RNAIndel: a machine-learning framework for discovery of somatic coding indels using tumor RNA-Seq data

Kohei Hagiwara[1], Liang Ding[1], Michael N. Edmonson[1], Stephen V. Rice[1], Scott Newman[1], Soheil Meshinchi[2], Rhonda E. Ries[2], Michael Rusch[1], Jinghui Zhang[1,*]

[1] Computational Biology, St Jude Children's Research Hospital, 262 Danny Thomas Place, Memphis, Tennessee, 38105, USA
[2] Clinical Research Division, Fred Hutchinson Cancer Research Center, 1101 Fairview Avenue North, Seattle, Washington, 98109, USA

* To whom correspondence should be addressed. Tel: +1(901) 595-7069; Fax: +1(901) 595-7100; Email: jinghui.zhang@stjude.org

ABSTRACT

Transcriptome sequencing (RNA-Seq) has been used for gene expression profiling and fusion detection but rarely for small insertion/deletion (indel) analysis due to the presence of artifacts generated during the PCR-based library preparation as well as alignment of spliced reads. Somatic indel calling is further challenged by the lack of matched normal RNA-Seq data. We present RNAIndel, a machine-learning based approach for classifying RNA-Seq indels into somatic, germline, and artifact by random forest models. RNAIndel was trained on tumor RNA-Seq of 330 pediatric cancer patients for whom whole exome and PCR-free whole-genome sequencing of paired tumor-normal DNA samples were also performed. Feature selection characterized somatic indels as those that were not explained by the strand-slippage model, a widely accepted hypothesis to explain indel generation during DNA replication. The method was tested on two independent RNA-Seq datasets with variable library protocols and RNA-Seq read lengths. Despite the heterogeneity in RNA-Seq data acquisition, RNAIndel robustly predicted 87–93% of somatic indels, recovering subclonal pathogenic indels that were missed by a 500× targeted sequencing of DNA samples. With RNAIndel, researchers can perform somatic indel calling in the transcriptome, expanding the utility of RNA-Seq and enhancing the interpretability of somatic indels.

INTRODUCTION

Transcriptome sequencing (RNA-Seq) is a versatile platform for performing a multitude of cancer genomic analyses such as gene expression profiling, allele specific expression, alternative splicing and fusion transcript detection. However, variant identification in RNA-Seq is not a common practice due to the presence of artifacts introduced in library preparation, the intrinsic complexity of splicing, and RNA editing (1). RNA-Seq data are predominantly generated from tumor-only samples as acquisition of a normal tissue with a comparable transcriptome is a rare practice. This lack of matching normal data further complicates somatic variant discovery in RNA-Seq. Despite these challenges, there are compelling reasons to explore RNA-Seq-based variant detection: 1) because RNA variants are expressed, they likely have a greater interpretability to the phenotype than DNA variants; 2) Since many tumor samples are analyzed by RNA-Seq alone, employing variant detection will make full use of existing resources. Thus, successful development of RNA-Seq variant calling tools will make this platform an interpretable and cost-effective alternative to DNA-based whole-genome or whole-exome sequencing (DNA-Seq), the current standard platform for variant detection.

Various single nucleotide variants (SNV) detection tools dedicated to RNA-Seq have been developed. SNPiR (1) calls true RNA-Seq SNV by hard-filtering calls in repetitive and low-quality regions, around splice sites, and at known RNA-editing sites. RVboost (2) is a machine-learning method to prioritize true SNVs trained on common SNPs in the input RNA-Seq data. eSNV-Detect (3) incorporates results generated from two mappers to confidently call expressed SNVs by removing mapping artifacts from individual mappers. Opossum (4) preprocesses RNA-Seq reads for SNV calling by splitting intron-spanning reads and removing spurious reads. The detection of indels in transcriptome is more challenging and has been largely unexplored (5). Even in DNA-Seq, indel discovery suffers from a high false discovery rate (6). In RNA-Seq, in addition to the artifacts from library preparation due to polymerase chain reaction (PCR), alignment of spliced reads containing an indel can introduce mapping artifacts. Indels are less common than SNVs in the genome with the ratio 1:7 and even more so in coding regions with 1:43 (7). This low prevalence is another reason for the difficult trade-off between sensitivity and false discovery rates in indel detection. In the context of cancer genomics, somatic indels should also be distinguished from germline indels. Somatic indel identification in tumor RNA-Seq can therefore be formulated as a three-class classification problem where somatic, germline, and artifact indels must be considered.

Here, we introduce RNAIndel, a novel computational method that takes a tumor RNA-Seq BAM file as input, calls and annotates coding indels, and classifies them into somatic, germline and artifacts by supervised learning. RNAIndel was developed by using 765,475 indels collected from 330 pediatric tumor transcriptomes. To demonstrate that the model is not sensitive to the particular error profiles in the training set, we tested RNAIndel against two RNA-Seq datasets which were generated by different library preparation protocols, read lengths, and sequencing platforms. Under these conditions, RNAIndel achieved 87–93% sensitivity with 7–10 somatic indels predicted from a total of 1,000–2,000 RNA-Seq indels per sample. RNAIndel is also flexibly designed to allow researchers to use the results from their own variant caller rather than the internally-implemented caller. By achieving high accuracy and sensitivity in somatic indel prediction, we expect RNAIndel will augment the range of RNA-Seq applications and facilitate the investigation of expressed somatic variants.

MATERIAL AND METHODS

Datasets

All RNA-Seq data sets in this study, which consisted of a training set and two test sets, were mapped by STAR (8) in 2-pass mode to GRCh38 (Supplementary Methods).

The training set was an internal dataset comprised of paired tumor (T) / normal (N) whole genome sequencing (WGS) and whole exome sequencing (WES), and tumor RNA-Seq generated from 330 pediatric cancer patients from 17 major cancer types (Table S1). Importantly, the WGS libraries were prepared by a PCR-free protocol to minimize PCR-artifacts. The details of nucleic acid extraction, library preparation, sequencing and variant detection were previously described (9).

Two public datasets were used as test sets. The first set (TestSet1) was comprised of 77 RNA-Seq samples of 20 tumor types (Table S2) previously used for developing a clinical pipeline (9) (St Jude Cloud data accession SJC-DS-1003 or EGA accession EGAS00001002217). Total RNA-Seq libraries were sequenced on Illumina HiSeq2000 or 2500 with 100-nt read length. The ground-truth somatic indels in TestSet1 were based on paired T/N DNA-Seq analysis performed on the samples (Table S3). The second test set (TestSet2) included 158 acute myeloid leukaemia samples by NCI TARGET project (https://ocg.cancer.gov/programs/target) (Table S4). The RNA-Seq data (dbGaP

study identifier phs000465) were generated at Canada's Michael Smith Genome Sciences Centre (BCCA-GSC) using strand-specific polyA+ RNA-Seq library preparation protocol and an Illumina HiSeq2000 sequencer with 75-nt read length. The ground-truth somatic indels in TestSet2 were compiled from the paired-T/N WGS analysis by the Complete Genomics, Inc. (CGI) Cancer Sequencing service pipeline version 2 (10) and targeted capture sequencing analysis (11) by Strelka (12) (Table S5).

Performance metrics

As performance metrics for overall classification are sensitive to class imbalance, we used class-specific metrics for evaluation. For class i, true positive (TP), true negative (TN), false positive (FP) and false negative (FN) are defined by:

$$TP = |predicted_class_i \cap true_class_i|$$

$$TN = |predicted_non_class_i \cap true_non_class_i|$$

$$FP = |predicted_class_i \cap true_non_class_i|$$

$$FN = |predicted_non_class_i \cap true_class_i|$$

-where $|\ \ |$ denotes the set size. True positive rate (TPR), false positive rate (FPR), and false discovery rate (FDR) are defined by TP/(TP+FN), FP/(FP+TN), and FP/(FP+TP), respectively.

A generalized F-score (F_β) is formulated as:

$$F_\beta = (1+\beta^2)\cdot TP/((1+\beta^2)\cdot TP + \beta^2 \cdot FN + FP)$$

-where β is a weighting parameter for TP. In this study, the parameter was set to $\beta = 10$ for somatic and 1 for germline and artifact prediction, respectively. The Hand-Till measure (M) is a generalization of the area under the curve (AUC) to multi-class classification (13). This quantifies the separation of classes in terms of prediction probability. For three classes i, j and k, the Hand-Till measure for class i is defined:

$$M = (p(i,j) + p(j,i) + p(i,k) + p(k,i))/4$$

-where $p(x, y)$ is the probability that a randomly chosen instance from class x has a higher prediction probability of being in class x than a randomly chosen class y instance's probability of being class x.

Cross validation and Feature selection

Models were evaluated in the actual imbalanced distribution by 5-fold cross-validation (5-fold CV). In the first fold, 20% of the data was held out unsampled as a validation set and the training set was downsampled from the remaining 80%. Trained models were evaluated using the unsampled validation set. This process was repeated to the fifth fold by rotating the validation set portion.

The optimal downsampling ratio was determined via a 5-fold CV for different artifact ratios from 1 to 20: $somatic: germline: artifact = 1: 1: x$, where $x \in [1,20]$. Each ratio was evaluated with a generalized F-score with β = 10 (F_{10}). We selected 1:1:18 for 1-nt indel and 1:1:6 for > 1-nt indel as a ratio achieving the highest value (Table S6).

Features were selected in a greedy best-first search by maximizing the true positive rate (TPR) for somatic indel. The search procedure began with an empty set of selected features (selected set) and the set of all features (feature set). In the first iteration, the model with each feature was separately evaluated in 5-fold CV. The feature that achieved the maximum TPR was added to the selected set and removed from the feature set. If features tied for the highest TPR, the feature with the minimum false discovery rate (FDR) was chosen. The next iteration examined the combination of the selected feature and one of the remaining features. The highest TPR feature was similarly added to the selected set and removed from the feature set. This procedure was continued while the model's TPR increased.

RNA-Seq variant calling by GATK

Preprocessing, variant calling and variant filtering of BAM files were performed by Haplotype Caller in the Genome Analysis Toolkit ver. 4.0.2.1. (GATK-HC) (14) The workflow closely followed the GATK Best Practice for calling variants in RNA-Seq (https://software.broadinstitute.org/gatk/documentation/article.php?id=3891). The Best Practice protocol filters spurious calls but does not distinguish somatic and germline calls. For a fair comparison with RNAIndel, which distinguishes somatic, germline and artifact calls, germline calls in the Best Practice protocol were filtered by matching the normal DNA-Seq data.

Indel recovery

RNAIndel uses actual indel alignments in the BAM file for feature calculation. For this calculation, indels reported from the caller are searched in the BAM file. RNAIndel first searches for indels equivalent to the reported indel and merges them (Fig. S1). If no equivalent indels are found, the nearest indel from the reported locus within a window of ±5-nt is used for analysis as proxy. These substituted cases are labelled as such in the output VCF file. Typically, 100% of indels reported from the built-in Bambino caller (15) are successfully recovered. When GATK-HC is used as an external caller, ~90% of indels are recovered; additional discussion can be found below (Results).

Software implementation

RNAIndel was implemented in Python 3 with pandas (https://pandas.pydata.org), numpy (http://www.numpy.org/), sklearn (16), pysam (https://github.com/pysam-developers/pysam) and pyvcf (https://github.com/jamescasbon/PyVCF). To facilitate the usage of the built-in Bambino variant caller which is implemented in Java, we packed Bambino into a Python package with optimizations for RNA-Seq variant calling. A Docker image with all the dependencies of Bambino and RNAIndel installed is available in Docker Hub (https://hub.docker.com/). We also developed workflows using Bash and Common Workflow Language (https://www.commonwl.org/) to bundle Bambino and RNAIndel.

RESULTS

Indel distribution in training data set

Each expressed coding indel in the training set was classified somatic, germline or artifact based on the paired T/N DNA-Seq by WGS and WES (Fig. 1A). Specifically, an indel in RNA-Seq was classified somatic if it matched to a somatic indel identified by the paired T/N analysis. Expressed germline indels were defined when they were supported by the germline WGS and WES. The remaining indels, present in RNA-Seq but absent in WGS and WES, were classified as artifacts. The resulting truth data, comprised of single nucleotide insertion/deletion (1-nt indel) and multi-nucleotide indels (>1-nt), showed distinct distributions in the three categories of somatic, germline, and artifact (Fig. 1B) where 1-nt indels are highly enriched in artifacts. Specifically, 1-nt indels accounted for 115, 12,529, and

616,121 of somatic, germline and artifacts, respectively while >1-nt indels accounted for 213, 45,098, and 91,399 of somatic, germline and artifacts, respectively.

Computational framework

The RNAIndel pipeline requires a RNA-Seq BAM file mapped by STAR as input. We specified STAR as the input-BAM aligner for its performance in RNA-Seq indel detection (5). Indel calling can be performed by the built-in Bambino caller or by an external caller by supplying calls in the Variant Call Format (VCF) (17). Indels are annotated using all RefSeq (18) isoforms containing coding exons; indels within a coding exon or in an intron region within 10 bases of a splice site (splice region) are considered coding indels and subjected to further analysis. For each coding indel, RNAIndel extracts reads covering the indel locus to retrieve the actual alignment pileup. This process also enables the incorporation of individual variations such as polymorphisms near the indel into the feature calculation. Matches to the dbSNP database (Build 151) (19) are also annotated as a feature; equivalent indels are also considered during matching (Fig. S1). For indels supported by ≥2 unique reads, prediction is made by classifiers specifically trained based on their size (i.e., 1-nt or >1-nt), which consist of an ensemble of random forest models. RNAIndel generates a VCF file where indel entries are parsimonious and left-aligned (21) to unify equivalent alignments of supporting reads (Fig. S1). Predicted class and probability are reported in the VCF INFO field as well as dbSNP membership, ClinVar annotation (22), variant effect annotation and calculated feature values.

Features

We extracted a total of 31 features from sequence/alignment, transcript/protein and database categories (Table 1 and Supplementary Method). Several features were selected based on the strand-slippage hypothesis, a widely accepted model for explaining the mechanism by which indels are generated in the process of DNA replication (23). Under this model, the DNA polymerase pauses synthesis in repetitive regions, and this delay in replication allows unpaired nucleotides to transiently anneal, leading to a misaligned replication. Thus, features governing sequence complexity (feature 1–3) and annealing temperature (feature 4–7) are expected to be important parameters of this model. By contrast, insertions or deletions that are dissimilar to the flanking sequences are unlikely to be caused by strand slippage (feature 8). Cancer-associated indels may be complex (Fig. S2) (24), so we define indel complexity based on misalignments near the indel site. Indel size (feature 10) is

negatively correlated to prevalence, and insertions (feature 11) are rarer than deletions in the human genome (25). Polymers of adenine or thymine are known to be more erroneous than guanine or cytosine polymers (feature 12–15) (6). Indels with sufficient read support are more likely to be true (feature 16–18). Mapping artifacts stem from ambiguous mapping (feature 19) or difficulty in mapping spliced reads (feature 20). Equivalent indels are alternative alignments of the identical indel sequence, a type of mapping artifact (feature 21). PCR-based genotyping can frequently create false multiallelic indels (feature 22) (26).

Indels were also characterized in terms of variant effect. In-frame indels are non-truncating unless they create a *de novo* stop codon. Indels in splicing regions may not affect splicing unless they destroy the splice motif (feature 23–25). Further, in-frame indels may be less deleterious if they occur outside of conserved domains (feature 26). The relative location of indels (feature 27) within proteins are bimodal around the N and C-termini. It has been postulated that transcripts with N-terminual indels can be rescued by an alternative downstream start codon, while a subset of C-terminus truncations may retain the all functional domains (7). Indels in the first and last exons are therefore known to be less sensitive to nonsense mediated decay (NMD) (feature 28) (27). We also hypothesized that the number of true indels in a gene is expected to be few. The number of detected indels in a gene were normalized by dividing by the coding sequence length (feature 29–30). Finally, indels present in dbSNP are considered more likely to be germline (feature 31).

Model development

Somatic indel discovery is the most likely application for RNAIndel, a context in which sensitivity is especially critical. For this reason, in model development, the true positive rate (TPR), or sensitivity, of somatic indels was favored over false discovery rate (FDR). As the training set revealed that artifacts were 3–4 magnitude higher than the true somatic indels (Fig. 1B), we determined the downsampling ratio for atifact with F_{10} as the performance metric to be maximized for the initial random forest models, which used a set of 27 features for 1-nt or >1-nt indel type (Methods and Table 1). The training sets were downsampled to somatic: germline: artifact = 1: 1: 18 for 1-nt with a F_{10} of 0.7 and 1: 1: 6 for >1-nt with a F_{10} of 0.79 (Table S6). The models were further refined by feature selection. Although the random forest algorithm can output the relative importance of features useful for feature selection, this importance is influenced by the classification performance for the dominant classes, i.e.

artifacts and germline indels. To overcome this issue, we selected features by a greedy best-first search for somatic TPR improvement (Methods), leading to a set of 11 features for 1-nt and 14 features for >1-nt indel type, respectively (Table 1). Performance metrics for each class are summarized (Table 2). Although optimized for somatic prediction, the final models accurately predicted germline and artifact indels as well.

The selected features provided distinct profiles of somatic, germline and artifact indels. Somatic indels were longer and dissimilar to the flanking sequences. Complex indels were rare but highly specific to the somatic class (Fig. 2). Somatic indels appeared at any variant allele frequency (VAF) below 0.5 (Fig. S3A). Germline indels were predominantly non-truncating or enriched for NMD-insensitivity if truncating, and were mostly found in the dbSNP Build 151 database (Fig. S3B). Artifacts were found in repetitive sequence with slightly lower nucleotide strength, a strong predictor of annealing temperature (28), and were predominantly in the lowest VAF range (Fig. S3A and C).

Somatic indel discovery from tumor RNA-Seq data

The training set was generated on an Illumina HiSeq 4000 using total RNA library with 125-nt read length. To test if the model was over-fitted to this particular specification, the developed model was evaluated using two unrelated datasets (TestSet1 and TestSet2) with varying library and sequencing specifications (Methods). Despite the technical differences between these datasets, RNAIndel was able to predict somatic indels at 87–93% TPR, performance comparable to the TPR obtained by the 5-fold CV of the training set (Table 2). Stringency was also reasonable: 3–6 indels were predicted to be somatic for 1-nt or >1-nt indel per sample (Table 3 and S2–S5). Detection of low-frequency indels (VAF<0.25) was also sensitive with a TPR of 82% (23/28) in TestSet1 and 88% (22/25) in TestSet2. Notably, in TestSet2, five previously unknown indels with low VAF in known cancer genes were predicted as somatic by RNAIndel, and ultimately validated by manual review of targeted capture exome sequencing (Table S7): *EP300* Y207fs (VAF: 0.1), *CEBPA* P23fs (0.17), *RAD21* D543fs (0.02), *KIT* Y418_D419>Y (0.15), and *CREBBP* S1767fs (0.01).

In addition to *de novo* indel analysis, RNAIndel also supports an optional user-provided custom indel panel, which is used to refine somatic prediction. Putative somatic indels matching this panel are assigned to germline or artifact, whichever has the higher probability (Fig. 3). This reclassification rule obviates the need to label non-somatic indels as germline or artifact, which can be

difficult in the absence of the DNA evidence. Such a non-somatic indel panel would ideally be compiled from a matching normal RNA-Seq data, which is rarely available; however, we found that germline and artifact indels misclassified as somatic frequently recur, which offers a practical strategy for compiling an exclusion panel if users have a dataset with DNA- and RNA-Seq performed. For example, we compiled an exclusion panel by collecting non-somatic indels in the training set that were misclassified as somatic in ≥3 samples in the 5-fold CV of the final model. After applying the exclusion panel to the predictions in the two independent test sets, the median number of putative indels were 3 (1-nt) and 2 (>1-nt) for TestSet1 and 1 (1-nt and >1-nt) for TestSet2.

Users can also prioritize the RNAIndel prediction for pathogenicity by uploading the output VCF file to the St Jude Cloud tool PeCan PIE (https://platform.stjude.cloud/tools/pecan_pie) (https://www.biorxiv.org/content/early/2018/06/06/340901), where the MedalCeremony algorithm (28) ranks variant pathogenicity into four tiers: gold, silver, bronze, and no medal. In the example workflow (Fig. 3), the *KIT* Y418_D419 deletion as somatic mutation was rated 'gold', the highest pathogenicity rank, while the other two indels did not receive a medal. Details of RNAIndel's performance on pathogenic indel prediction are presented below.

Working with an external variant caller— A GATK example

In addition to calling indels from a BAM file, RNAIndel can accept a VCF file containing calls made by an external variant caller (Fig. 1C). We chose GATK-HC to illustrate this feature for two reasons. First, GATK Best Practice of RNA-Seq variant calling has been documented in detail for GATK-HC with STAR (Methods). Second, the STAR/GATK-HC pipeline performed the best in a recent comparison study where combinations of a RNA-Seq mapper and a variant caller were tested for detecting known somatic *EGFR* indels in lung cancer (5). Following this procedure, we evaluated the performance of three approaches for detecting expressed pathogenic indels: RNAIndel with the Bambino caller, RNAIndel with GATK-HC, and the Best Practice-based approach without RNAIndel (Fig. 4A). We used TestSet1 which contained 21 pathogenic indels classified by our MedalCeremony algorithm (Table S8) and the RNA-Seq data were analyzed by the three approaches followed by the MedalCeremony classification. RNAIndel with the built-in caller achieved the highest sensitivity (20/21) with one misclassifed as artifact (Table S8). The prediction from the combination of RNAIndel and GATK-HC had the fewest false positives (a total of 7 artifacts were predicted), but this was

achieved at a cost of sensitivity: only 14 of the 21 pathogenic indels were detected. However, for true indels predicted to be somatic by this approach, the predictions coincided with those made by RNAIndel with Bambino (14 somatic and 3 germline), suggesting that the prediction was robust between the callers. The Best Practice prediction, i.e., the STAR/GATK-HC pipeline, was the noisiest with 442 artifacts, a rate 10 times higher than the first two approaches and also the lowest in sensitivity: only 13 of the 21 indels were correctly predicted with 1 filtered and 7 undetected (Fig. 4B, and Table S8). While the RNAIndel-based approaches predicted 3 germline indels as somatic, two were curated as cancer predisposition mutations on ClinVar (Table S9): *RAD50 K994_E995fs* (rs587780154) and *BRCA2 P1062_Q1063fs* (rs80359374).

Because GATK-HC locally assembles haplotypes while detecting variants, it may not report indels as aligned in the BAM file, and so these may not be recovered by RNAIndel's indel search (Methods). In TestSet1, nearly half the unrecovered indels (48.6%) were located in reads spanning intron-exon boundaries. Unrecovered exonic indels were most frequently located in human leukocyte antigen (HLA) loci (12.5%). For further investigation, 100 unrecovered indels were randomly selected and reviewed: 47 cases were in intron-exon boundaries, 14 at HLA loci and 39 non-HLA exon cases (Table S10). Remapping to the genome by the BLAT algorithm (30) confirmed that 44 of the 47 intron cases were mis-alignments over splice junctions. The remaining 3 cases were ambiguous, with low quality mappings. The high degree of genetic diversity at HLA loci likely explains the unsuccessful recovery. The 39 non-HLA exonic indels were also re-mapped by BLAT, revealing 17 mis-mappings. Of the remaining 22 exonic indels, 15 were found in repeats where somatic indels are less frequent as characterized in this study, 4 were in non-repetitive regions and 3 cases should have been realigned as a SNV or a cluster of SNVs rather than an indel. Considering that a majority of these indels were mismappings or in HLA loci, the impact of the underreporting may be marginal to the somatic prediction; indeed, all of the 21 somatic indels were successfully recovered in TestSet1.

DISCUSSION

We developed RNAIndel, a machine-learning based method to classify coding indel calls in RNA-Seq data. The performance of a machine learning method is constrained by the quality of the training set. In our initial assessment, nearly one third of RNA-Seq indels supported by WES were not validated by

further investigations, suggesting that they were PCR artifacts common to both platforms. To construct a high-quality training set, we employed three-platform clinical sequencing of WGS, WES and RNA-Seq with WGS data generated from a PCR-free library protocol (9). Previously, we have shown that this approach achieved a positive predictive rate of 98.8% and a sensitivity of 94.3% for exonic indels, assuring the quality of the training set (9). A possible limitation of the training set is that indels in the set were predominantly small indels with the maximum length being 23-nt. Analysis of larger indels may require tools designed to detect structural variations. This method was optimized and evaluated for somatic indel prediction using two unrelated test datasets generated under varying conditions. RNAIndel robustly predicted somatic indels at with a TPR of ~0.9. The ability to predict low-frequency variants is particularly important as tumor heterogeneity is the source of tumor clonal evolution causing resistance to cancer treatment. However, false negatives are frequent in these low VAF variants because artifacts are predominantly distributed in this VAF range (Fig. S3A). RNAIndel is able to discern subclonal indels with VAF <0.25 from artifacts at a TPR >0.8. This sensitivity enabled the discovery of novel somatic indels in TestSet2 which had not been identified by 500× targeted sequencing, possibly due to subclonality.

Features assigned to variants by the software are used to inform its classification logic and at times may unveil interesting biological insights. For example, the strand-slippage model predicts a simple indel with the flanking sequence. Two of our assigned features capture this state; the feature 'dissimilarity' is set to zero if the indel matches its flanking sequence, and the feature 'indel_complexity' is set to zero for simple indels. Interestingly, somatic indels frequently deviated from the slippage model, while germline and artifact indels were generally consistent with it (Fig. 2). This may be because the mechanism of somatic indel acquisition differs from that of germline due to the instability of cancer genome. Alternatively, one may speculate that a high proportion of indels compatible with the slippage model tend to be under weaker selection pressure and may hence appear in germline. For example, a triplet indel in a tandem tri-nucleotide repeat region is a pattern typical of germline indels and can be explained by strand slippage. This type of indel may have limited impact on protein function due to the possible redundancy of the repetitive amino acid molecules.

Like other variant discovery tools, RNAIndel does expect users to review the predicted outputs. In our review, single nucleotide indels (1-nt) with 'alt_count' being 2 or 3 appeared as a major

source of misclassification. As a strategy to review such indels, we recommend applying a panel of non-somatic indels (Fig. 3), gene prioritization (Fig. 3), using another indel caller (Fig. 4) to confirm consensus, visual inspection of the indel alignments, and reasoning by user's domain knowledge of the disease being investigated. In conclusion, RNAIndel is a sensitive, robust, and practical tool to predict somatic indels from tumor RNA-Seq alone.

AVAILABILITY

RNAIndel is freely available in St Jude GitHub repository (https://github.com/stjude/RNAIndel).

SUPPLEMENTARY DATA

Supplementary Figures S1–S3.

Supplementary Tables S1–S10 in SupplementaryData.xlsx

SupplementaryMethods.docx

FUNDING

This work was funded by the American Lebanese Syrian Associated Charities of St. Jude Children's Research Hospital. JZ and KH were also supported in part by grant P50GM115279-03 from National Institute of General Medical Sciences.

CONFLICT OF INTEREST

None declared.

REFERENCES

1. Piskol, R., Ramaswami, G., and Li, J.B. (2013) Reliable identification of genomic variants from RNA-Seq data. *Am J Hum Genet.*, 93, 641–651.

2. Wang., C., Davila, J.I., Baheti, S., Bhagwate, A.V., Wang, X., Kocher, J.P., Slager, S.L., Feldman, A.L., Novak, A.J., Cerhan, J.R., et al. (2014) RVboost: RNA-seq variant prioritization using a boosting method. *Bioinformatics*, 30, 3414–3416.

3. Tang, X., Baheti, S., Shameer, K., Thompson, K.J., Wills, Q., Niu, N., Holcomb, I.N., Boutet, S.C., Ramakrishnan, R., Kachergus, J.M., et al. (2016) The eSNV-detect: a computational system to identify expressed single nucleotide variants from transcriptome sequencing data. *Nucleic Acids Res*., 42, e172.

4. Oikkonen, L. and Lise, S. (2017) Making the most of RNA-seq: pre-processing sequencing data with Opossum for reliable SNP variant detection. *Wellcome Open Res*., 10.12688/wellcomeopenres.10501.2.

5. Sun, Z., Bhagwate, A., Prodduturi, N., Yang, P. and Kocher J.A. (2017) Indel detection from RNA-seq data: tool evaluation and strategies for accurate detection of actionable mutations. *Brief Bioinform*., 18, 973–983.

6. Fang, H., Wu, Y., Narzisi, G., O'Rawe, J.A., Barrón, L.T., Rosenbaum, J., Ronemus, M., Iossifov, I., Schatz, M.C. and Lyon, G.J. (2014) Reducing INDEL calling errors in whole genome and exome sequencing data. *Genome Med*., 10.1186/s13073-014-0089-z.

7. Ng, P.C., Levy, S., Huang, J., Stockwell, T.B., Walenz, B.P., Li, K., Axelrod, N., Busam, D.A., Strausberg, R.L. and Venter, J.C. (2008) Genetic variation in an individual human exome. *PLoS Genet*., 4, e1000160.

8. Dobin, A., Davis, C.A., Schlesinger, F., Drenkow, J., Zaleski, C., Jha, S., Batut, P., Chaisson, M. and Gingeras, T.R. (2013) STAR: ultrafast universal RNA-seq aligner. *Bioinformatics*, 29, 15–21.

9. Rusch, M., Nakitandwe, J., Shurtleff, S., Newman, S., Zhang, Z., Edmonson, M.N., Parker, M., Jiao, Y., Ma, X., Liu, Y., et al. (2018) Clinical cancer genomic profiling by three-platform sequencing of whole genome, whole exome and transcriptome. *Nat Commun*., 10.1038/s41467-018-06485-7.

10. Ma, X., Liu, Y., Liu, Y., Alexandrov, L.B., Edmonson, M.N., Gawad, C., Zhou, X., Li, Y., Rusch, M.C., Easton, J., et al. (2018) Pan-cancer genome and transcriptome analyses of 1,699 paediatric leukaemias and solid tumours. *Nature*, 555, 371–376.

11. Bolouri, H., Farrar, J.E., Triche, T Jr., Ries, R.E., Lim, E.L., Alonzo, T.A., Ma, Y., Moore, R., Mungall, A.J., Marra, M.A., et al. (2018) The molecular landscape of pediatric acute myeloid leukemia reveals recurrent structural alterations and age-specific mutational interactions. *Nat Med*., 1, 103–112.

12. Saunders, C.T., Wong, W.S., Swamy, S., Becq, J., Murray, L.J. and Cheetham, R.K. (2012) Strelka: accurate somatic small-variant calling from sequenced tumor-normal sample pairs. *Bioinformatics*, 28, 1811–1817.

13. Hand, D.J. and Till, R.J. (2001) A simple generalisation of the area under the ROC curve for multiple class classification problems. *Machine Learning*, 45, 171–186.

14. DePristo, M.A., Banks, E., Poplin, R., Garimella, K.V., Maguire, J.R., Hartl, C., Philippakis, A.A., del Angel, G., Rivas, M.A., Hanna, M. et al. (2011) A framework for variation discovery and genotyping using next-generation DNA sequencing data. *Nat Genet*., 10.1038/ng.806.

15. Edmonson, M.N., Zhang, J., Yan, C., Finney, R.P., Meerzaman, D.M. and Buetow, K.H. (2011) Bambino: a variant detector and alignment viewer for next-generation sequencing data in the SAM/BAM format. *Bioinformatics*, 27, 865–866.

16. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011) Scikit-learn: Machine Learning in Python. *JMLR*., 12, 2825-2830.

17. Danecek, P., Auton, A., Abecasis, G., Albers, C.A., Banks, E., DePristo, M.A., Handsaker, R.E., Lunter, G., Marth, G.T., Sherry, S.T., et al. (2011) The variant call format and VCFtools. *Bioinformatics*, 27, 2156–2158.

18. Pruitt, K.D., Tatusova, T. and Maglott, D.R. (2005) NCBI Reference Sequence (RefSeq): a curated non-redundant sequence database of genomes, transcripts and proteins. *Nucleic Acids Res*., 33, D501-504.

19. Sherry, S.T., Ward, M.H., Kholodov, M., Baker, J., Phan, L., Smigielski, E.M. and Sirotkin, K. (2001) dbSNP: the NCBI database of genetic variation. *Nucleic Acids Res*., 29, 308-311.

20. Breiman, L. (2001). Random Forests. *Machine Learning*, 45, 5–32.

21. Tan, A., Abecasis, G.R. and Kang, H.M. (2015) Unified representation of genetic variants. *Bioinformatics*, 31, 2202–2204.

22. Landrum, M.J., Lee, J.M., Riley, G.R., Jang, W., Rubinstein, W.S., Church, D.M. and Maglott, D.R. (2014) ClinVar: public archive of relationships among sequence variation and human phenotype. *Nucleic Acids Res*., 42, D980–985.

23. Garcia-Diaz, M. and Kunkel, T.A. (2006) Mechanism of a genetic glissando: structural biology of indel mutations. *Trends Biochem Sci*., 31, 206–214.

24. Ye, K., Wang, J., Jayasinghe, R., Lameijer, E.W., McMichael, J.F., Ning, J., McLellan, M.D., Xie, M., Cao, S., Yellapantula, V., et al. (2016) Systematic discovery of complex insertions and deletions in human cancers. *Nat Med*., 22, 97–104.

25. Zhang, Z. and Gerstein, M. (2003) Patterns of nucleotide substitution, insertion and deletion in the human genome inferred from pseudogenes. *Nucleic Acids Res*., 15, 5338–5348.

26. Weber, J.L., David, D., Heil, J., Fan, Y., Zhao, C. and Marth, G. (2002) Human diallelic insertion/deletion polymorphisms. *Am J Hum Genet*., 71, 854–862.

27. Popp, M.W. and Maquat, L.E. (2016) Leveraging Rules of Nonsense-Mediated mRNA Decay for Genome Engineering and Personalized Medicine. *Cell*, 165, 1319–1322.

28. Khandelwal, G. and Bhyravabhotla, J. (2010) A phenomenological model for predicting melting temperatures of DNA sequences. *PLoS One*, 10.1371/journal.pone.0012433.

29. Zhang, J., Walsh, M.F., Wu, G., Edmonson, M.N., Gruber, T.A., Easton, J., Hedges, D., Ma, X., Zhou, X., Yergeau, D.A., et al. (2015) Germline Mutations in Predisposition Genes in Pediatric Cancer. *N Engl J Med.*, 373, 2336–2346.

30. Kent, W.J. (2002) BLAT--the BLAST-like alignment tool. *Genome Res.*, 12, 656–664.

TABLE AND FIGURES LEGENDS

Table 1

			Initial model		Final model	
	Category	Feature	1-nt	>1-nt	1-nt	>1-nt
1	Sequence/Alignment	repeat	✓	✓	✓	
2		lc	✓	✓		
3		local_lc	✓	✓		
4		gc	✓	✓		
5		local_gc	✓	✓		
6		strength	✓	✓		
7		local_strength	✓	✓		✓
8		dissimilarity		✓		✓
9		indel_complexity	✓	✓	✓	✓
10		indel_size		✓		✓
11		is_ins	✓	✓		✓
12		is_at_ins	✓		✓	
13		is_at_del	✓		✓	
14		is_gc_ins	✓			
15		is_gc_del	✓			
16		ref_count	✓	✓	✓	✓
17		alt_count	✓	✓	✓	✓
18		is_bidirectional	✓	✓		✓
19		is_uniq_mapped	✓	✓	✓	✓
20		is_near_exon_boundary	✓	✓	✓	✓
21		equivalence_exists	✓	✓		
22		is_multiallelic	✓	✓		✓
23	Transcript/Protein	is_inframe		✓		
24		is_splice	✓	✓		
25		is_truncating	✓	✓		✓
26		is_in_cdd		✓		
27		indel_location	✓	✓		
28		is_nmd_insensitive	✓	✓	✓	
29		indels_per_gene	✓	✓	✓	✓
30		coding_seq_length	✓	✓		
31	DB	is_on_dbsnp	✓	✓	✓	✓

Table 1. Features used by RNAIndel falling into three categories. Features 1–22 characterize sequence and/or alignment properties. Features 23–30 are related to variant effects on gene products. Membership in the dbSNP database Build 151 (DB) is Feature 31. Features used before (Initial model) or after feature selection (Final model) are checked for single-nucleotide (1-nt) and multi-nucleotide indels (>1-nt). The definitions and computations of features are described in Supplementary Methods.

Table 2

Somatic	TPR	FPR	F_β	Hand-Till
1-nt	0.870	0.003	0.741	0.979
> 1-nt	0.897	0.015	0.818	0.982
Germline				
1-nt	0.952	0.001	0.954	0.985
> 1-nt	0.958	0.005	0.964	0.985
Artifact				
1-nt	0.996	0.031	0.997	0.993
> 1-nt	0.985	0.016	0.986	0.995

Table 2. Performance of the final model for 1-nt and >1-nt indel prediction. TPR (sensitivity): true positive rate; FPR (1-specificity): false positive rate; F_β: generalized F-score. β=10 for somatic, 1 otherwise. Hand-Till: Hand-Till's measure, a generalization of AUC to multi-class problem (Methods).

Table 3

	Library	Read Length	Sequencer	TPR		Median indels predicted somatic per sample	Median indels per sample
TestSet1 N = 77	Total RNA	100	HiSeq2000 or 2500	1-nt	0.882 (15/17)	6	1823
				>1-nt	0.868 (33/38)	3	326
TestSet2 N = 158	Poly A	75	HiSeq2000	1-nt	0.909 (20/22)	3	1036
				>1-nt	0.934 (57/61)	4	202

Table 3. Somatic prediction in test datasets with different conditions.

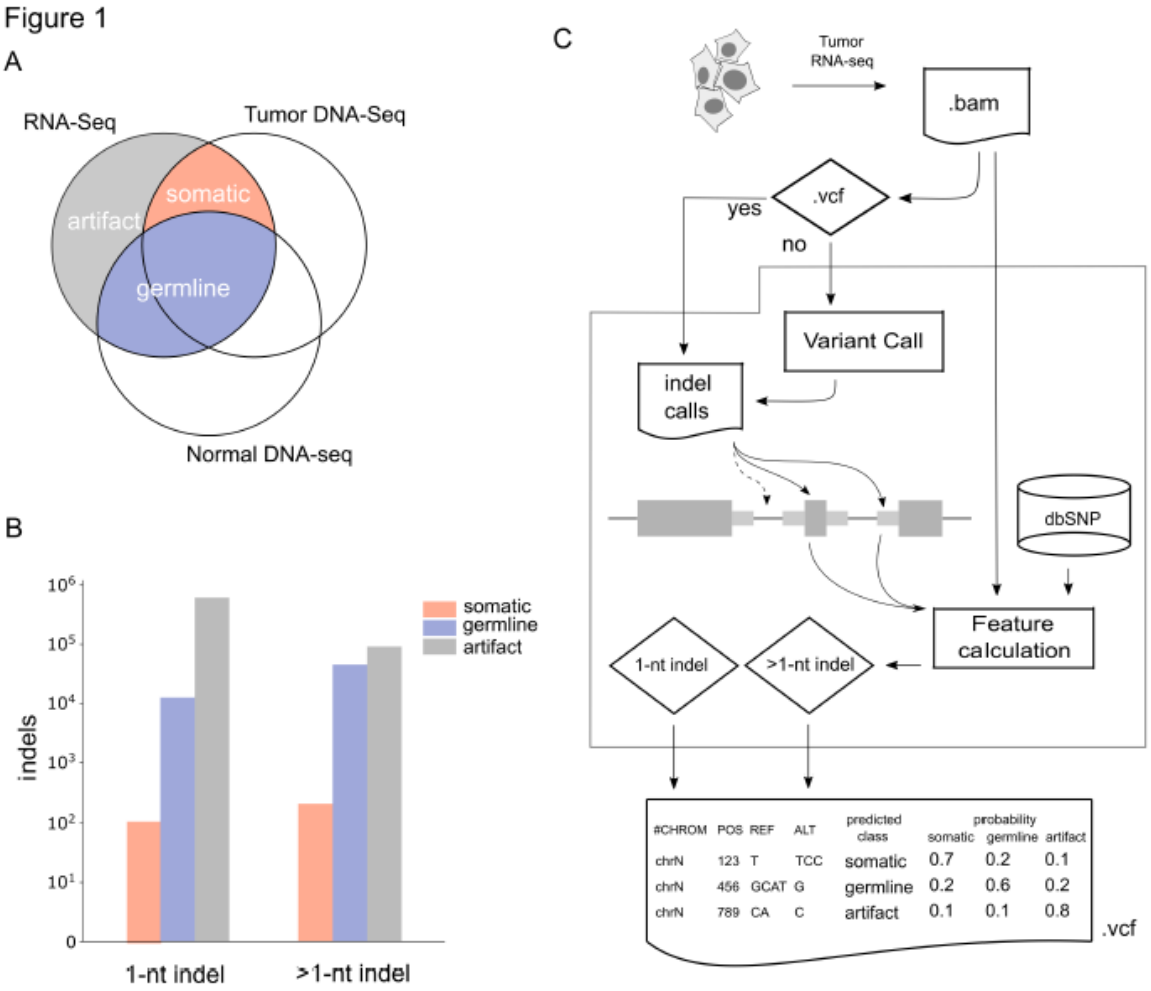

Figure 1. Computational framework and training dataset construction. **A**. Training set generated from 330 cases. Indel calls in RNA-Seq were classified somatic, germline and artifact by matching with tumor-normal paired WGS and WES (DNA-Seq) data: 328 somatic, 57,627 germline and 707,520 artifact indels. **B**. The 1-nt and >1-nt indel distribution in the categories of somatic, germline and artifacts. The class distribution of each dataset is shown in logarithm scale. **C**. Workflow of RNAIndel. A tumor RNA-Seq BAM file is a required input. If a VCF file from an external variant caller is supplied ('.vcf=yes'), indel calls in the files are used for prediction. Otherwise ('.vcf=no'), indel calling is performed on the input BAM file using Bambino. Indel calls are annotated for coding exon (grey thick box in the gene model) and splice region (light grey thin box), defined as an intronic region within 10 nt of the exon boundary. Indels annotated elsewhere will be removed from the analysis (dashed arrow). Features are calculated from the input BAM file and dbSNP database. Single nucleotide indels (1-nt) and multi-nucleotide indels (>1-nt) are predicted separately by an ensemble of random forest classifiers specifically trained for each type. Predicted class is based on the highest probability of being somatic, germine or artifact. The pipeline outputs a fully annotated VCF file.

Figure. 2

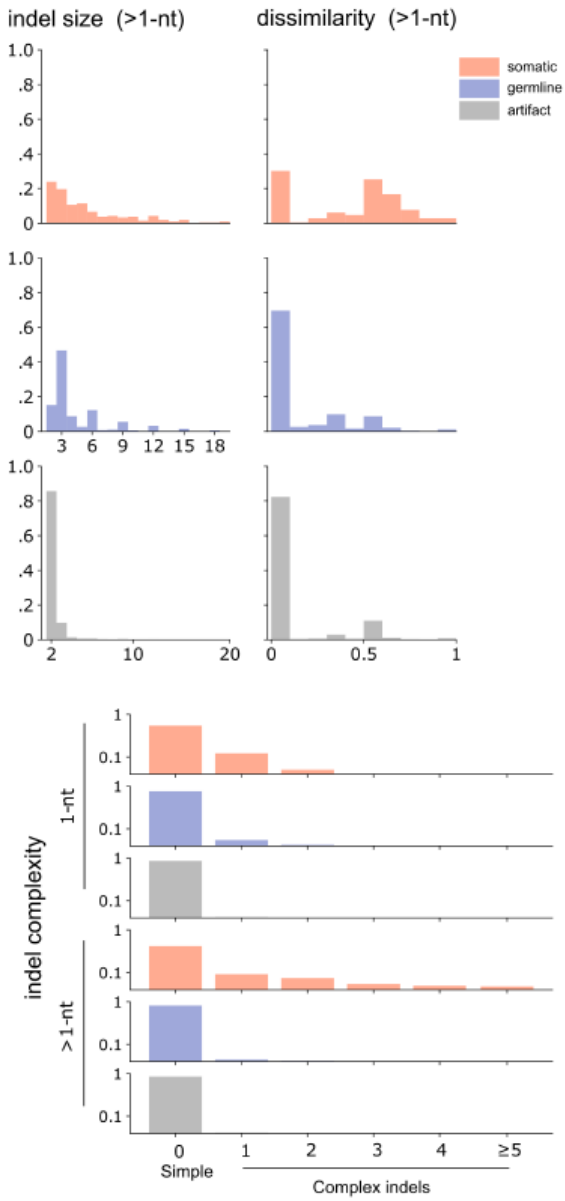

Figure 2. Somatic indel features. Features 'indel size' and 'dissimilarity' for >1-nt indel and 'indel complexity' for both 1-nt and >1-nt indels were selected in the final model. Somatic indels appeared longer and their lengths were frequently not divisible by three as opposed to germline indels. The inserted or deleted sequences in somatic indels were more dissimilar to the flanking sequences. Complex indels (Fig. S2) were almost exclusively found in the somatic class.

Figure 3

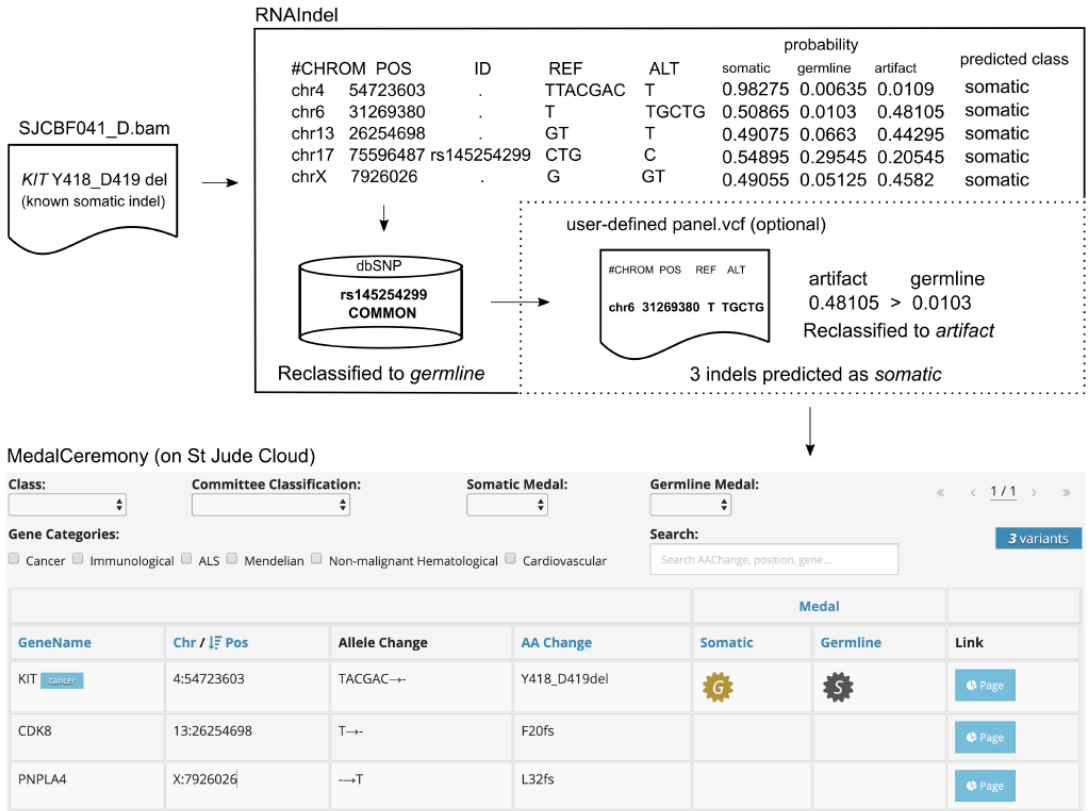

Figure 3. Pathogenic indel discovery from RNA-Seq. An example workflow is shown with a leukemia sample SJCBF04_D which harbors a single activating deletion in the *KIT* oncogene based on WGS and WES analysis. Before reclassification, five indels were predicted as somatic out of 2559 coding indels. The 2-nt deletion at chrX:7926026 matched rs14525299, a common indel in the dbSNP database, was reclassified as germline. Somatic prediction can be further refined by optionally supplying a user-defined exclusion panel (circumscribed in dashed line). Putative somatic indels in the panel are reclassified to germline or artifact class, whichever has the higher probability. Here, the indel on chromosome 6 matched a recurrent non-somatic indel in the panel prepared from the training set and was reclassified to artifact whose probability (0.48105) was higher than germline (0.0103). By uploading the output VCF file to the St Jude Cloud, the MedalCeremony algorithm prioritizes indels by pathogenicity. The *KIT* indel was identified as a somatic indel with the highest pathogenicity rank 'gold' (the 'G' symbol).

Figure 4

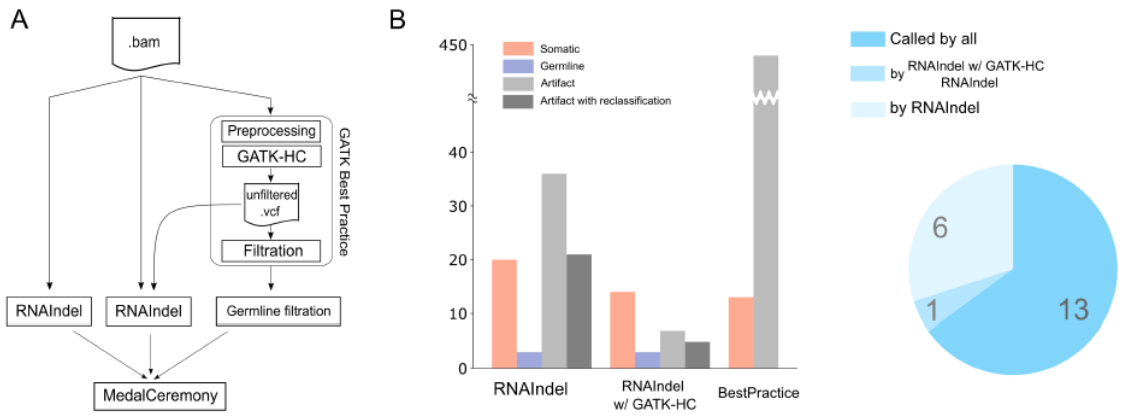

Figure 4. Workflow and performance using GATK as the external variant caller. RNAIndel is flexibly designed to work with an external variant caller via a user-provided VCF file. GATK-HaplotypeCaller (GATK-HC) was used as an example. **A**. TestSet1 (n=77) was screened for pathogenic somatic indels by three approaches. RNAIndel with built-in caller (left) used a RNA-Seq BAM file as input. RNAIndel with GATK-HC (middle) accepted the unfiltered VCF file from GATK-HC and the RNA-Seq BAM file as input. Best Practice-based approach (right) filters artifacts but does not discriminate somatic and germline calls. For a fair comparison, germline calls in this approach were removed by matching the normal WGS data (Germline filtration). Indels predicted as somatic by each method were prioritized by the MedalCeremony algorithm. Predicted somatic indels were defined pathogenic if prioritized as such by MedalCeremony. **B**. Predicted pathogenic indels by these approaches are shown in the bar plot. The number of artifacts is presented without (gray) and with reclassification by the panel of non-somatic indels prepared from the training set (darkgray). Of 21 known pathogenic indels in this dataset, 20 indels were predicted by these approaches. The overlap of the prediction is shown in the pie chart.

Figure S1

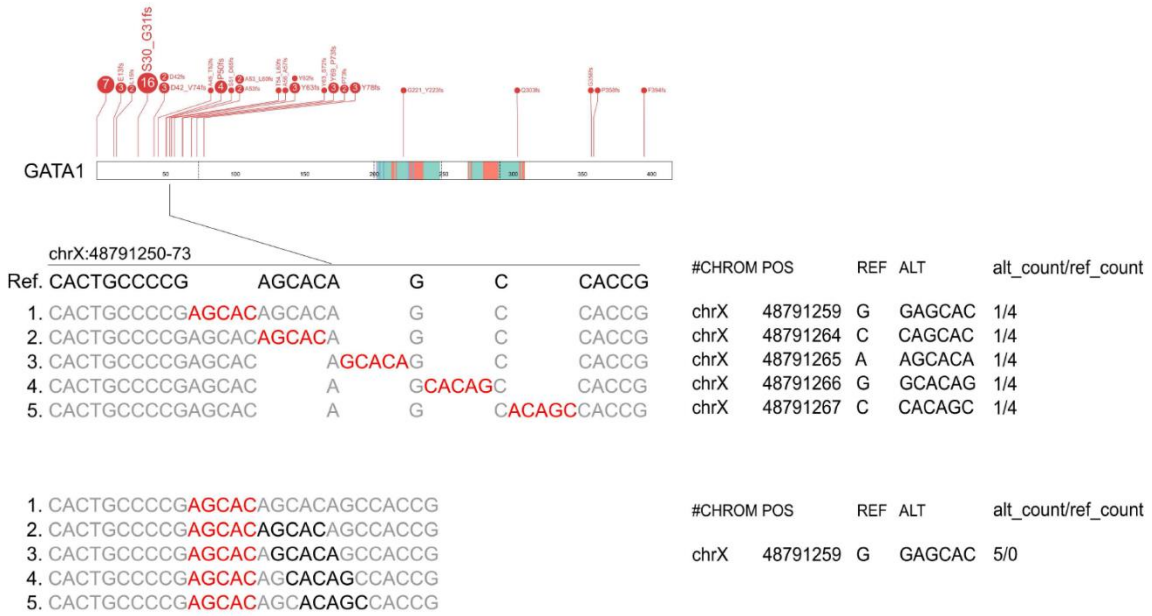

Figure S1. Equivalent indels. A 5-nt insertion in the *GATA1* gene that can be aligned to the reference in five different ways (*upper*). These alignments are left-aligned in red with the original inserted sequences highlighted in black (*lower*). Such alternative indel alignments are called equivalent and considered identical throughout the tool.

Figure S2

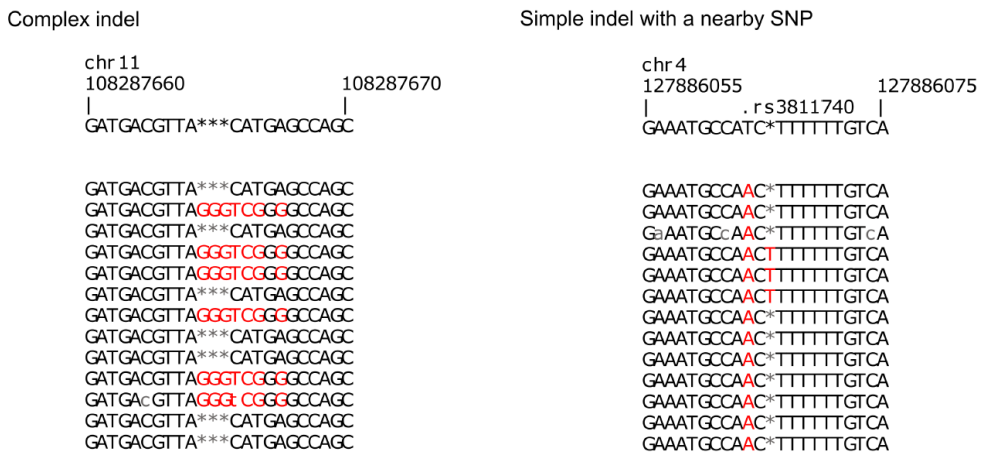

Figure S2. Complex indel. The inserted 'GGGT' is compounded with a cluster of misalignments (left). Note that the inserted sequence and additional misalignments are exclusively on the same reads. The inserted 'T' in the right is also associated with an additional misalignment (T>A), however this is a simple indel with a nearby single nucleotide polymorphism (rs3811740).

Figure.S3

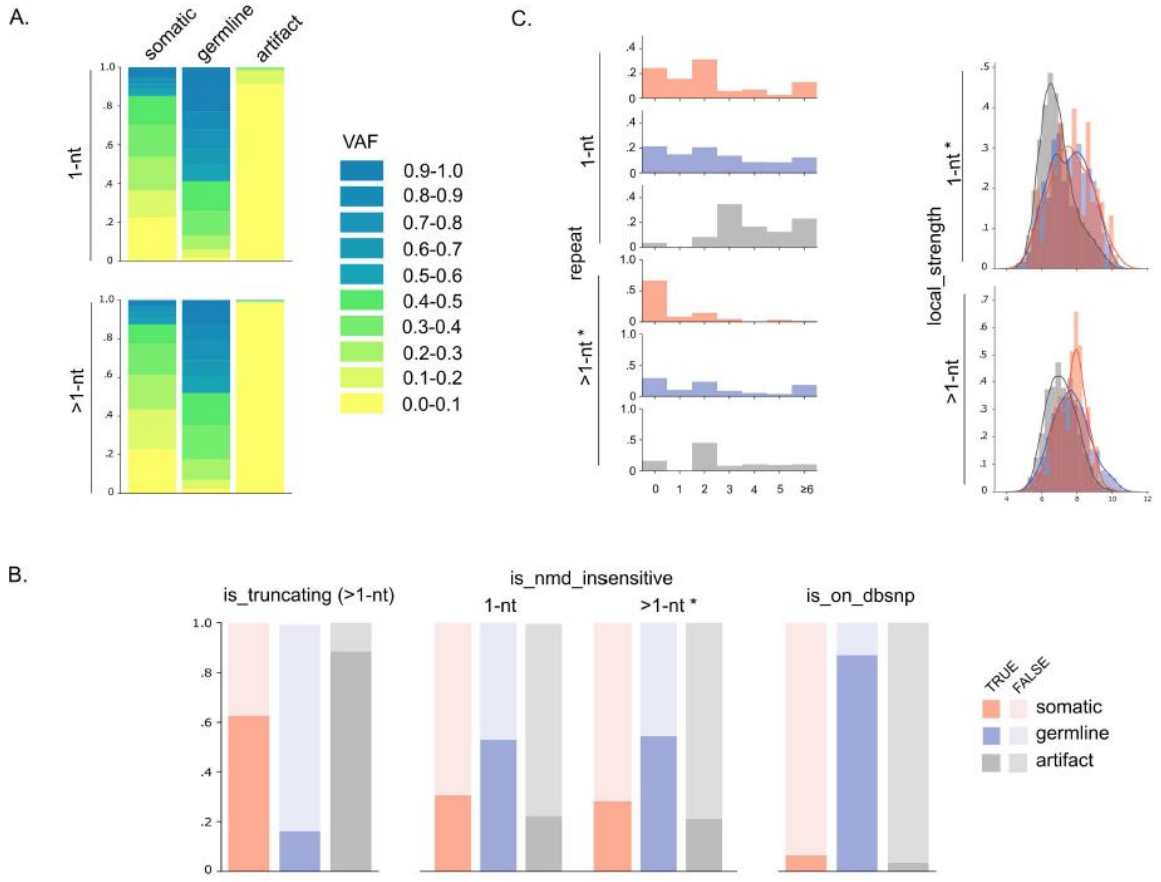

Figure S3. VAF spectrum, and germline and artifact features. **A**. The VAF spectrum of each class is visualized as a summary statistic of the features 'ref_count' and 'alt_count'. **B**. Visualization of features characteristic of germline indels. The distribution of truncating indels was visualized for the feature 'is_nmd_insensitive'. **C**. Visualization of features characteristic of artifacts. Asterisks indicate that the feature was shown for comparison, while not selected in the final model for the indel type (e.g., >1-nt * in 'is_nmd_insensitive').