



# HHS Public Access

Author manuscript

*Nat Methods*. Author manuscript; available in PMC 2015 April 01.

Published in final edited form as:

*Nat Methods*. 2014 October ; 11(10): 1033–1036. doi:10.1038/nmeth.3069.

## Accurate detection of *de novo* and transmitted indels within exome-capture data using micro-assembly

Giuseppe Narzisi<sup>1</sup>, Jason A. O'Rawe<sup>2,3</sup>, Ivan Iossifov<sup>1</sup>, Han Fang<sup>2,3</sup>, Yoon-ha Lee<sup>1</sup>, Zihua Wang<sup>1</sup>, Yiyang Wu<sup>2,3</sup>, Gholson J. Lyon<sup>2,3</sup>, Michael Wigler<sup>1</sup>, and Michael C. Schatz<sup>1</sup>

<sup>1</sup>Simons Center for Quantitative Biology, One Bungtown Road, Cold Spring Harbor Laboratory, NY, USA

<sup>2</sup>Stanley Institute for Cognitive Genomics, One Bungtown Road, Cold Spring Harbor Laboratory, NY, USA

<sup>3</sup>Stony Brook University, 100 Nicolls Rd, Stony Brook, NY, USA

### Abstract

We present an open-source algorithm, Scalpel, which combines mapping and assembly for sensitive and specific discovery of indels in exome-capture data. A detailed repeat analysis coupled with a self-tuning *k*-mer strategy allows Scalpel to outperform other state-of-the-art approaches for indel discovery, particularly in regions containing near-perfect repeats. We analyze 593 families from the Simons Simplex Collection and demonstrate Scalpel's power to detect long (> 20bp) transmitted events, and enrichment for *de novo* likely gene-disrupting indels in autistic children.

---

While the analysis of Single Nucleotide Variations (SNVs) has become a standard technique to study human genetics<sup>1</sup>, insertions and deletions in DNA sequences (indels) cannot be detected as reliably<sup>2,3</sup>. Indels are the second most common sources of variation in human genomes and the most common structural variant<sup>4</sup>. Within microsatellites (simple sequence repeats, SSRs, of 1 to 6bp motifs), indels alter the length of the repeat motif and have been linked to more than 40 neurological diseases<sup>5</sup>. Indels also play an important genetic component in autism: *de novo* indels that are likely to disrupt the encoded protein are nearly twice as abundant in affected children than in their unaffected siblings<sup>6</sup>.

---

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](http://www.nature.com/authors/editorial_policies/license.html#terms)

Corresponding author: G.N. (gnarzisi@cshl.edu).

**Authors contributions:** G.N. developed the software and conducted the computational experiments. G.N. and M.C.S. designed and analyzed the experiments. Y.W. assisted in designing the primers and performed the MiSeq validation experiments. J.A.O. designed the primers and analyzed the MiSeq data. H.F. and J.A.O. assisted with the computational experiments for the comparative analysis between different variant detection pipelines. G.J.L. planned and supervised the experimental design for indel validation. Z.W. designed the primers and performed experiments for the validation of *de novo* and transmitted indels in the SSC. I.I., Y.L., and M.W. assisted with the analysis of the SSC. G.N. and M.C.S. wrote the manuscript with input from all authors. All of the authors have read and approved the final manuscript.

**Competing Financial Interests:** The authors declare no competing financial interests.

Detecting indels is challenging for several reasons: (1) reads overlapping the indel sequence are more difficult to map<sup>7</sup> and may be aligned with multiple mismatches rather than with a gap; (2) irregularity in capture efficiency and non-uniform read distribution increase the number of false positives; (3) increased error rates makes their detection very difficult within microsatellites; and, as shown in this study, (4) localization, near identical repetitive sequences can create high rates of false positives. For these reasons, the size of indels detectable by available software tools has been relatively small, rarely more than a few dozen base pairs<sup>8</sup>.

Two major paradigms are currently used for detecting indels. The first and most common approach is to map all of the input reads to the reference genome using a read mapper (e.g., BWA, Bowtie, Novoalign), although the available algorithms are not as effective for mapping across indels of more than a few bases. Advanced approaches exploit paired-end information to perform local realignments to detect longer mutations (e.g., GATK UnifiedGenotyper<sup>1</sup> and Dindel<sup>9</sup>), although, in practice, their sensitivity is greatly reduced for longer variants (> 20bp). Split-read methods (e.g., Pindel<sup>10</sup> and Splitread<sup>11</sup>) can theoretically find deletions of any size, but they have limited power to detect insertions due to the short read length of current sequencing technologies. The second paradigm consists of performing *de novo* whole-genome assembly of the reads, and detecting variations between the assembled contigs and the reference genome<sup>12,13</sup>. While having the potential to detect larger mutations, in practice this paradigm is less sensitive since detecting indels requires a fine-grained and localized analysis to correctly report homozygous and heterozygous mutations. Recently, three approaches have been developed that use *de novo* assembly for variation discovery: GATK HaplotypeCaller, SOAPindel<sup>14</sup>, and Cortex<sup>15</sup>. An other recent approach, TIGRA<sup>16</sup>, also uses localized assembly, but it has been tailored for breakpoints detection, without reporting the indel sequence.

We present a DNA sequence micro-assembly pipeline, Scalpel, for detecting indels within exome-capture data (Fig. 1). By combining the power of mapping and assembly, Scalpel carefully searches the de Bruijn graph for sequence paths (contigs) that span each exon. The algorithm includes an on-the-fly repeat composition analysis of each exon, coupled with a self-tuning *k*-mer strategy.

We confirm previous findings that nine standard algorithms have reduced power to detect large (> 20bp) indels<sup>14,15</sup> using simulated reads: Scalpel, SOAPindel<sup>14</sup>, GATK-HaplotypeCaller, GATK-UnifiedGenotyper, SAMtools<sup>17</sup>, FreeBayes<sup>18</sup>, Platypus ([www.well.ox.ac.uk/platypus](http://www.well.ox.ac.uk/platypus)), lobSTR<sup>19</sup>, and RepeatSeq<sup>20</sup> (Supplementary Note 1 and 2, Supplementary Fig. 1 and 2, Supplementary Tables 1–5). We also performed a large-scale validation experiment involving ~1000 indels from one single exome. The individual has a severe case of Tourette Syndrome and obsessive compulsive disorder (sample ID: K8101-49685), and was sequenced to 20× coverage over 80% of the exome target using the Agilent 44MB SureSelect capture protocol and Illumina HiSeq2000 paired-end reads, averaging 90bp in length, after trimming. Indels were called using the three pipelines that had performed best with our simulated reads: Scalpel v0.1.1 beta, SOAPindel v2.0.1 and GATK HaplotypeCaller v2.4.3 (Online Methods). Interestingly, there is only ~37% concordance among calls made by all of the pipelines, and each method reports hundreds of

indels unique to that pipeline (Fig. 2a), which is in close agreement with a recent analysis<sup>2</sup>. An update to GATK to version 3.0 was released after our initial validation experiments, but we also assessed its accuracy with a second blinded re-sequencing experiment (Fig. 2b, Supplementary Note 3, Supplementary Fig. 3, Supplementary Table 6 and 7).

From the concordance rate alone, it is hard to judge the quality of indels unique to each pipeline, as these could either represent superior sensitivity or poor specificity. The size distribution of indels called by the HaplotypeCaller (v2.4.3) has a bias towards deletions whereas SOAPindel has a bias towards insertions (Fig. 2b). Scalpel and HaplotypeCaller (v3.0) instead show a well-balanced distribution, in agreement with other studies of human indel mutations<sup>8</sup>.

We further investigated the performance of the algorithms by a focused re-sequencing of a representative sample of indels using the more recent 250bp Illumina MiSeq sequencing protocol (Online Methods). Based on the data depicted in Figure 2a, we selected a total of 1,000 indels according to the following categories: (1) 200 random indels from the intersection of all pipelines; (2) 200 random indels only found by HaplotypeCaller (v2.3.4); (3) 200 random indels only found by SOAPindel; (4) 200 random indels only found by Scalpel; (5) 200 random indels of size  $\leq 30$ bp from the union of all three algorithms.

Due to possibly ambiguous representation, indels positions are “left-normalized”<sup>2</sup>. However, some ambiguity can still remain, especially within microsatellites, so we computed validation rates using two different approaches. (1) *Position-based*: an indel is considered valid if a mutation with the same coordinate exists in the validation data (Fig. 3a). (2) *Exact-match*: an indel is considered valid if there is a mutation with the same coordinate and sequence in the validation data (Fig. 3b).

As expected, indels detected by all pipelines have a high validation rate and their sizes follow a lognormal distribution (Supplementary Fig. 4). However, the validation rate varies dramatically for each tool. Respectively, only 22% and 55% of the HaplotypeCaller (2.4.3) and SOAPindel specific indels could be validated even when the less strict position-based approach was used, whereas 77% of Scalpel's specific indels are true positive. For the long indels: less than 10% called by SOAPindel and HaplotypeCaller passed validation. (Figure 3c and Supplementary Table 8). The new version of GATK (v3.0) has largely removed the bias towards deletions (Fig 2b), but find that Scalpel still outperforms HaplotypeCaller (Supplementary Note 3). Scalpel shows substantially higher validation rate (76%) for longer indels (>5bp) compared to HaplotypeCaller v3.0 (27%).

We further divide the results to separately report the validation rate for indels within microsatellites. SOAPindel shows an appreciably higher rate of false-positives within microsatellites (“SSRs-only” in Fig. 3a and 3b). When microsatellites are excluded (“no-SSRs” in Fig. 3a and 3b), the performance of SOAPindel and HaplotypeCaller decline, while Scalpel's validation rate is only slightly reduced. Figure 3a and Figure 3b also illustrate the relative abundance of indels within microsatellites called by each tool, although HaplotypeCaller seems to filter against these. Finally, when switching from position-based to exact-match, indels within microsatellites show notable reduction in validation rate. This

phenomenon is due to their high instability and higher error rates, and in fact it is not unusual to have more than one candidate mutation at a microsatellite locus.

We further inspected the sequence composition of all false-positive long indels. Specifically, we reanalyzed the 129 SOAPindel invalid long mutations using Scalpel. The majority of these mutations (115) overlap repeat structures where the reference contains a perfect or near-perfect repeat (Supplementary Fig. 5). In contrast, of the 62 false-positive long indels from HaplotypeCaller, only 16 overlap a repeat. The remaining false positive deletions appear to be due to an aggressive approach used by the algorithm when processing soft-clipped sequences. The soft-clipped reads in false positive indels for HaplotypeCaller are highly variable, and are conjectured to be mapping artifacts of reads from different genomic locations (Supplementary Fig. 6). Finally we investigated the relationship between the false-discovery rate (FDR) and characteristic features (e.g., chi-square score and coverage) for 614 indels detected by Scalpel and validated by re-sequencing (Supplementary Note 4, Supplementary Fig. 7 and 8). In addition to highlighting the common trends, this analysis provides recommendation on how to select a chi-square score cutoff to achieve a given FDR.

Using Scalpel we detected a total of 3.3 million indels in exomes from 593 families from the Simons Simplex Collection, corresponding to an average of  $\sim 1,400$  ( $=3,388,139/(4*593)$ ) mutations per individual. Accounting for population frequencies, there were 27,795 distinct transmitted indels across the exomes. We find close agreement to the size distributions reported by Montgomery *et al.*<sup>8</sup> using low coverage whole-genome data from 179 individuals (Supplementary Fig. 9a). Direct comparison to those detected by the GATK-UnifiedGenotyper based mapping pipeline used by Iossifov *et al.*<sup>6</sup> shows that Scalpel has superior power to detect longer insertions (Supplementary Fig. 10). To estimate Scalpel's ability to discover transmitted mutations, we performed targeted re-sequencing of 31 long ( $> 29$ bp) transmitted indels. Excluding indels that failed to sequence (4), 21 passed validation (out of 27), which gives a 78% true positive rate. Three of the indels that did not pass validation were indeterminate with ambiguous alignments because they were either too long ( $> 70$ bp) or embedded in a repetitive region.

Within the coding sequence (CDS), frame-preserving indels are more abundant than frame-shifts (Supplementary Fig. 9b). In agreement with MacArthur *et al.*<sup>21</sup>, we detect a large number of transmitted loss-of-function (LOF) variants in protein-coding genes. Frame-shift mutations are found at lower frequency in the population when located in protein-coding sequences compared to intronic regions (Supplementary Fig 9c). Finally, we observe an enrichment of deletions over insertions (Supplementary Table 9 and Supplementary Fig. 11), with an overall 2:1 ratio across all annotation categories. Similar trends were reported in previous studies<sup>8,22</sup>.

Here we reanalyzed the data on autistic children and unaffected siblings<sup>6,23,24</sup> with Scalpel with the goals of examining *de novo* likely gene disrupting (LGD) mutations. We confirm an overabundance of frame-shift mutations in autistic patients<sup>6</sup> predict additional candidates, and extend the analysis to a larger number of families. Our re-analysis of a previous study with 200 SSC families<sup>23</sup> reports an enrichment of 11 LGD indels in autistic children

compared to 4 in their healthy siblings. In targeted re-sequencing of 102 candidate indels we confirmed 84 as *de novo* mutations, invalidated 11 and failed to sequence 7, giving an 88% *de novo* positive predictive rate.

In order to focus the list of candidate genes, we excluded mutations that are common in the population, and used stringent coverage filters (Online Methods) to select a total of 97 high quality *de novo* indels (Supplementary Data 1). Even after extending the population size from 343<sup>6</sup> to 593, the same 2:1 enrichment for LGD mutations is confirmed: 35 frame shifts in autistic children vs. 16 in siblings (*p*-value 0.01097) (Supplementary Table 10 and 11, Supplementary Note 5, Supplementary Fig. 12-16), other smaller studies came to similar conclusions<sup>23,24,25</sup> This result also holds for a larger collection of 1303 SSC families (not presented in this study). All together, in agreement with the previously reported results<sup>6</sup>, we find a significant overlap between the LGD target genes and the 842 FMRP-associated genes<sup>26</sup>. Specifically, 8 out of 35 LGDs in autistic children overlap with the 842 FMRP-associated genes.

## Online Methods

### The Scalpel pipeline

Scalpel is designed to perform localized micro-assembly of specific regions of interest in a genome with the goal of detecting insertions and deletions with high accuracy. It is based on the de Bruijn graph assembly paradigm where the reads are decomposed into overlapping *k*-mers, and directed edges are added between *k*-mers that are consecutive within any read<sup>27</sup>. Figure 1 shows the high level structure of the pipeline. (1) The pipeline begins with a fast alignment of the reads to the reference genome using BWA<sup>28,29</sup>. Importantly, these alignments are not directly used to call variations, but only to localize the analysis by identifying all the reads that have similarity to a given locus. Reads are then extracted in the region of interest (e.g., exon) including: (i) *well-mapped* reads, (ii) *soft-clipped* reads, and (iii) reads that *fail to map*, but are anchored by their mate. The latter two classes correspond to locations where the mapper encountered trouble aligning the reads, especially because of the large indels present, so it's necessary to include them in the assembly. (2) Once localized, the algorithm computes an on-the-fly assembly of the reads in the current region using the de Bruijn graph paradigm, specifically, reads are decomposed into overlapping *k*-mers (starting with a default *k*=25) and the associated graph is constructed. (3) Using the reference sequence, one source node and one sink node are selected according to the procedure described later in the “Graph traversal” section. (4) An on-the-fly analysis of the repeats in each region is used to automatically select the *k*-mer size to be used for the assembly, described in section “Repeat analysis”. (5) The graph is then exhaustively examined to find end-to-end paths that span the region. (6) After the sequences are assembled, they are aligned to the reference to detect candidate mutations using a sensitive gapped sequence aligner based on the Smith Waterman algorithm<sup>30</sup> targeted at the reference window. Finally, the above assembly process is applied using a sliding window approach over each target region. By default a window size of 400bp is used with a sliding factor of 100bp. The sliding window strategy is fundamental to handle the highly non-uniform read distribution across the target (Supplementary Fig. 17). A window size of 400bp is large

enough to assemble the majority of the exons into a single contig: ~95% of the human exon-targets are shorter than 400bp (Supplementary Fig. 18), however each assembly task is small enough for using in-depth techniques to optimize the assembly.

### Graph construction

Two critical components of the Scalpel algorithm are (i) construction of the de Bruijn graph and (ii) detection of sequence paths spanning the targeted region. Reads aligning to the region are extracted and decomposed into overlapping  $k$ -mers. In order to model the double stranded nature of the DNA, a bidirected de Bruijn graph is constructed<sup>31,32</sup>. The graph is then compressed by merging all non-branching chains of  $k$ -mers into a single node. Tips and low coverage nodes are removed according to input threshold parameters to remove obvious sequencing errors. Note that, differently from a traditional de Bruijn graph assembler, Scalpel does not use any threading strategy to resolve collapsed repeats. In principle, threading would allow resolution of repeats whose lengths are between  $k$  and the read length. However we observed in both real and simulated data that, due to the localized graph construction, if a bubble were not covered end-to-end by the reads, threading would either disconnect the graph or introduce errors. Repeats are instead handled differently, as explained in the next section,

### Repeat Analysis

Due to the highly non-uniform read depth distribution across the targeted region and the presence of near-perfect repeats that can mislead the assembly (Supplementary Note 6 and Supplementary Fig. 19), Scalpel implements a detailed repeat composition analysis coupled with a self-tuning  $k$ -mer strategy. Specifically, when assembling a window, Scalpel inspects both the base pair composition of the corresponding reference sequence as well as the resulting de Bruijn graph for the presence of cycles in the graph or near-perfect repeats in the assembled sequences. If a repeat structure is detected, the graph is discarded and a larger  $k$ -mer is selected. This process continues until a maximum  $k$ -mer length is reached, which is a function of the read length. If no  $k$ -mer value can be chosen to avoid the presence of repeats, the region is skipped and the next available region from the sliding window scheme is analyzed. This conservative strategy reduces the number of false-positive calls in highly repetitive regions, and, skips less than 2% of possible windows in the human exome. Note also that, once  $k$  is selected by the self-tuning  $k$ -mer strategy, the graph is “repeat free”, and there is no need to use threading to resolve small repeats.

The proposed self-tuning  $k$ -mer strategy is similar to the dynamic approach used by SOAPindel and TIGRA to reconnect a broken path in low coverage regions. However, SOAPindel searches for unused reads with gradually shorter  $k$ -mers until a path is formed or the lower bound on  $k$ -mer length has been reached; In TIGRA the user can specify the list of  $k$ -mers to use (by default only two: 15 and 25). Scalpel instead starts from a small  $k$ -mer value (input parameter) first and then gradually increases it, such that the smallest possible  $k$ -mer value that generates a “repeat free” graph is used for each region. This strategy has the advantage of better handling repetitive sequences, highly polymorphic regions, and sequencing errors: source and sink have higher chance to be selected (see section “Graph

traversal”) and a smaller  $k$ -mer reduces the chance of fragmented assembly in low coverage regions.

### Graph traversal

Once a valid de Bruijn graph is constructed, Scalpel examines the graph to find end-to-end sequence paths that span the target window. Because the coverage from exome capture data is highly variable, a special selection algorithm is used to find the edges of each window where coverage is present. First, two nodes in the graph are labeled as *source* and *sink* according to the following procedure: the reference sequence of the target region is scanned left-to-right to detect the first sequence of  $k$  bases that exactly matches one of the  $k$ -mers from the nodes in the graph, this node will be marked as the source. In a similar fashion the sink node is detected scanning the reference sequence right-to-left. Since every region is first inspected for repeats, source and sink can be safely selected at this stage. The automated strategy used by Scalpel to select the boundaries of the reference sequence improves over TIGRA's approach, where the reference region is selected based only on input parameters. After the source and sink nodes are identified, all possible source-to-sink paths are enumerated up to a max number (default 100,000) using a depth-first search (DFS) traversal of the graph, similarly to Sutta assembly algorithm<sup>33</sup>. Note that since the regions to assemble are very small, time and space computational complexities associated with large-scale whole-genome assembly are not relevant and an exact brute-force strategy can be efficiently applied.

If there are no repeat structures in the graph, all the candidate paths are enumerated and aligned to the portion of the reference sequence delimited by source and sink  $k$ -mers using the standard Smith-Waterman-Gotoh alignment algorithm with affine gap penalties. The list of candidate mutations is then generated. Under typical conditions, the assembler reports a single path for homozygous mutations and two paths for heterozygous mutations. For example, if the sample has an insertion in only one of the two haplotypes, the assembler would discover the indel and also the unmodified reference sequence. Note that a traditional sequence assembler would have selected only one of these two paths (usually with higher coverage) and discarded the other one. Scalpel instead examines both paths to distinguish, for example, between homozygous and heterozygous mutations. However, in practice, various factors in real data complicate the detection process and, sometimes, multiple paths are reported in the case of more exotic variations. For example, the Illumina sequencing platform is particularly error prone around microsatellites (e.g., homopolymer runs) and, as a consequence, multiple candidate alleles are elucidated by the data at these loci. Highly polymorphic regions are also prone to generate multiple paths and could be computationally demanding: if the distance between multiple nearby mutations is larger than the (automatically) selected  $k$ -mer value, each of the associated bubbles in the graph will give rise to two different paths. Finally it is important to note that SNVs are also computed by Scalpel but they are not reported in output. SNVs are used internally for important downstream analysis of the variants called in order, for example, to compute coverage around the indels and to correctly characterize the zygosity of the mutations.

## Exome capture data

Exome capture for the sample K8101-49685 was carried out using the Agilent 44MB SureSelect protocol and then sequenced on Illumina HiSeq2000 with average read length of 100bp. More than 80% of the target region was covered with depth of 20 reads or more. All of the HiSeq data have been submitted to the Sequence Read Archive (<http://www.ncbi.nlm.nih.gov/sra>) under project accession number SRX265476.

## MiSeq validation

A total of 1,400 indels were selected for MiSeq validation over the course of this study. MiSeq validation was initially performed on 1,000 indels prior to the release of GATK v3.0 (see the Results section for detailed selection criteria). After the release of GATK v3.0, we selected an additional 400 indels for MiSeq validation (see Supplementary Note 3 for detailed selection criteria). Out of these 400 indels, 215 were covered with more than 1,000 reads in the initial MiSeq dataset or in another MiSeq dataset reported by O'Rawe *et al.*<sup>2</sup> Thus, the second MiSeq validation experiment was performed on the remaining 185 indels. For both of the MiSeq validation experiments performed during the course of this study, PCR primers were designed using Primer 3 (<http://primer3.sourceforge.net>) to produce amplicons ranging in size from 200 to 350 bp, with indels of interest located approximately in the center of each amplicon. Primers were obtained from Sigma-Aldrich® in 96-well mixed plate format, 10 µmol/L dilution in Tris per Oligo. Upon arrival, all primers were tested for PCR efficiency using a HAPMAP DNA sample (Catalog ID NA12864, Coriell Institute for Medical Research, Camden, NJ, USA) and LongAmp® Taq DNA Polymerase (New England Biolabs, Beverly, MA, USA). PCR products were visually inspected for amplification efficiency using agarose gel electrophoresis. For the validation experiment, this same PCR was performed using sample K8101-49685 genomic DNA as template. PCR product was verified on E-Gel® 96 gels (Invitrogen Corp., Carlsbad, CA, USA) and subsequently pooled for ExoSAP-IT® (Affymetrix Inc., Santa Clara, CA, USA) cleanup. The cleanup product was further purified using QIAquick PCR Purification Kit (QIAGEN Inc., Valencia, CA, USA) and quantified by Qubit® dsDNA BR Assay Kit (Invitrogen Corp.). Library construction for the MiSeq Personal Sequencer platform (Illumina Inc.) was performed based on Illumina® TruSeq® DNA Sample Prep LS protocol (for the initial 1,000 indels) and TruSeq® Nano DNA Sample Preparation Guide (for the additional 185 indels), omitting the DNA fragmentation step. Finally, before being loaded onto the MiSeq machine, the quality and quantity of the sample was again verified using the Agilent DNA 1,000 Kit on the Agilent Bioanalyzer and with quantitative PCR (Kapa Biosystems Inc., Woburn, MA, USA). This protocol generated high quality 250 bp reads (paired end) with an average coverage of 47,018X (Supplementary Fig. 20). The reads were aligned with BWA-MEM (v0.7.5a) to the reference human genome hg19. The alignment was sorted with SAMtools (v0.1.18) and PCR duplicates were marked with the Picard tool set (v1.91). Indels were realigned with the GATK (version v2.6-4) using the IndelRealigner and base quality scores were recalibrated. Variants were then called with GATK UnifiedGenotyper. All of the MiSeq data have been submitted to the Sequence Read Archive (<http://www.ncbi.nlm.nih.gov/sra>) under project accession number SRX386284.



## Alignment

Sequencing reads from K8101-49685 exome-capture data were aligned using BWA (v0.6.2-r126) with default parameters to the human reference hg19. Alignments were converted from SAM format to sorted, indexed BAM files with SAMtools (v0.1.18). The Picard tool set (v1.91) was used to remove duplicate reads. These BAM files were used as input for all the indel callers used in this study. Reads coming from the re-sequencing experiments were also aligned using BWA. However, if the indel approaches half the size of the read length, even after target re-sequencing, mapping the reads containing the indel is problematic. The problem is emphasized if the indel is located towards the ends of the read (instead of in the middle). To avoid this problem we aligned sequencing reads containing long indels (> 30 bp) using Bowtie2<sup>34</sup> instead of BWA. Bowtie2 offers an end-to-end alignment mode that searches for alignments involving all of the read characters, also called an “untrimmed” or “unclipped” alignment. Specifically, we used the following parameter settings: “--end-to-end --very-sensitive --score-min L,-0.6,-0.6 -rdg 8,1 -rfg 8,1 --mp 20,20”.

## Variant Calling

indels for K8101-49685 were called using Scalpel, GATK HaplotypeCaller and SOAPindel as follows:

**Scalpel**—Scalpel (v0.1.1 beta) was run on the indexed BAM using the following parameter setting: “--single --lowcov 1 --mincov 3 --outratio 0.1 --intarget”. Indels showing high coverage unbalance were then removed (chi-square *k*-mer score > 20).

**GATK**—GATK software tools (v2.4-3 and v3.0) were used for improvement of alignments and genotype calling and refining with recommended parameters. BAM files were re-aligned with the GATK IndelRealigner, and base quality scores were re-calibrated by the GATK base quality recalibration tool. Genotypes were called by the GATK UnifiedGenotyper and HaplotypeCaller. According to the GATK recommendations, the Variant Quality Score Recalibration (VQSR) was not used for the K8101-49685 single exome experiment. Also hard filtering criteria (such as “QD < 2.0, ReadPosRankSum < -20.0 FS > 200.0”) were not used for v2.4.3 as they were aggressively removing long indels from the HaplotypeCaller calls (Supplementary Note 3), but were used instead for the more recent HaplotypeCaller from GATK v3.0. It is important to report that the bias towards deletions for HaplotypeCaller has been extensively reduced with the release of GATK 2.8 and 3.0 in December 2013 and March 2014 after our initial MiSeq re-sequencing experiments were completed (Supplementary Note 3). However, many research groups have already employed the older version of HaplotypeCaller for genetic studies and they are still extensively using it. We expect GATK, along with our own and other algorithms, to improve over time as new insights are made into the mutation mechanisms and error profiles.

**SOAPindel**—SOAPindel (v2.0.1) was run on the indexed BAM file using default parameters. According to SOAPindel documentation, putative indels are initially assumed to be located near the unmapped reads whose mates mapped to the reference genome. SOAPindel then executes a local assembly (*k*-mer=25 by default) on the clusters of

unmapped reads. The assembly results were aligned to reference in order to find the potential indels. To distinguish true and false positive indels, SOAPindel generates Phred quality scores, which take into consideration the depth of coverage, indel size, number of neighboring variants, distance to the edge of contig, and position of the second different base pair. Only those indels filtered by internal threshold are retained in the final indel call set.

Finally, for all pipelines we selected only indels located within the regions targeted by the exome capture protocol.

### The Simons Simplex Collection

The Simons Simplex Collection<sup>35</sup> (SSC) is a permanent repository of genetic samples from 2,700 families operated by SFARI (<http://sfari.org>) in collaboration with 12 university-affiliated research clinics. Each simplex family has one child affected with autism spectrum disorder, and unaffected siblings. Each genetic sample also has an associated collection of phenotype measurements and assays. The results presented in this work are based on a subset of the SSC composed of 593 families (2,372 individuals). Specifically this subset of the SSC collection corresponds to families that have been examined in three recent studies: 343 families from Iossifov et al.<sup>6</sup> (CSHL), 200 families from Sanders et al.<sup>23</sup> (Yale), and 50 families from O’Roak et al.<sup>24</sup> (University of Washington). We selected only family units of four individuals (father, mother, proband, one unaffected sibling), referred to as “quads,” for all analyses in this study.

### Analysis of *de novo* indels related to Autism

After eliminating all candidate positions that are common in the population, and thus unlikely to be related to the disorder, we re-assembled each region associated with the candidates indels across the family members using a more sensitive parameter setting for Scalpel. Specifically we reduce the starting  $k$ -mer value to 10 and turned off the removal of low coverage nodes. This step was important to adjust for possible allele imbalance favoring the reference allele over the mutation in the parents, but was impractical to do initially for the whole collection: lowering the  $k$ -mer and keeping all the nodes in the graph significantly increase the computation complexity of the algorithm. Then we selected *de novo* indels with chi-square  $k$ -mer score  $> 10.84$ . The chi-square  $k$ -mer score is computed using the standard formula for the chi-square test statistics ( $X^2$ ) but applied to the  $k$ -mer coverage of the reference and alternative alleles for the mutation according to the following formula:

$$X^2 = \frac{(C_o^R - C_e^R)^2}{C_e^R} + \frac{(C_o^A - C_e^A)^2}{C_e^A}$$

Where  $C_o^R$  and  $C_o^A$  are the observed  $k$ -mer coverage for the reference and alternative alleles, and  $C_e^R$  and  $C_e^A$  are the expected coverage such that  $C_e^R = C_e^A = \text{totCov}/2$ . Finally we enforced parents to have at least a  $k$ -mer coverage of 15 over the assembled region.

## System requirements and software availability

Scalpel is written in C++ and Perl. The source code is freely available as an open-source software project on the SourceForge website at <http://scalpel.sourceforge.net>. It usually takes 2-3 hours to process one exome-capture data set (80% of target at 20X) using 10 cores and requiring a minimum of 3GB of RAM.

## Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

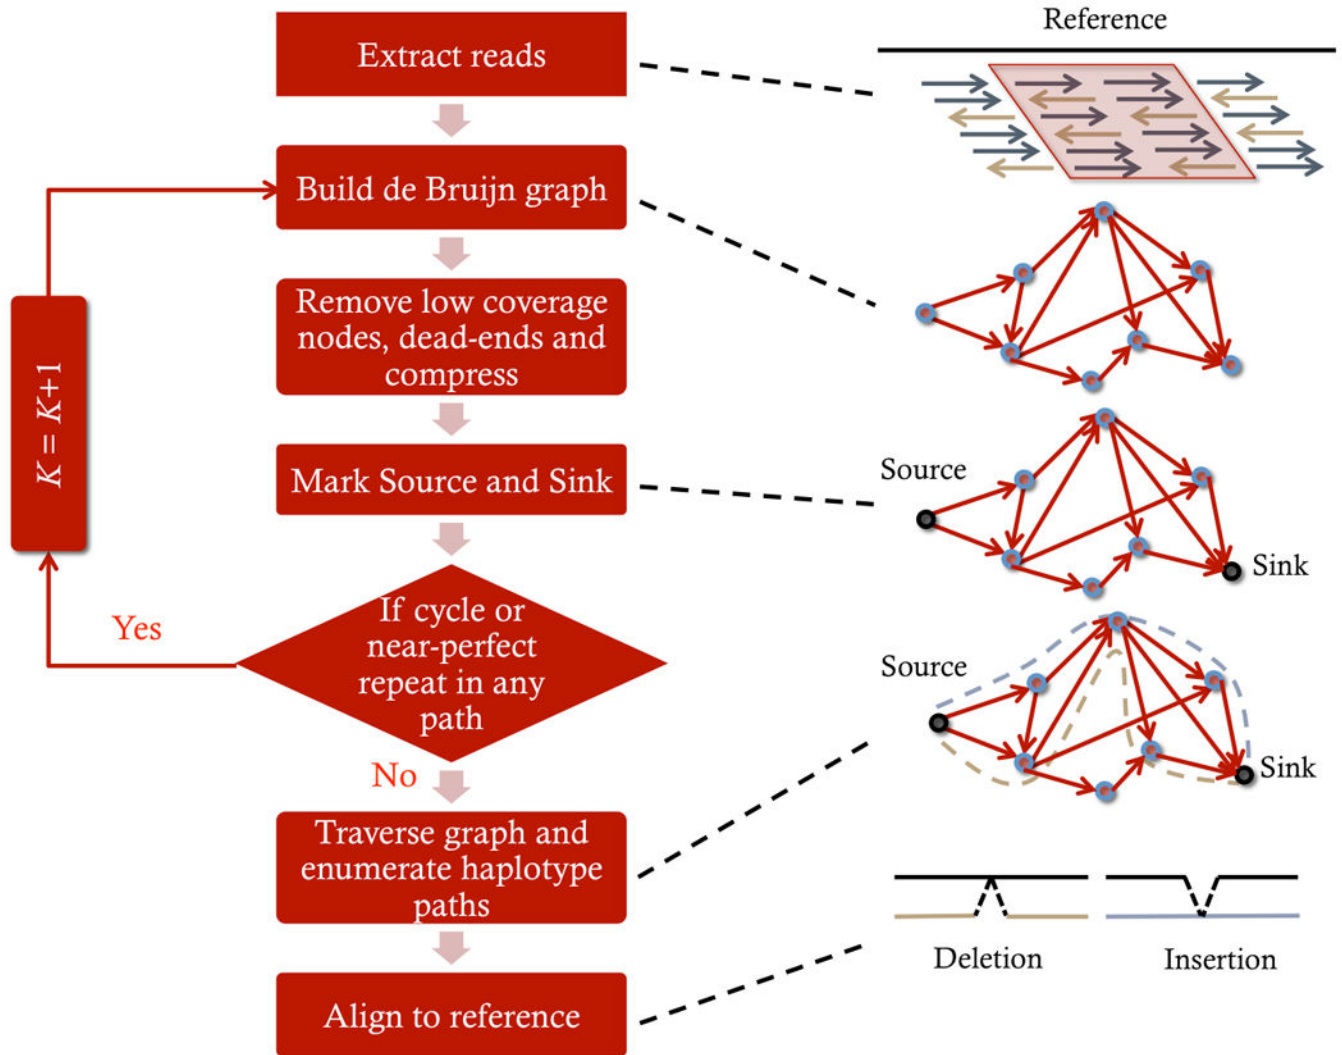
## Acknowledgments

The project was supported in part by National Institutes of Health award (R01-HG006677) and National Science Foundation award (DBI-1350041) to M.C.S., the CSHL Cancer Center Support Grant (5P30CA045508), the Stanley Institute for Cognitive Genomics, and the Simons Foundation (SF51 and SF235988) to M.W. The DNA samples used in this work are included within SSC release 13. Approved researchers can obtain the SSC population dataset described in this study by applying at <https://base.sfn.org>. We thank S. Eskipehliyan for the technical assistance with the MiSeq validation experiments. We thank M. Bekritsky, S. Neuburgerand, M. Ronemus, D. Levy, B. Yamron, and B. Mishra for helpful discussions and comments on the paper. We thank R. Aboukhalil for testing the software.

## References

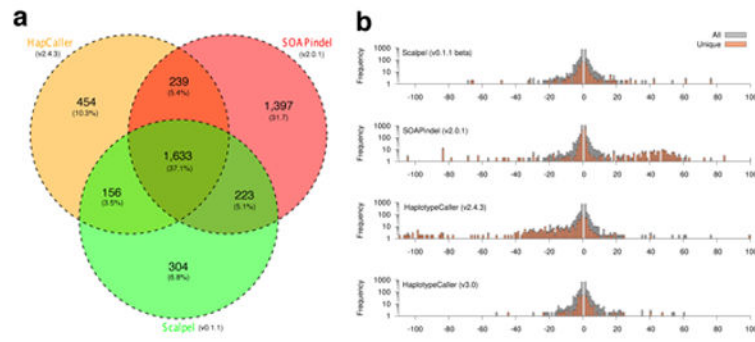
1. DePristo M, et al. *Nature Genetics*. 2011; 43:491–498. [PubMed: 21478889]
2. O'Rawe J, et al. *Genome Medicine*. 2013; 5:28. [PubMed: 23537139]
3. Zook JM, et al. *Nature Biotechnology*. 2014; 32:246–251.
4. Mullaney JM, Mills RE, Pittard WS, Devine SE. *Hum Mol Genet*. 2010; 19:131–6.
5. Pearson CE, Edamura NK, Cleary JD. *Nat Rev Gen*. 2005; 6:729–742.
6. Iossifov I, et al. *Neuron*. 2012; 74:285–299. [PubMed: 22542183]
7. Li H, Ruan J, Durbin R. *Genome Res*. 2008; 18:1851–1858. [PubMed: 18714091]
8. Montgomery SB, et al. *Genome Res*. 2013; 23:749–761. [PubMed: 23478400]
9. Albers CA, et al. *Genome Res*. 2011; 21:961–973. [PubMed: 20980555]
10. Ye K, Schulz MH, Long Q, Apweiler R, Ning Z. *Bioinformatics*. 2009; 25:2865–2871. [PubMed: 19561018]
11. Karakoc E, et al. *Nat Methods*. 2011; 9:176–8. [PubMed: 22179552]
12. Li Y, et al. *Nature Biotechnology*. 2011; 29:723–730.
13. Heng L. *Bioinformatics*. 2012; 28:1838–1844. [PubMed: 22569178]
14. Li S, et al. *Genome Res*. 2012; 23:195–200. [PubMed: 22972939]
15. Iqbal Z, Caccamo M, Turner I, Flicek P, McVean G. *Nat, Genet*. 2012; 44:226–232. [PubMed: 22231483]
16. Chen K, et al. *Genome Res*. 2014; 24:310–317. [PubMed: 24307552]
17. Li H, et al. *Bioinformatics*. 2009; 25:2078–9. [PubMed: 19505943]
18. Garrison E, Marth G. *arXiv:1207.3907v2* [q-bio GN].
19. Gymrek M, Golan D, Rosset S, Erlich Y. *Genome Research*. 2012; 22:1154–1162. [PubMed: 22522390]
20. Highnam G, et al. *Nucl Acids Res*. 2013; 41:e32. [PubMed: 23090981]
21. MacArthur DG, Chris, Tyler-Smith C. *Hum Mol Genet*. 2010; 19:R125–R130. [PubMed: 20805107]
22. Sjödin P, Bataillon T, Schierup MH. *PLoS ONE*. 2010; 5:e8650. [PubMed: 20098729]
23. Sanders SJ, et al. *Nature*. 2012; 485:237–241. [PubMed: 22495306]
24. O'Roak BJ, et al. *Nature*. 2012; 485:246–250. [PubMed: 22495309]

25. Neale BM, et al. *Nature*. 2012; 485:242–245. [PubMed: 22495311]
26. Darnell JC, et al. *Cell*. 2011; 146:247–261. [PubMed: 21784246]
27. Nagarajan N, Pop M. *Nature Reviews Genetics*. 2013; 14:157–167.
28. Li H, Durbin R. *Bioinformatics*. 2010; 26(5):589–595. [PubMed: 20080505]
29. Li H. arXiv:1303.3997v1 [q-bio GN]. 2013
30. Smith TF, Waterman MS. *Journal of Molecular Biology*. 1981; 147:195–197. [PubMed: 7265238]
31. Medvedev P, Georgiou K, Myers G, Brudno M. *Lecture Notes in Computer Science*. 2007; 4645:289–301.
32. Jackson, BG.; Aluru, S. *Parallel Processing*, 2008. ICPP '08. 37th International Conference on, 346-353; 2008;
33. Narzisi G, Mishra B. *Bioinformatics*. 2011; 27:153–160. [PubMed: 21088026]
34. Langmead B, Salzberg S. *Nature Methods*. 2012; 9:357–359. [PubMed: 22388286]
35. Fischbach GD, Lord C. *Neuron*. 2010; 68:192–195. [PubMed: 20955926]



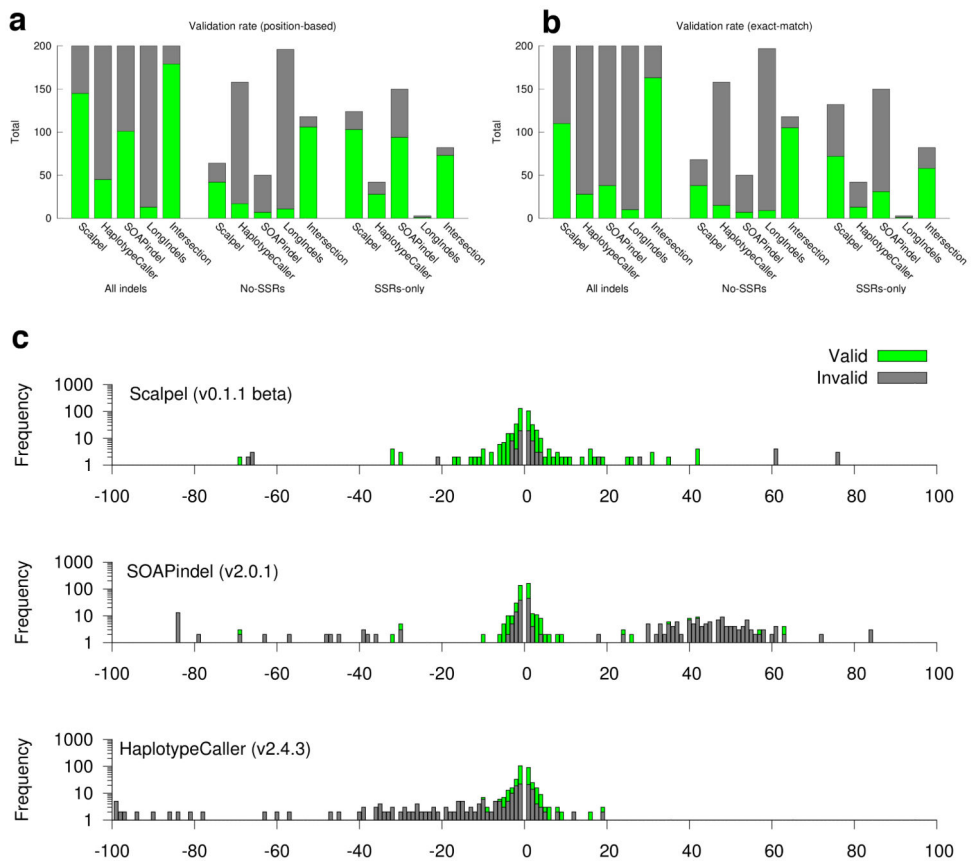
**Figure 1. Overview of the Scalpel algorithm work flow**

Extracted reads include: well-mapped reads, soft-clipped reads, and reads that fail to map, but are anchored by their mate. The assembled sequences are aligned to a reference using the standard Smith-Waterman-Gotoh alignment algorithm with affine gap penalties.



**Figure 2. Concordance of indels between pipelines**

(a) Venn Diagram showing the percentage of indels shared between the three pipelines. (b) Size distribution for indels called by each pipeline. 1,000 indels from five categories were analyzed by focused resequencing.



**Figure 3. MiSeq validation**

Ratio of valid (green) and false (grey) indel calls based on position-based matches (a) or exact matches (b) for the indicated tools, for indels of size  $\geq 30$ bp from the union of the mutations detected by all three pipelines (“Long Indels”), and for indels in the intersection (“Intersection”). Validation for all indels (“All indels”), validation only for indels within microsatellites (“SSRs-only”), and validation for indels that are not within microsatellites (“No-SSRs”). (c) Stacked histogram of validation rate by indel size for each variant caller.