



HHS Public Access

Author manuscript

Nat Methods. Author manuscript; available in PMC 2016 April 01.

Published in final edited form as:

Nat Methods. 2015 October ; 12(10): 966–968. doi:10.1038/nmeth.3505.

SpeedSeq: Ultra-fast personal genome analysis and interpretation

Colby Chiang^{1,2}, Ryan M. Layer^{3,4}, Gregory G. Faust², Michael R. Lindberg², David B. Rose², Erik P. Garrison^{5,4}, Gabor T. Marth^{3,4}, Aaron R. Quinlan^{3,4}, and Ira M. Hall^{1,6,*}

¹McDonnell Genome Institute, Washington University School of Medicine, St. Louis, MO, USA

²Department of Biochemistry and Molecular Genetics, University of Virginia School of Medicine, Charlottesville, VA, USA

³Department of Human Genetics, University of Utah School of Medicine, Salt Lake City, UT, USA

⁴USTAR Center for Genetic Discovery, University of Utah School of Medicine, Salt Lake City, UT, USA

⁵Wellcome Trust Sanger Institute, Hinxton, UK

⁶Department of Medicine, Washington University School of Medicine, St. Louis, MO, USA

Abstract

SpeedSeq is an open-source genome analysis platform that accomplishes alignment, variant detection and functional annotation of a 50× human genome in 13 hours on a low-cost server, alleviating a bioinformatics bottleneck that typically demands weeks of computation with extensive hands-on expert involvement. SpeedSeq offers competitive or superior performance to current methods for detecting germline and somatic single nucleotide variants, indels, and structural variants, and includes novel functionality for streamlined interpretation.

Technical advances in second-generation DNA sequencing technologies have reduced both the cost and time required to generate whole-genome sequencing (WGS) data, creating opportunities in healthcare and academic research to survey the human genome with unprecedented depth and scope. However, bottlenecks in computational processing and variant interpretation have hindered adoption of these technologies for time-sensitive and large-scale projects. Using a standard pipeline based on BWA¹, GATK², SAMtools³, and Picard, processing a 50× human genome from raw sequence data to variant calls on a 32-thread server requires 60-70 hours (Supplementary Note 1). Furthermore, distinguishing

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use:http://www.nature.com/authors/editorial_policies/license.html#terms

*Corresponding Author, ihall@genome.wustl.edu.

Author contributions

C.C. wrote SpeedSeq and analyzed the data. R.M.L. advised on LUMPY implementation, G.G.F. contributed SAMBLASTER features, M.R.L. assisted with cloud implementation, and D.B.R. parallelized CNVnator. E.P.G. and G.T.M. advised on implementing FreeBayes. A.R.Q. contributed GEMINI features and advised on software design. I.M.H. conceived and designed the study. C.C. and I.M.H. wrote the manuscript.

Competing financial interests

The authors declare no competing financial interests.

pathogenic from benign mutations is a labor-intensive process that can take hours or days of manual curation per patient⁴.

SpeedSeq is a suite of open-source software designed for rapid whole-genome variant detection and interpretation. Its modular architecture and universal formats confer adaptability to a variety of experimental designs and compatibility with standard industry software (Fig. 1a). It achieves superior processing efficiency through rapid duplicate marking with SAMBLASTER⁵, through balanced parallelization of external applications, and by executing non-dependent pipeline components simultaneously. SpeedSeq translates raw 50× WGS data into prioritized single nucleotide variants (SNVs), short insertions and deletions (indels), and structural variants (SVs) in 13 hours on a single 16-core server with 128 GB of RAM (current cost: <\$10,000), and our accelerated implementations show little to no difference compared to original software (Supplementary Note 1). This represents at minimum a several-fold speed increase over current practices using typical computing resources.

We assessed the accuracy of SpeedSeq's SNV and indel calls against the Genome in a Bottle Consortium (GIAB) truth set in the well-studied human sample, NA12878 (2,803,144 SNVs and 364,031 indels)⁶. SpeedSeq achieved a sensitivity of 99.9% and 89.9% for germline SNVs and indels respectively, with acceptably low false discovery rates (FDR) (0.4% and 1.1%) (Fig. 1b,c). These detection rates exceeded the GATK Unified Genotyper's (GATK-UG) sensitivity (SNVs: 99.7%, indels: 89.0%) with a similar FDR (SNVs: 0.5%, indels: 1.0%;). The GATK Haplotype Caller (GATK-HC) showed superior indel detection sensitivity (SNVs: 99.8%, indels: 95.7%) with lower FDR for both variant types (SNVs: 0.2%, indels: 0.6%). SpeedSeq's implementation of FreeBayes therefore exhibits comparable – albeit slightly inferior – performance to GATK-HC when tested on the GIAB callset⁷. However, the GIAB truth set is biased towards GATK because it was primarily derived from GATK-based analyses. We therefore assessed SpeedSeq's performance against an unbiased truth set of 689,788 SNVs at 2,177,040 sites (Illumina Omni 2.5) in which SpeedSeq attained the highest sensitivity at the minor expense of specificity compared to GATK-UG or GATK-HC (Supplementary Fig. 1). Miscalled variants were enriched in repetitive regions of the genome and adjacent to assembly gaps (Supplementary Note 2 and Supplementary Table 1). SpeedSeq also supports joint multi-sample variant calling and *de novo* germline mutation detection in families (Supplementary Note 3), which is crucial for clinical applications such as rapid newborn diagnosis⁸.

Cancer genome analysis is a common WGS application in research and clinical environments, and can be a time-sensitive component of patient care. To emulate a WGS dataset from a heterogeneous tumor-normal pair, we defined NA12877 as the “normal” sample and pooled raw data from his 11 children in equal proportions to generate a single 50× “tumor” sample. The 875,206 SNVs present in the mother (NA12878) but absent from the father (NA12877) were defined as somatic mutations, with variant allele frequencies (VAFs) ranging from 0.05 to 0.5 (Supplementary Fig. 2a). Using this evaluation paradigm, we compared SpeedSeq's performance to three other leading somatic variant calling tools: MuTect⁹, SomaticSniper¹⁰, and VarScan 2¹¹. SpeedSeq recalled 96.6% of the somatic variants in the “tumor” with a FDR of 3.3%, outperforming SomaticSniper in both

sensitivity and specificity, and delivering competitive performance against MuTect and VarScan 2 (Fig. 1d, Supplementary Fig. 2b,c).

To test SpeedSeq's performance on real cancer data, we obtained WGS reads (50× tumor, 30× normal) from five tumor-normal pairs with validated somatic mutations ascertained by deep exome sequencing from The Cancer Genome Atlas (TCGA). SpeedSeq recalled 96.4% of 2,746 orthogonally validated mutations across all five datasets including 98.8% of mutations in genes that have been causally implicated in cancer¹² (Supplementary Table 2).

Ascertainment of structural variants – copy number variants (CNVs), balanced rearrangements and mobile element insertions – is a critical component of comprehensive genome analysis. SV detection poses two key technical challenges. First, SVs are extremely difficult to detect reliably¹³. Second, functional interpretation of SVs requires specialized logic due to their variable size and diverse configurations, and because SV breakpoints are often mapped imprecisely. Due to these challenges, few established genome analysis pipelines attempt to rigorously detect and interpret SVs.

SpeedSeq achieves comprehensive SV analysis with a suite of three complementary tools that are sensitive to a range of SV signals. At its core is LUMPY, a state-of-the-art breakpoint detection tool that integrates split-read and discordant paired-end data¹⁴. Next, a custom parallelized implementation of CNVnator uses read-depth analysis to detect CNVs that may be invisible to LUMPY due unmappable or repetitive sequence at their breakpoints¹⁵. Finally, SpeedSeq genotypes SVs with SVTyper, a novel Bayesian likelihood algorithm that can operate on copy-neutral events such as inversions and translocations as well as CNVs. This step produces SV genotypes that are crucial for meaningful variant interpretation, as well as quantitative estimates of breakpoint allele frequencies that allow inference of the fraction of tumor cells that carry a particular variant.

Measuring SV detection performance on real data is difficult due to the lack of established truth sets. If we accept the 1000 Genomes Project (1KGP) deletion callset for NA12878 as ground truth^{16,17}, SpeedSeq achieves a sensitivity of 61.9% (2089/3376) and positive predictive value of 60.8% (2089/3438) for detecting deletions, which is consistent with our recent comparative performance tests for LUMPY¹⁴ and by inference shows that SpeedSeq achieves state-of-the-art SV detection relative to other tools. However, this test underestimates absolute performance due to known false positives and negatives in the 1KGP callset. We therefore developed a composite strategy in which SVs in NA12878 could be validated either by overlap with split-read mapping of deep (30×) long-read data from PacBio and Illumina Moleculo platforms or by overlap with 1KGP. Based on this hybrid approach, SVs with quality scores of 100 or greater show a positive predictive value of 86.0% (2823/3282) (Fig. 1e, Supplementary Fig. 3). Virtually none of these SVs are likely to have validated by random chance, as 100 permutations of the callset resulted in a validation rate of 0.073% ($\pm 6.1E-3$, 95% CI). Moreover, SVTyper's quality scores provide a tunable parameter for refining callsets to a desired confidence threshold. By requiring both paired-end and split-read support, users may generate an extremely high confidence callset of 1,663 SVs with a 97.8% validation rate.

As an independent measure of SV detection and genotyping performance, we developed a haplotype-based test that exploits the structure of the CEPH 1463 pedigree. First, we phased the pedigree by SNV transmission to produce haplotype lineage maps, allowing us to attribute an average of 63.0% of the mappable genome of each F₂ individual to a particular founding grandparent (Fig. 1f). Next, we performed joint SV detection on the pedigree to generate 1,722 high-confidence autosomal SVs that could be assigned to a founding grandparent by transmission, resulting in a truth set of 8,397 predicted SV observations across the 11 grandchildren with known genotypes. SpeedSeq showed a detection sensitivity of 90.2% (7,578/8,397) for these predicted SVs, encompassing 1,660 of the 1,722 unique variants (Supplementary Table 3). Among the SVs that were detected, SVTyper reported the correct genotype at 96.6% (6,845/7,083) of heterozygous variants and 72.3% (358/495) of homozygous variants. Moreover, the high specificity of this callset is apparent from the infrequency of Mendelian violations (5.0%) and the consistent co-segregation of SVs with SNV-based haplotypes (93.8%) (Supplementary Table 4).

Results from SpeedSeq seamlessly integrate into the GEMINI variant interpretation framework, which annotates calls with information from external databases including dbSNP, ENCODE, ClinVar, CADD, ESP, and ExAC for efficient filtering with command line queries or a graphical browser interface¹⁸. In concert with SpeedSeq, we have made numerous enhancements to GEMINI, particularly in handling structural variants and interpreting somatic mutations. Users can rapidly prioritize somatic mutations through queries on two newly added databases: the COSMIC catalogue of somatic mutations in cancer¹² and DGIdb, the Drug-Gene Interaction database¹⁹. In addition, GEMINI can now identify structural variants that alter gene dosage or interrupt transcripts, as well as putative somatic gene fusions affecting COSMIC cancer genes.

Finally, to provide an example of a typical cancer analysis interpretation, we performed somatic variant calling on the tumor-normal pair of an invasive breast carcinoma from TCGA that carries a known gene fusion²⁰. With four concise commands and less than an hour of computation, we loaded the VCF file into GEMINI, filtered variant calls for high-confidence, clinically informative somatic mutations, and predicted gene fusion events (Fig. 2). These analyses demonstrate the ease with which high impact somatic point mutations and genomic rearrangements can be identified using the SpeedSeq framework.

Online methods

Software availability

The SpeedSeq v0.0.3a source code, documentation, and example data files are available in Supplementary Software, as well as at <https://github.com/cc2qe/speedseq>.

Hardware

All timings reported herein were performed on a single machine with 128 GB RAM and two Intel Xeon E5-2670 processors, each with 16 threads.

Data

We benchmarked SpeedSeq's processing time using the NA12878 genome from the Illumina Platinum Genomes dataset (European Nucleotide Archive: ERP001960), which comprises 50× WGS datasets for each of the 17 members of the three-generation CEPH 1463 pedigree (Supplementary Fig. 4).

Whole-genome sequencing data from five matched tumor-normal pairs and their orthogonally validated somatic mutations were obtained from The Cancer Genome Atlas (TCGA). These included three colorectal tumors (TCGA-A6-6141, TCGA-CA-6718, TCGA-D5-6540), one ovarian tumor (TCGA-13-0751), and one breast tumor (TCGA-B6-A0I6). Raw FASTQ reads were down-sampled to 50× coverage in the tumor and 30× coverage in the normal sample. Samples were processed with SpeedSeq for alignment, somatic mutations, and structural variants using default parameters and then loaded into GEMINI for variant interpretation. We also analyzed WGS data from a tumor-normal pair (63× tumor, 49× normal coverage) of a patient with an invasive breast carcinoma (TCGA-E2-A14P) containing a previously reported gene fusion between TBL1XR1 and PIK3CA²⁰.

FASTQ alignment and BAM processing

SpeedSeq aligns paired-end FASTQ files to the human GRCh37 reference genome with BWA-MEM, using the “-M” flag to mark shorter alignments as secondary. Aligned reads are streamed directly into SAMBLASTER⁵, which seizes idle CPU cycles that are periodically liberated each time BWA reads a FASTQ data chunk into the buffer. Marking duplicates on the pre-sorted BAM file allows simultaneous extraction of discordant read-pairs and split-read alignments, followed by rapid sorting and BAM compression with Sambamba²¹.

SNV and indel detection strategy

SpeedSeq runs FreeBayes version 0.9.16 with “--min-repeat-entropy 1” and “--experimental-gls” parameters for germline variant calling⁷. To increase specificity, SpeedSeq also requires at least one read on both the left and the right to support the variant allele. For somatic variant detection, SpeedSeq uses parameters tuned to increase sensitivity over low frequency variants (--pooled-discrete --genotype-qualities --min-alternate-fraction 0.05 --min-alternate-count 2 --min-repeat-entropy 1), and reports a somatic score (SSC) to estimate the confidence of each variant. The somatic score is the sum of the log odds ratios of the tumor (LOD_T) and normal (LOD_N) based on the genotype likelihood probabilities from FreeBayes (P_T and P_N for tumor and normal genotype probabilities respectively). The SSC is the preferred tuning parameter since it is robust to sequencing depth by design, however, the minimum alternate fraction and minimum alternate count may also be adjusted on the SpeedSeq command line.

$$LOD_T = \log \frac{P_T(\text{alternate})}{P_T(\text{reference})} \quad (1)$$

$$LOD_N = \log \frac{P_N(\text{alternate})}{P_N(\text{reference})} \quad (2)$$

$$SSC = LOD_T + LOD_N \quad (3)$$

SpeedSeq's implementation of FreeBayes is parallelized over 34,123 windowed regions of the GRCh37 genome using GNU Parallel²². We generated these regions, which average 84 kb in length, by partitioning the genome into bins of approximately equal numbers of reads based upon the aggregate coverage depth of all 17 members of the CEPH 1463 family pedigree and excluding high depth sequences (Supplementary Note 4 and Supplementary Fig. 5). This binning scheme balances the computational load over the FreeBayes instances by allocating processors based on the quantity of expected input data. It is 13.3-fold faster than the single-threaded version and 34.9% faster than naïve parallelization over each chromosome (Supplementary Note 1).

Structural variation detection and genotyping strategy

SpeedSeq runs LUMPY with “-msw 4 -tt 0 min_clip 20 min_non_overlap 101 min_mapping_threshold 20 discordant_z 5 back_distance 10”, and weights of 1 for both paired-end and split-read evidence. SpeedSeq's implementation of CNVnator parallelizes the genome by chromosome and performs copy number segmentation with a window size of 100 bp.

SVTyper is a maximum-likelihood Bayesian classification algorithm that infers an underlying genotype at each SV. Alignments at SV breakpoints either support the alternate allele with discordant or split-reads, or they support the reference allele with concordant reads/read-pairs that span the breakpoint. The ratio and quantity of these observations allow probabilistic inference of genotype likelihood. Under the assumption of diploidy, the set of possible genotypes at any locus is $G = \{\text{reference}, \text{heterozygous}, \text{homozygous}\}$. We defined the function S , where $S(g)$ is the prior probability of observing a variant read in a single trial given a genotype g at any locus. These priors were set to 0.1, 0.4, 0.8 for reference, heterozygous, and homozygous deletions respectively. Assuming a random sampling of reads, the number of observed alternate (A) and reference (R) reads (scaled by mapping quality, $10^{-(\text{mapq}/10)}$) will follow a binomial distribution $B(A+R, S(g'))$, where $g' \in G$ is the true underlying genotype. Using Bayes's theorem we can derive the conditional probability of each underlying genotype state from the observed read counts (Eq. 4), assuming an *a priori* probability $P(g)$ of 1/3 for each genotype. Finally, we calculate \hat{g} as the inferred genotype for the variant. Since the algorithm only interrogates SVs in the VCF file that have passed LUMPY filters as non-reference, it reports the more likely genotype of heterozygous or homozygous alternate states.

$$S(g) = \begin{cases} 0.1 & \text{if } g = \text{homozygous reference} \\ 0.4 & \text{if } g = \text{heterozygous} \\ 0.8 & \text{if } g = \text{homozygous alternate} \end{cases} \quad (4)$$

$$P(A, R|g) = \binom{A+R}{A} \cdot S(g)^A \cdot (1 - S(g))^R \quad (5)$$

$$\begin{aligned} P(g|A, R) &= \frac{P(A, R|g) \cdot P(g)}{P(A, R)} \\ &= \frac{P(A, R|g) \cdot P(g)}{\sum_{g \in G} P(A, R|g) \cdot P(g)} \quad (6) \end{aligned}$$

$$\hat{g} = \underset{g \in G}{\operatorname{argmax}} P(g|A, R) \quad (7)$$

SNV and indel evaluation

We compared SpeedSeq's germline SNV and indel variant calling against two independent truth sets for NA12878, one derived from the Genome in a Bottle (GIAB) NA12878 gold standard calls and the other based on Omni microarray data from the 1000 Genomes Project (1KGP). The GIAB 2.17 truth set contained 2,803,144 SNVs and 364,031 indels within highly confident regions (excluding segmental duplications, simple repeats, decoy sequence, and CNVs), spanning 2.2 Gb (77.6% of the mappable genome) for which non-variant sites could be confidently considered homozygous reference. The Omni microarray truth set contained 2,177,040 informative SNVs of which 689,788 were non-reference in NA12878, excluding markers within 50 bp of known indels.

We aligned NA12878 raw reads from the Illumina Platinum data with SpeedSeq, and then called germline SNVs and indels using SpeedSeq default parameters. To evaluate SpeedSeq's performance against other standard tools, we also processed the aligned BAM files according to the Genome Analysis Toolkit (version 3.2-2-gec30cee) best practices workflow, including realignment around indels, base recalibration, and variant calling with Unified Genotyper (GATK-UG) and Haplotype Caller (GATK-HC). Variant quality score recalibration was performed on the GATK results using a passing tranche filter of <99%. We normalized and compared variant calls according to the GIAB protocol, with vcfallelicprimatives, GATK's LeftAlignAndTrimVariants, and VcfComparator^{2,6}. We filtered variants for sensitivity and FDR against the GIAB truth set using a minimum quality score of 100 for GATK tools, and 1 for SpeedSeq (open circles, Fig. 1b,c).

To evaluate performance in detecting somatic variants, we generated a simulated tumor-normal matched pair from the CEPH 1463 family Illumina Platinum data. The "tumor" dataset was an equal mixture of all 11 members of the F₂ generation, down-sampled to 50× coverage and aligned with SpeedSeq. The father of the F₂ generation (NA12877) represented the 50× matched normal sample. For inclusion in the somatic SNV truth set, we required a variant to be diallelic, autosomal, in the NA12878 GIAB truth set, and called by Real Time Genomics (RTG) as non-reference in NA12878 and reference in NA12877^{6,23}. Additionally, variants were disqualified from the truth set if they violated Mendelian inheritance patterns. These criteria resulted in a set of 875,206 high confidence SNVs covering 77.6% of the mappable genome. The truth set of variants in the chimeric tumor

followed the expected binomial pattern of inheritance in her children, with a peak at 0.5 VAF from homozygous SNVs in NA12878 (Supplementary Fig. 2a).

We processed the simulated tumor data with SpeedSeq, MuTect 1.1.4, SomaticSniper, and VarScan 2 using parameters designed to target variants as low as 5% variant allele fraction. Receiver operating characteristic (ROC) curves were generated by varying somatic score (SSC) for SpeedSeq, SomaticSniper, and VarScan 2. For MuTect, which does not produce a single quality score for somatic variants, we varied the t_lod_fstar value to construct the ROC curve.

Structural variant evaluation

We constructed the 1KGP truth set by integrating deletions from the Pilot and Phase 1 callsets^{16,17}. For long-read validation of SV breakpoints, we obtained 30× PacBio (ftp://ftp-trace.ncbi.nih.gov/1000genomes/ftp/technical/working/20131209_na12878_moleculo/) and Moleculo (ftp://ftp-trace.ncbi.nih.gov/1000genomes/ftp/technical/working/20131209_na12878_pacbio/Schadt) data from 1KGP. We realigned the PacBio data with BWA-MEM 0.7.10 using the `-x pacbio` flag for consistency with the Moleculo alignments. Validations were performed according to previously published methods¹⁴. Custom scripts for this analysis are available at <https://github.com/hall-lab/long-read-validation>. To construct haplotype maps of the CEPH 1463 F₂ genomes, we called SNVs with SpeedSeq on the entire 17-member pedigree, and phased SNVs by transmission at polymorphic sites in the parents. We smoothed the chromosomes for contiguous blocks of inheritance by selecting informative bases where 95% of each run of 101 SNVs reported a consistent parent-of-origin. We then merged regions that shared inheritance and were within 100 kb of each other. This allowed us to trace an average of 1.8 Gb (63.4%) of each F₂ chromosome back to a particular grandparent, encapsulating meiotic crossovers that occurred in the F₁ germline (Fig. 1f). We then used SpeedSeq to jointly call structural variants on the entire pedigree, filtering for deletions that had at least seven pieces of support in at least one member of the pedigree, had legal Mendelian transmission, and whose origin could be unambiguously attributed to a single grandparent. Variants for which the founding grandparent by SV transmission agreed with the founding grandparent by SNV phasing were considered to be concordant, with strong supporting evidence for their authenticity. To test whether the 1,722 informative SVs were representative of the dataset as a whole, and not of misleadingly high quality due to their ascertainment criteria, we assessed their validation rate as above using the 1KGP callset and long-read sequencing (Supplementary Table 4). The 1,722 informative SVs had a similar validation rate as the remaining 6,734 SVs, suggesting that they are representative of overall callset quality.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank A. Abyzov for helpful discussions about CNVnator. This work was supported by a US National Institutes of Health (NIH) training grant to C.C. (T32 GM007267), a NIH/NHGRI grant to A.R.Q.

(R01HG006693), a NIH/NHGRI center grant (U54 HG003079), and a NIH New Innovator Award (DP2OD006493-01) and Burroughs Wellcome Fund Career Award to I.M.H.

References

1. Li, H. 2013. [arXiv.org](https://arxiv.org/abs/1303.3997)1303.3997
2. DePristo MA, et al. *Nature Genetics*. 2011; 43:491–498. [PubMed: 21478889]
3. Li H, et al. *Bioinformatics*. 2009; 25:2078–2079. [PubMed: 19505943]
4. Dewey FE, et al. *JAMA*. 2014; 311:1035–1045. [PubMed: 24618965]
5. Faust GG, Hall IM. *Bioinformatics*. 2014; 30:2503–2505. [PubMed: 24812344]
6. Zook JM, et al. *Nat Biotechnol*. 2014; 32:246–251. [PubMed: 24531798]
7. Garrison, E.; Marth, G. 2012. [arXiv.org](https://arxiv.org/abs/1207.3907)1207.3907
8. Kingsmore SF, Saunders CJ. *Science Translational Medicine*. 2011; 3:87ps23.
9. Cibulskis K, et al. *Nat Biotechnol*. 2013; 31:213–219. [PubMed: 23396013]
10. Larson DE, et al. *Bioinformatics*. 2012; 28:311–317. [PubMed: 22155872]
11. Koboldt DC, et al. *Genome Research*. 2012; 22:568–576. [PubMed: 22300766]
12. Futreal PA, et al. *Nat. Rev. Cancer*. 2004; 4:177–183. [PubMed: 14993899]
13. Alkan C, Coe BP, Eichler EE. *Nat. Rev. Genet*. 2011; 12:363–375. [PubMed: 21358748]
14. Layer RM, Chiang C, Quinlan AR, Hall IM. *Genome Biol*. 2014; 15:R84. [PubMed: 24970577]
15. Abyzov A, Urban AE, Snyder M, Gerstein M. *Genome Research*. 2011; 21:974–984. [PubMed: 21324876]
16. 1000 Genomes Project Consortium. et al. *Nature*. 2010; 467:1061–1073. [PubMed: 20981092]
17. 1000 Genomes Project Consortium. et al. *Nature*. 2012; 491:56–65. [PubMed: 23128226]
18. Paila U, Chapman BA, Kirchner R, Quinlan AR. *PLoS Comput Biol*. 2013; 9:e1003153. [PubMed: 23874191]
19. Griffith M, et al. *Nat Meth*. 2013; 10:1209–1210.
20. Stransky N, Cerami E, Schalm S, Kim JL, Lengauer C. *Nat Comms*. 2014; 5:4846.

Online methods references

21. Tarasov A, Vilella AJ, Cuppen E, Nijman IJ, Prins P. *Bioinformatics*. 2015 doi:10.1093/bioinformatics/btv098.
22. Tange O. *The USENIX Magazine*. 2011; 36:42–47.
23. Cleary JG, et al. *J. Comput. Biol*. 2014; 21:405–419. [PubMed: 24874280]

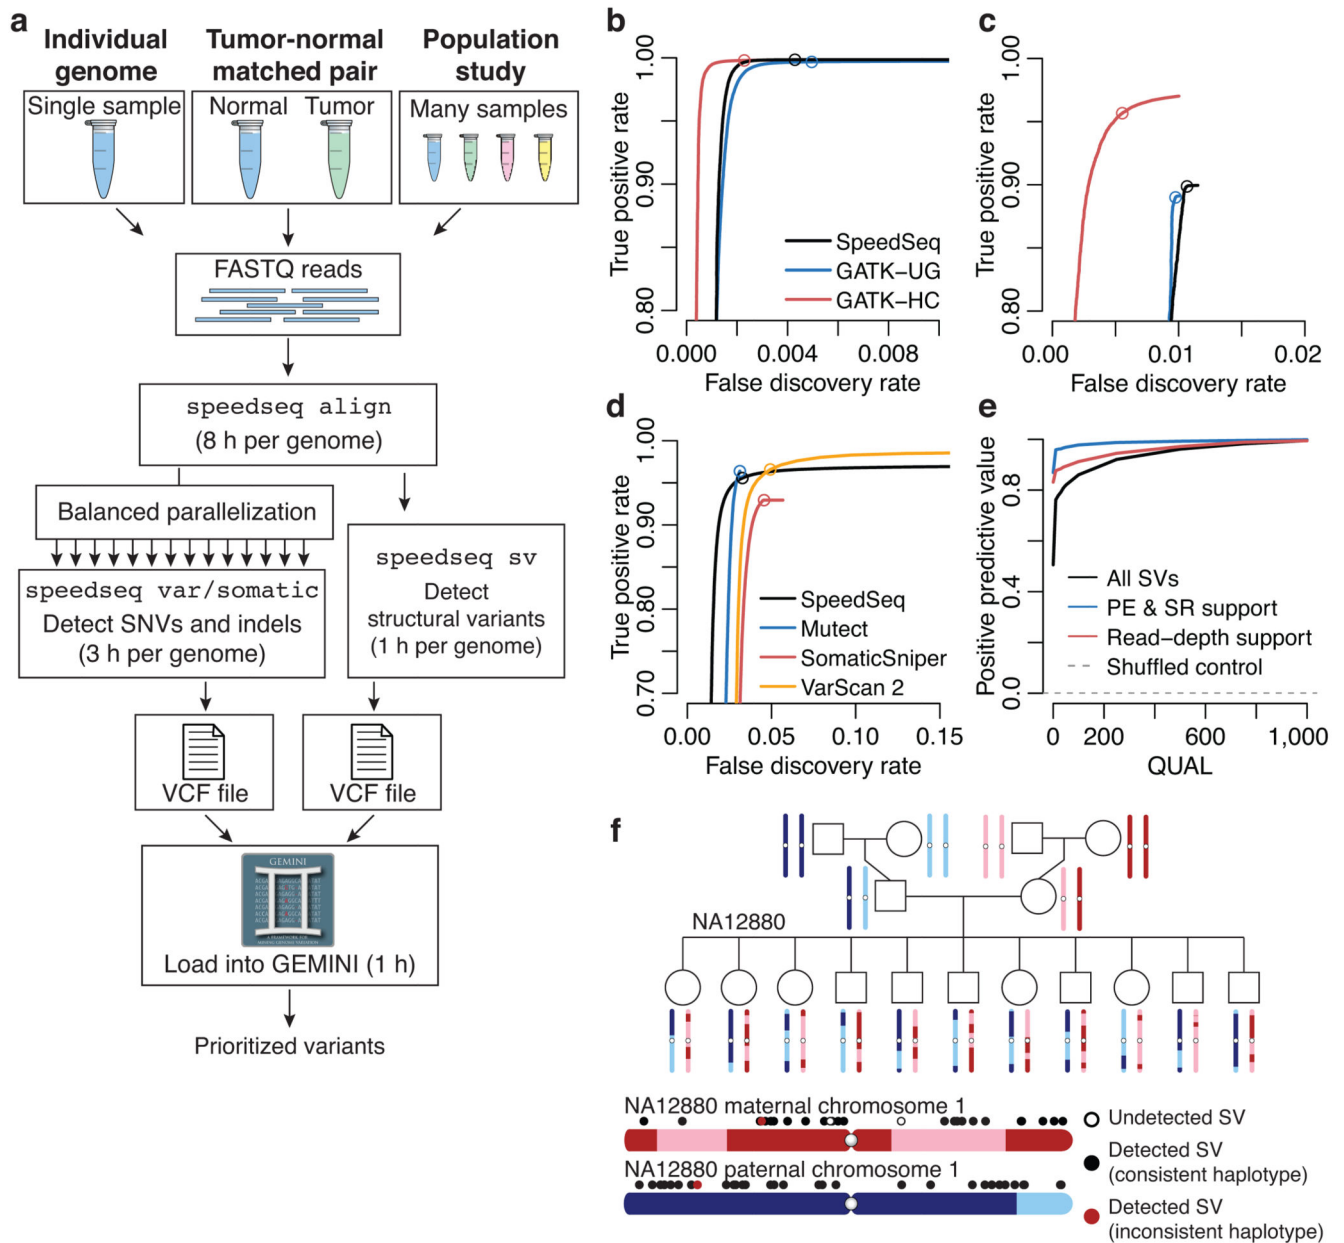


Figure 1. SpeedSeq workflow

(a) SpeedSeq converts raw reads into prioritized variants in 13 hours for a 50× human dataset. (b) Germline SNV ($N=2,803,144$) and (c) indel ($N=364,031$) receiver operating characteristic (ROC) curves over the Genome in a Bottle (GIAB) truth set for SpeedSeq, GATK Unified Genotyper (GATK-UG) and Haplotype Caller (GATK-HC). (d) Somatic SNV detection ROC curves for a simulated 50× tumor-normal pair using SpeedSeq and three other tools ($N=875,206$). Open circles (b-d) denote the data points reported in the main text. (e) SpeedSeq's SV detection performance by quality score of all SVs (black), those with split-read and paired-end support (blue), and those with read-depth support from CNVnator (red), as validated by either PacBio/Moleculo long-reads or the 1000 Genomes


Project. **(f)** Schematic of haplotype-based SV validation showing undetected (open circles), consistently segregating (black circles), and inconsistently segregating (red circles) SVs through the CEPH 1463 pedigree.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript



```
$ gemini load -v tumor.vcf.gz -t VEP -p sample.ped tumor.db
```

```
$ gemini set_somatic tumor.db
```

```
$ gemini actionable_mutations tumor.db
```

```
chr17 7578460 7578461 TP53 non_syn_coding
chr17 59861630 59861641 BRIP1 inframe_codon_loss
chr2 128046288 128046289 ERCC3 non_syn_coding
chrX 132838304 132838328 GPC3 splice_region
```

```
$ gemini fusions --in_cosmic_census --min_qual 1 tumor.db
```

```
chr3 178905990 chr3 176909982 - + complex PIK3CA TBL1XR1
```



Figure 2. Case study in a tumor-normal pair

A SpeedSeq workflow demonstrating the seven succinct commands required to process a tumor-normal pair (TCGA-E2-A14P) from raw FASTQ reads to clinically actionable somatic mutations with predicted damaging consequences. In this tumor, SpeedSeq detected a previously reported somatic gene fusion product between exon 1 of TBL1XR1 and exon 2 of PIK3CA²⁰.