi:[10.1093/bioinformatics/btl647](https://doi.org/10.1093/bioinformatics/btl647)
- Cormen, T., Leiserson, C., Rivest, R., & Stein, C. (2009). *Introduction to Algorithms* (3rd ed.). MIT Press. ISBN: [978-0-262-03384-8](#)
- Dobin, A., Davis, C., Schlesinger, F., Drenkow, J., Zaleski, C., Jha, S., & others. (2013). STAR: ultrafast universal RNA-seq aligner. *Bioinformatics*. doi:[10.1093/bioinformatics/bts635](https://doi.org/10.1093/bioinformatics/bts635)
- Dunnen, J. den, Dalgleish, R., Maglott, D., Hart, R., Greenblatt, M., McGowan-Jordan, J., & others. (2016). HGVS Recommendations for the Description of Sequence Variants: 2016 Update. *Human Mutation*. doi:[10.1002/humu.22981](https://doi.org/10.1002/humu.22981)
- Engel, M., Efsun, A., Mignardi, M., Beausang, J., Bottino, R., & Quake, S. (2017). Single-Cell Analysis of Human Pancreas Reveals Transcriptional Signatures of Aging and Somatic Mutation Patterns. *Cell*. doi:[10.1016/j.cell.2017.09.004](https://doi.org/10.1016/j.cell.2017.09.004)
- Garrison, E., & Marth, G. (2012). Haplotype-based variant detection from short-read sequencing. *arXiv*. Retrieved from <https://arxiv.org/abs/1207.3907>
- Holtgrewe, M., & Beule, D. (2016). VCFPy: a Python 3 library with good support for both reading and writing VCF. *Journal of Open Source Software*. doi:[10.21105/joss.00085](https://doi.org/10.21105/joss.00085)
- Huang, Y., & Sanguinetti, G. (2017). BRIE: transcriptome-wide splicing quantification in single cells. *Genome Biology*. doi:[10.1186/s13059-017-1248-5](https://doi.org/10.1186/s13059-017-1248-5)
- Li, H., & Durbin, R. (2009). Fast and accurate short read alignment with Burrows-Wheeler Transform. *Bioinformatics*. doi:[10.1093/bioinformatics/btp324](https://doi.org/10.1093/bioinformatics/btp324)
- Maynard, A., McCoach, C., Rotow, J., Harris, L., Haderk, F., Kerr, L., & others. (2020). Therapy-Induced Evolution of Human Lung Cancer Revealed by Single-Cell RNA Sequencing. *Cell*. doi:[10.1016/j.cell.2020.07.017](https://doi.org/10.1016/j.cell.2020.07.017)
- McKerns, M., Strand, L., Sullivan, A., Fang, A., & Aivazis, M. (2011). Building a framework for predictive science. *Proceedings of the 10th Python in Science Conference*. Retrieved from <http://arxiv.org/pdf/1202.1056>
- Poplin, R., Ruano-Rubio, V., DePristo, M., Fennell, T., Carneiro, M., Van der Auwera, G., & others. (2018). Scaling accurate genetic variant discovery to tens of thousands of samples. *bioRxiv*. doi:[10.1101/201178](https://doi.org/10.1101/201178)

- Shirley, M., Zhaorong, M., Pedersen, B., & Wheelan, S. (2015). Efficient "pythonic" access to FASTA files using pyfaidx. *PeerJ PrePrints*. doi:[10.7287/peerj.preprints.970v1](https://doi.org/10.7287/peerj.preprints.970v1)
- Tate, J., Bamford, S., Jubb, H., Sondka, Z., Beare, D., Bindal, N., & others. (2018). COSMIC: the Catalogue Of Somatic Mutations In Cancer. *Nucleic Acids Research*. doi:[10.1093/nar/gky1015](https://doi.org/10.1093/nar/gky1015)
- Yates, A., Achuthan, P., Akanni, W., Allen, J., Allen, J., Alvarez-Jarreta, J., & others. (2019). Ensembl 2020. *Nucleic Acids Research*. doi:[10.1093/nar/gkz966](https://doi.org/10.1093/nar/gkz966)