Title

A Rigorous Interlaboratory Examination of the Need to Confirm NGS-Detected Variants with an Orthogonal Method in Clinical Genetic Testing

Authors

Stephen E. Lincoln [1]
Rebecca Truty [1]
Chiao-Feng Lin [2,3]
Justin M. Zook [4]
Joshua Paul [1]
Vincent H. Ramey [1]
Marc Salit [4,5]
Heidi L. Rehm [2,3,6,7,8]
Robert L. Nussbaum [1,9]
Matthew S. Lebo [2,3,7,8]

1. Invitae, San Francisco, CA
2. Partners HealthCare Laboratory for Molecular Medicine, Cambridge, MA
3. Brigham and Women's Hospital, Boston, MA
4. National Institute of Standards and Technology, Gaithersburg, MD
5. Joint Initiative for Metrology in Biology, Stanford, CA
6. Center for Genomic Medicine, Massachusetts General Hospital, Boston, MA
7. Harvard Medical School, Boston, MA
8. The Broad Institute of MIT and Harvard, Cambridge, MA
9. Department of Medicine, University of California San Francisco, San Francisco, CA

Corresponding Author

Stephen E. Lincoln
1400 16th Street, San Francisco, CA, 94103, USA
Mobile: 301-312-1725
email: steve.lincoln@me.com

Contents - This Document

Abstract
Main Text
- Introduction
- Methods
- Results
- Discussion
Acknowledgements
Disclosures
Oversight
Main Figure Legends
References
Main Figures 1–4
Main Table 1

Contents - Other Documents

Supplemental Text
Supplemental Figures 1–16
Supplemental Figure Legends
Supplemental Tables 1,2

Word counts:

Abstract: 217
Main text: 4355
- Introduction: 551
- Methods: 496
- Results: 1559
- Discussion: 1749

Abstract

Background: Orthogonal confirmation of NGS-detected germline variants has been standard practice, although published studies have concluded that confirmation of the highest quality calls may not always be necessary. The key question is how laboratories can establish criteria that consistently identify those NGS calls that require confirmation. Prior studies have limitations in addressing this: They are generally small, omit statistical justification, and explore limited aspects of the underlying data. The rigorous definition of criteria that separate high accuracy NGS calls from those that may or may not be correct remains a critical issue.

Methods: Five reference samples and over 80,000 patient specimens from two laboratories were analyzed. We examined quality metrics for almost 200,000 NGS calls with orthogonal data, including 1662 false-positives.

Results: A classification algorithm used these data to identify a battery of criteria that flag 100% of false positives as requiring confirmation (CI lower bound: 98.5–99.8% depending on variant type) while minimizing the number of flagged true positives. These criteria identify false positives that the previously published criteria miss. Sampling analysis showed that smaller datasets resulted in less effective criteria.

Discussion: Our methodology for determining test and laboratory-specific criteria can be generalized into a practical approach that many laboratories can use to help reduce the cost and time burden of confirmation without impacting clinical accuracy.

Introduction

The use of orthogonal assays (e.g., Sanger sequencing) to confirm variants identified by next-generation sequencing (NGS) is standard practice in many laboratories to reduce the risk of delivering false positive (FP) results. Clinical NGS tests can inform significant medical decisions [1,2], and thus confirmation is recommended by medical practice guidelines [3,4], although the details are generally left up to the laboratory [3,5]. Because clinical NGS methods often emphasize sensitivity (to avoid missing clinically important variants), FP rates can be elevated compared with those in research NGS [6]. Moreover, pathogenic variants are often technically challenging (e.g., many are located within repetitive sequences or segmental duplications), which can further increase FP rates [7–9]. Confirmation assays have a monetary cost, however, and can increase the time needed to deliver results, a critical factor in many clinical situations.

Published studies examining this issue have concluded that confirmation of the highest quality NGS calls may not always be necessary [10–13]. Some of these studies [10,12] propose specific criteria to separate high-confidence true-positive (TP) variant calls from those that are possibly FPs (Figure 1). These criteria differ from those used in filtering, the separate process of removing calls confidently believed to be false or unsupportable. The remaining intermediate-confidence calls are those that indeed benefit from confirmatory assays, manual data review, or both to determine which are TPs and which are FPs. Unfortunately, these prior studies are generally small, in some cases proposing criteria using only 1 or 5 example FPs (Table 1, Figure 1). The presence of few FPs may seem reassuring but leads to significant limitations in these studies. First, because their statistical power to characterize the FP population is limited, these studies do not address the question of whether future FPs are likely to resemble the few observed in the study. Quite possibly, additional FPs could be different and thus missed by the proposed criteria. Second, these studies use identical datasets for training and evaluating the proposed criteria, likely making the results subject to overfitting [14]. Few of these prior studies provide statistical justification for

their proposed criteria. Finally, these prior studies use data from individual laboratories and do not examine whether the methodologies can be generalized.

We set out to examine the role of confirmation using a much larger set of variant calls than has been published previously. We combined five reference samples characterized by the Genome in a Bottle Consortium (GIAB) [15–17] with confirmatory data from over 80,000 clinical tests. Our methodology was applied in two clinical laboratories that use similar but not identical NGS methods. Similar to prior studies, we found it possible to identify high-confidence NGS calls that do not benefit from confirmation. However, we found that a battery of criteria was necessary to capture all FPs, in contrast with the one or two metrics used in prior studies. Indeed we found that the specific criteria proposed by prior studies would miss FPs in our much larger dataset. We also found that observations of a variant as a TP can say little about its chance of being an FP in a different run, which indicates that historical performance can be an ineffective quality metric. Approaches such as ours can be used by many laboratories to provide efficient, effective, and statistically justified criteria for prioritizing variant calls for confirmation.

Methods

We compiled eight component datasets from two laboratories (Table 1) following the process illustrated in Figure 1. Five GIAB specimens were used: NA12878, NA24385, NA24143, NA24149, and NA24631 (Coriell Institute, Camden, NJ). NGS was performed using Illumina (San Diego, CA) 2x150 base pair (bp) reads as described previously [8,18,19]. Seven (lab 1) and three (lab 2) custom hybridization-based targeting assays were used, each targeting the exons of 100–1000 genes. Average coverage across targeted regions was 300–1000x or more depending on the assay and laboratory. Replicates of the GIAB samples were included.

Both laboratories used the Genome Analysis Toolkit (GATK) Haplotype Caller [20,21], although there were differences between the two bioinformatics pipelines. Variant filtering and manual review were performed according to the clinically validated criteria of each laboratory, and filtered or removed variant calls were excluded from this analysis. Copy number and structural variants were also excluded. Confirmation was performed using Sanger (Thermo Fisher, Waltham, MA) or Pacific Biosciences (Menlo Park, CA) sequencing. Multiple rounds of confirmation using different primers were performed as needed.

VCF (variant call format) files were annotated using the stratification BED (browser extensible data) files produced by the Global Alliance for Genomics and Health benchmarking workgroup [22] (Supplemental Table 1). These files delineate genomic regions by technical complexity (e.g., repetitive sequences). These regions were padded by 10 bp on each side to ensure that all affected positions were appropriately annotated. In the GIAB samples, we restricted analysis to the sample-specific high-confidence regions (GIAB-HC) [17]. In patient specimens, the union of the five GIAB-HC files was used to determine regions that were considered GIAB-HC—i.e., if a region was considered high confidence in any of the five GIAB reference datasets, it

was considered to be GIAB-HC in all patients. This approach prevented specific low-confidence variant calls in a particular GIAB specimen from inappropriately annotating that site as low confidence in general. Both laboratories defined clinical target regions to include, with some exceptions, protein-coding exons plus immediate flanking sequences (10–20 bp on each side).

GIAB reference calls version 3.3.2 (www.genomeinabottle.org) were used as confirmation data for variants in the GIAB specimens. Manual review was not performed for these calls, although analysis was restricted to the GIAB-HC regions. VCFeval [23] was used to compare our variant calls to the reference data and determine which calls were TPs and which were FPs. VCFeval can match variant calls even when the same diploid sequence is represented in different ways, an important factor when comparing indels and complex variants [22,23]. VCFeval also detects partial matches—i.e., zygosity errors (falsely calling a homozygous variant as heterozygous or vice versa) or sites at which one of two called alleles matches and one does not. These cases, collectively reported by VCFeval as "FP_CA," were considered FPs in our study. This study did not examine false negatives in detail.

Confidence intervals (CIs) were computed at 95% using the Jeffreys method. The Wilson score method was also used in Figure 2D.

Results

Our datasets included (a) almost 14,000 NGS variant calls subjected to confirmation during clinical testing, and (b) more than 184,000 calls with high-quality reference data from the GIAB samples (Table 1). We observed a total of 1662 FPs. We calculated the false discovery rate (FDR) of calls in each set, defined as the number of FPs divided by the number of FPs plus TPs. Note that FDR is one minus the analytic positive predictive value (aPPV), a metric recommended by the AMP/CAP guidelines, which consider aPPV preferable to specificity (i.e., the specific calculation) for describing multi-gene sequencing tests [24]. As expected, the FDR of indels was consistently higher than that of single-nucleotide variants (SNVs). In the clinical data, lab 2 applied less stringent filtering and manual review than did lab 1, which was reflected in higher FDRs. By contrast, FDRs in the GIAB data were roughly comparable between the two laboratories. Indel FDRs were particularly high in the GIAB data because (a) manual review was not performed, and (b) off-target regions containing low-coverage, repetitive, and complex sequences were included in the GIAB data. Such regions are present but less common in clinical targets.

Analysis of individual quality metrics from lab 1

We examined a variety of metrics to determine which were informative for identifying (or "flagging") FPs (Supplemental Figures 1–4). We considered a variety of thresholds for each metric and computed the fraction of FPs and TPs flagged at each threshold—these indicate the "value" and "cost," respectively, of each potential flag. We also computed FDR, which indicates how likely flagged variants are to be FPs.

We identified a number of informative metrics, although no single metric proved adequate alone. For example, consider the quality score (QUAL) for SNVs in the GIAB data (Figure 2, Supplemental Figure 1G). This same metric was used by Strom et al.

[10] to identify SNVs that require confirmation, although Strom had only a single FP to use in defining thresholds. In our dataset, no threshold value for QUAL could be chosen that flagged all, or even most, of the FPs without also flagging many of the TPs. Indeed, some FPs had quite high QUAL scores. A QUAL cutoff of 800, for example, flagged only 56% of FPs but also flagged 17% of TPs. These flagged variants had an FDR of 0.28%, indicating that this subset was not highly enriched for errors compared with the overall FDR of 0.15%. Indels fared similarly: 55% of FPs and 18% of TPs had QUAL<800, with an FDR of 16% compared with the baseline rate of 6.3% (Supplemental Figure 2G).

This observation was also true of our clinical dataset (Supplemental Figures 3G and 4G) in which 39% of FP SNVs and 2.9% of TPs had QUAL<800. These variants had a 3.6% FDR compared with 0.16% overall—20-fold enrichment—although most low-QUAL calls were still TPs. There is little precision in these measurements, however, because the clinical dataset contained only 10 FP SNVs, and only four had QUAL<800. Flagging all 10 FP SNVs required a QUAL threshold of 6400 (using logarithmic binning), which would flag almost 30% of TPs. Overfitting would certainly be an issue when analyzing the clinical dataset alone: If the single FP SNV with the highest QUAL score (5516) had been absent from these data, then thresholds would be driven by the next, much lower, score (2567). There is little statistical confidence in any such threshold: If 10 of 10 FPs are flagged, the point estimate of sensitivity for FPs would be 100%, but the lower bound of the CI is only 78%, demonstrating that thresholds determined using such a small dataset could miss many FPs.

We similarly analyzed other metrics, finding that allele balance was the most informative. Other informative metrics included read depth, strand bias, repetitive genomic context, the presence of nearby variant calls, and the type of the variant itself. All of these criteria had limitations, however. They either (a) were strongly indicative of

a call being an FP but flagged few such variants—stand bias was one example of this. More commonly, they (b) captured many TPs in addition to FPs yet still missed FPs at useful thresholds, similar to QUAL described above. More information on these findings is provided in the supplemental text.

Of the variants observed more than once, approximately half of those with at least one FP call were also correctly called as TP in different run(s) (Figure 3, Supplemental Figures 1–4S). This was true for both SNVs and indels. It was even sometimes true across replicates of the same sample, a situation made possible because partial matches were considered FPs in our analysis. Most examples were found in the GIAB specimens, although we observed examples in our smaller clinical datasets despite manual review having removed many FPs from those datasets. We calculated the observed FDR for each variant with multiple calls, where 100% indicates a fully systematic FP and 0% a consistent TP (Supplemental Figures 1–4T). Variants with FDRs between 1% and 40% were prevalent. Indeed, we suspect that many more such variants existed, as our power to detect low FDR variants was limited by the small number of GIAB specimens available.

Combining metrics

Because no single metric proved adequate, we investigated whether combining metrics might work better. One precedent for this approach is the Mu et al. study [12] which suggested that requiring depth>100 and allele balance between 40–60% would identify variants that do not require confirmation. In our much larger dataset, this is not the case: 29 FP indels and seven FP SNVs failed to be flagged as requiring confirmation using Mu's criteria. Nevertheless, we suspected that a larger battery of metrics might prove effective.

A heuristic algorithm was developed to explore this hypothesis. Briefly, this algorithm incrementally adds flags to a proposed set with the primary aim of capturing 100% of FPs using the combination of flags. This algorithm secondarily prioritizes minimizing the number of TPs also captured. Criteria were separately chosen for SNVs and indels and analysis was restricted to the GIAB-HC regions (see Discussion). To avoid overfitting, we ran the algorithm on hundreds of randomly chosen data subsets, in each case evaluating the chosen flags on the remaining data. This particular analysis suggested important refinements to our approach, which we implemented. It also demonstrated that smaller datasets would result in the selection of ineffective criteria, consistent with our statistical calculations. Further details are provided in the supplement.

The criteria chosen by this process are shown in Figure 4, Supplemental Figures 9–10, and Supplemental Table 2. The selected criteria included read depth, allele balance, quality scores, strand bias, nearby variants, and others. Combined, these criteria flagged 100% of all 201 FP SNVs (CI 98.9–100%) and 100% of all 987 FP indels (CI 99.8–100%) contained within the GIAB-HC regions. Only 5.8% of clinical TP SNVs and 6.8% of TP indels in GIAB-HC regions were flagged by the same criteria. Most FPs were captured by two or more flags, which we found reassuring. If we maximized such double-flagging, the fraction of TPs captured increased to 14.9% of SNVs and 18.1% of indels.

Application of these methods to data from lab 2

We examined whether a similar approach would work for the data from lab 2, even though these data were produced using somewhat different NGS methods and bioinformatics. As of this analysis, lab 2 had sequenced only one GIAB sample (NA12878) in replicate in addition to compiling clinical confirmation data.

We first analyzed data for individual metrics (Supplemental Figures 5–8) and subsequently combined metrics (Supplemental Figures 11, 12; Supplemental Table 2) using the approach detailed above. We found general consistency with the lab 1 criteria albeit with some differences. Using these criteria, 100% of FPs in the GIAB-HC regions were flagged, with 13.2% of TP SNVs and 15.4% of TP indels also flagged. CIs for the fraction of FPs flagged were 98.5–100% for SNVs and 99.1–100% for indels.

In comparison with lab 1, the criteria chosen for lab 2 appeared equally effective although somewhat less efficient – a greater fraction of TPs were flagged as requiring confirmation. We found multiple reasons for this, one of which was that fewer "bad" variant calls had been removed prior to confirmation compared with lab 1. This does not present an accuracy problem for the test but did cause our algorithm to broaden criteria as it strives to flag the full spectrum of FPs with which it is presented. Tightening up the lab 2 filtering thresholds, where possible, and increasing manual review would further reduce the confirmation workload. The CIs in this analysis would become wider as the number of FPs dropped, although these CIs would become more representative of the intermediate confidence variant population.

Another difference was that clinical FPs, rather than GIAB FPs, composed a larger fraction of the lab 2 dataset and played a substantial role in determining both criteria and CIs (Table 1). We do not consider this problematic, although we note that clinical data can have significant biases, such as over-representation of specific relatively frequent pathogenic variants. High diversity of variants in a study such as ours is certainly advantageous. The unique variant count for lab 2 will increase substantially, and the CIs will tighten when the analysis of the four additional GIAB samples is completed.

Discussion

Our study investigated whether a large and diverse dataset could be used to develop statistically robust criteria to guide the application of confirmation assays in clinical genetic testing. Our datasets combined clinical confirmation results with data from the sequencing of five GIAB specimens. Similar to prior studies [10–12], we found it possible to separate intermediate-confidence calls that require confirmation (i.e., to determine which are TPs and which are FPs) from calls that are high-confidence TPs which do not benefit from confirmation. We cannot guarantee that the criteria chosen by our method (or any such method) will capture all FPs. However, we can state that the probability of missing an FP is below a measurable level in datasets containing, collectively, almost 200,000 diverse variant calls with confirmatory data.

Our results differ from those of prior studies in important ways. Neither quality score (as suggested by Strom et al. [10]) nor the combination of allele balance and read depth (as suggested by Mu et al. [12]) efficiently captured all of the FPs in our datasets even if we re-optimized the thresholds. Although we acknowledge that criteria must be established independently for each NGS workflow—and indeed, workflows varied among these studies—we believe that this discrepancy results from the small datasets used in these prior studies, which may have left them underpowered and subject to overfitting (Table 1, Figure 2). A study by Baudhuin et al. [11] reported no FPs, a seemingly excellent result. However, our data (Supplemental Figures 1–8) and other studies [6,12] show that FPs are abundant at the threshold boundaries of many NGS quality indicators and that this marginal population contains a significant number of TPs as well. Thus, with current NGS methods, achieving high specificity with a single set of quality thresholds must sacrifice sensitivity. Indeed, a separate study showed that the specific methods used by Baudhuin had certain sensitivity limitations [7].

In our data, a battery of criteria based on read depth, strand bias, quality score, allele balance, genomic context, nearby calls, and other indicators was required to identify all FPs. This result is intuitive, given that a variety of different underlying factors can result in NGS false positives. It is reassuring that our list of metrics closely follows those recommended by the most recent AMP/CAP guidelines [24]. Indeed, many of these metrics are commonly used in variant filtering, albeit with different thresholds that emphasize sensitivity rather than specificity. In theory, a single quality score that captures all of these factors would be simpler to use than a battery of criteria. Unfortunately, we know of no current approach for computing such a score that will identify all FPs without capturing many TPs.

There were numerous benefits to using the large datasets we generated. These datasets not only provided a diverse set of FP calls for training (i.e., selecting criteria) but also allowed us to establish separate test sets to validate these criteria, identify outliers and minimize overfitting. Furthermore, we could compute statistical measures of confidence, a crucial part of any laboratory validation study, albeit one that is unfortunately not always used [24]. The statistical metric we used (CI lower bound given number of FPs) is rigorous but simple. More sophisticated approaches are feasible and ideal topics for future work. Prior studies generally report net PPV (without statistics) a metric that can be misleading because FPs near the edge of quality thresholds can be diluted by the large number of very high confidence TPs. Marginal error rates are far more informative and were extensively used in our study.

A key question is how large a dataset is adequate? Basic guidance can be found using the calculations described in recent guidelines [24,25]. For example, if certain criteria are shown to flag 59 of 59 FP variants in a study, then the point estimate is that the criteria are 100% effective. However, these criteria have only been rigorously demonstrated to capture ≥95% of FPs at $p = 0.05$ (i.e., 19 out of 20 times). Similar

values are shown in Figure 2. Our training/test set analyses demonstrated that this issue is not hypothetical: Small datasets indeed resulted in poorly performing criteria. Whether demonstrating 95% at p = 0.05 is adequate is not a statistical question but a medical policy question, and the answer could vary among clinical applications. Our laboratories aimed higher and achieved CI lower bounds between 98.5% and 99.8% (Table 1).

Obtaining a large number of FPs for study can be challenging, and we were careful not to do so in artificial and problematic ways. For example, additional FPs could be collected by lowering filtering thresholds. However, the resulting criteria might be good only at flagging known erroneous calls as opposed to defining the intermediate-confidence set (Figure 1) for which confirmation matters most. Another approach would be to run many samples containing the same FP variants, either replicates or different patients with systematic errors. In this case, the resulting criteria might be good only for limited classes of FPs. To avoid these problems, we used the same filtering thresholds that we use in clinical practice and we considered both the number of unique variants and the number of variant calls when designing our study (see Table 1). Running the set of five GIAB specimens greatly increased the diversity of variants with confirmatory data.

One might argue that our study artificially increased the number of FPs by including the off-target regions in the GIAB specimens. To the contrary, we consider it valuable to deliberately challenge our approach by including these regions. Doing so increased our ability to observe different types of FPs, many of which have features that are infrequent, yet certainly present, in coding exons (e.g., low depth, repeats, and complex variants). Because our algorithm required 100% of the combined FPs (clinical and GIAB, on and off target) to be flagged, we know that the criteria it selects will be effective on any subset of these FPs. We also ensured the validity of our results by

examining GIAB quality metrics in clinical target regions alone (supplement), observing results generally similar to those we observed in the full datasets, albeit with lower confidence and detail owing to the much smaller dataset sizes.

Our study design has important limitations. We used reference calls for confirmation in the GIAB specimens – these calls are highly accurate but imperfect [17]. For example, false negative variant clusters are known to exist in the GIAB reference data. These will incorrectly appear as FPs in our laboratory data. In addition, when the Coriell GIAB specimens are sequenced, there can be genetic differences compared to the original (NIST) DNA samples used to develop the GIAB reference data. In general, these two issues will cause our methodology to establish broader criteria than might otherwise be necessary. Such errors do not increase the risk of an FP being missed, but they do increase the number of flagged TPs. Given our study's objective, this is the "better" problem to have. In our study, outliers (FPs in the GIAB data with unusually high confidence in our data) were examined in detail to ensure that the results appeared correct. We did not confirm any sites in the GIAB specimens using an orthogonal assay, although studies following ours may find it valuable to do so.

A further limitation results from our focus on high-confidence regions in the GIAB reference call sets. We have little data regarding our ability to flag FPs outside of the GIAB-HC regions, which are heterogeneous and present diverse issues that can increase FP rates [15]. Thus, the criteria selected by our approach may inadequately capture FPs in these regions. We conclude that clinical variants outside of the GIAB-HC regions should always require confirmation unless additional, statistically valid studies are performed, likely on a variant-specific basis. In the lab 1 dataset, 7.0% of clinical SNVs and 15.4% of indels were outside of the GIAB-HC regions. In combination with the flagged variants in GIAB-HC regions, a total of 12.8% of SNVs and 22.2% of indels would thus require confirmation.

The ACMG guidelines for NGS recommend that laboratories have "extensive experience with NGS technology and be sufficiently aware of the pitfalls … before deciding that result confirmation with orthogonal technology can be eliminated" [3]. Our methodology provides a practical and rigorous way to follow this recommendation. Of course, changes in NGS processes, whether biochemical or bioinformatic, could change the criteria for selecting variants for confirmation. Per guidelines, the process described in this manuscript would need to be repeated as part of the validation of any significant process change.

Our results have specific implications for the current New York state guidelines, under which confirmation may be waived after "at least 10 positive samples per target gene" have been confirmed [4]. In our data, approximately half of all variants that appeared as FPs also appeared as TPs. These variants run a high risk of being confirmed as TPs in a series of tests and then, after confirmation is no longer considered necessary, being called falsely. Statistically, this is likely to happen with multi-gene panel, exome, or genome sequencing, which can detect and report many different variants. Moreover, in our datasets, different variants within a gene often exhibited remarkably different properties that correlated with remarkably different FDRs. Observing some of these variants as TPs provides little information about whether the other variants in that gene are FPs. In summary, our results argue against using the New York criteria for confirmation with multi-gene sequencing tests.

Our methodology does not address other roles that confirmation assays serve—e.g., verifying the identity of a specimen or determining the exact structure of indels and complex variants. Our framework also does not address the critical issue of how to handle conflicting results between NGS and confirmation assays. As Beck et al. [13] elegantly showed, naively assuming that a confirmation assay is always correct can

introduce more errors into NGS data than it fixes. Laboratories need to address these issues separately.

Our data show that confirmation assays can be limited to a small fraction of variants without any measurable effect on analytic specificity. They also show that an adequately large and diverse dataset is required to do this with confidence. Combining GIAB and clinical specimen data is a practical way to create a large dataset, and the analysis of such data is straightforward and highly informative. Limiting confirmation assays in this careful manner may help reduce costs and improve the turnaround time of clinical genetic tests without compromising quality.

Acknowledgements

We thank the Invitae and Partners HealthCare Laboratory for Molecular Medicine laboratory staff for producing the data used in this study. We thank all of the collaborators in the Genome in a Bottle Consortium for their excellent work on the reference specimens we used. We thank Nancy Jacoby for editorial help on this manuscript. All of the authors' work on this project was funded by their respective employers.

Disclosures

Authors CL, ML, SL, RN, JP, HR, VR, and RT work for laboratories that provide clinical genetic testing services. SL, RN, JP, VR, and RT own stock and/or stock options in Invitae. SL owns stock in Illumina and Thermo Fisher.

Certain commercial equipment, instruments, or materials are identified in this paper only to specify the experimental procedure adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the materials or equipment identified are necessarily the best available for the purpose.

Oversight

The use of de-identified clinical data in this study was subject to the requirements of the institutional review board of each laboratory. The GIAB specimens were used under the terms of their material transfer agreements.

Main Figure Legends

Figure 1: Study methodology.

(A) Variant calls can be separated into high-confidence true positive (green) and intermediate-confidence (yellow) categories using strict thresholds intended to maximize the specificity of the high confidence set. The intermediate set will contain a mixture of true and false positive calls. Variant calls that are confidently false positives (red) are typically filtered out using different criteria that emphasize sensitivity. The analogy to a traffic light is illustrative.

(B) Process for collecting true positive and false positive calls for both clinical and Genome in a Bottle specimens used in our study.

Figure 2: Quality score data and statistics.

(A) Variant call quality scores (QUAL, y-axis) for true positive (TP, blue) and false positive (FP, red) single-nucleotide variant (SNV) calls in the lab 1 Genome in a Bottle (GIAB) dataset. The x-axis position of each point is randomly assigned. To make the point density changes visible, we plotted a random selection of one-fifth of the TP calls along with all FPs. Corresponding histograms using the full dataset and no downsampling are shown in Supplemental Figure 1G.

(B) Random sample of 1000 data points from the lab 1 GIAB SNV data. All points are displayed. The arrows indicate the two FPs present. A thousand such random samples were generated, and many would lead to very different conclusions about the minimal effective QUAL thresholds (dashed line) compared to analyses of the full dataset.

(C) Full dataset from Strom et al. Note that quality scores are not precisely calibrated between these data and plots (A) and (B), although no Strom-like threshold would prove effective at separating TPs and FPs in plot (A).

(D) Lower bound of the 95% confidence interval (CI) on the fraction of FPs captured (y-axis) as a function of number of FPs used to determine criteria (x-axis) assuming 100% success is observed. This calculation used the Jeffreys method (blue) and the Wilson score method (red), which both produce similar but not identical results to those produced with the tolerance calculation used in the AMP/CAP guidelines [24] (this method produces a value between that of Wilson and Jeffreys). Specific points highlighted show 95%, 98%, and 99% of FPs captured. These levels require at least 49, 124, and 249 FPs, respectively, using the Jeffreys method. Many prior studies did not achieve such significance (Table 1).

Figure 3: Historical performance as a quality metric.

These two plots examine all variants (i) that were observed more than once in our lab 1 dataset, and (ii) for which one or more of those calls was a false positive (FP). SNVs and indels are combined. Each unique variant (for example, a particular base change) is counted once. A detailed breakdown of these data and the corresponding lab 2 results are in the supplement.

(A) Approximately half of all variants that were falsely called were also correctly called as positives in different specimens or runs.

(B) Many of these variants have low false discovery rates (FDRs). Calls of a low-FDR variant will usually but not always be correct, and repeated TP observations of such a variant provide little information about the accuracy of any following observation of that

variant. Our study is underpowered to measure FDRs near 0% or near 100% owing to the small number of GIAB specimens, and many more such variants may exist than are shown here. We conclude that historical performance is a ineffective quality metric in our datasets.

Figure 4: Combining flags.

These plots show the cumulative effect (top to bottom) of sequentially combining flags chosen by the heuristic algorithm. Plot (A) shows single-nucleotide variants (SNVs) and plot (B) shows indels in the lab 1 dataset. Red bars show the fraction of all false positives (FPs) captured and blue bars the fraction of clinical true positives (TPs). Grey bars show the fraction of FPs captured by two or more flags. The dotted lines illustrate the flags needed to capture 100% of the FPs. Details and lab 2 results can be found in the supplement.

The flags include the following:
DP: read depth, specifically GT_DP
DPDP: ratio of GT_DP to INFO_DP
SB5, SB25: strand bias metrics
QUAL: quality score
QD: quality-depth score
ABnorm: Allele balance for heterozygotes, normalized to be within 0.0–0.5
ABhom: Allele balance for homozygotes
HetAltorUnk: Het-alt call or 3+ alleles called in sample
Repeat: variant call within a homopolymer or short tandem repeat

References

1. Korf BR, Rehm HL. New approaches to molecular diagnosis. JAMA. 2013;309: 1511–1521.

2. Rehm HL. Disease-targeted sequencing: a cornerstone in the clinic. Nat Rev Genet. 2013;14: 295–300.

3. Rehm HL, Bale SJ, Bayrak-Toydemir P, Berg JS, Brown KK, Deignan JL, et al. ACMG clinical laboratory standards for next-generation sequencing. Genet Med. 2013;15: 733–747.

4. New York State Department of Health. Guidelines for Validation Submissions of Next Generation Sequencing (NGS) assays under the NYS Testing Category of Genetic Testing [Internet]. 2015. Available: health.ny.gov

5. Aziz N, Zhao Q, Bry L, Driscoll DK, Funke B, Gibson JS, et al. College of American Pathologists' laboratory standards for next-generation sequencing clinical tests. Arch Pathol Lab Med. 2015;139: 481–493.

6. DePristo MA, Banks E, Poplin R, Garimella KV, Maguire JR, Hartl C, et al. A framework for variation discovery and genotyping using next-generation DNA sequencing data. Nat Genet. 2011;43: 491–498.

7. Lincoln SE, Zook JM, Chowdhury S, Mahamdallie S, Fellowes A, Klee EW, et al. An interlaboratory study of complex variant detection [Internet]. bioRxiv. 2017. p. 218529. doi:10.1101/218529

8. Lincoln SE, Kobayashi Y, Anderson MJ, Yang S, Desmond AJ, Mills MA, et al. A Systematic Comparison of Traditional and Multigene Panel Testing for Hereditary Breast and Ovarian Cancer Genes in More Than 1000 Patients. J Mol Diagn. 2015;17: 533–544.

9. Mandelker D, Schmidt RJ, Ankala A, McDonald Gibson K, Bowser M, Sharma H, et al. Navigating highly homologous genes in a molecular diagnostic setting: a resource for clinical next-generation sequencing. Genet Med. 2016;18: 1282–1289.

10. Strom SP, Lee H, Das K, Vilain E, Nelson SF, Grody WW, et al. Assessing the necessity of confirmatory testing for exome-sequencing results in a clinical molecular diagnostic laboratory. Genet Med. 2014;16: 510–515.

11. Baudhuin LM, Lagerstedt SA, Klee EW, Fadra N, Oglesbee D, Ferber MJ. Confirming Variants in Next-Generation Sequencing Panel Testing by Sanger Sequencing. J Mol Diagn. 2015;17: 456–461.

12. Mu W, Lu H-M, Chen J, Li S, Elliott AM. Sanger Confirmation Is Required to Achieve Optimal Sensitivity and Specificity in Next-Generation Sequencing Panel Testing. J Mol Diagn. 2016;18: 923–932.

13. Beck TF, Mullikin JC, NISC Comparative Sequencing Program, Biesecker LG. Systematic Evaluation of Sanger Validation of Next-Generation Sequencing Variants. Clin Chem. 2016;62: 647–654.

14. Alpaydin E. Introduction to Machine Learning. MIT Press; 2014.

15. Zook JM, Chapman B, Wang J, Mittelman D, Hofmann O, Hide W, et al. Integrating human sequence data sets provides a resource of benchmark SNP and indel genotype calls. Nat Biotechnol. 2014;32: 246–251.

16. Zook JM, Catoe D, McDaniel J, Vang L, Spies N, Sidow A, et al. Extensive sequencing of seven human genomes to characterize benchmark reference materials. Sci Data. 2016;3: 160025.

17. Zook J, McDaniel J, Parikh H, Heaton H, Irvine SA, Trigg L, et al. Reproducible integration of multiple sequencing datasets to form high-confidence SNP, indel, and reference calls for five human genome reference materials [Internet]. bioRxiv. 2018. p. 281006. doi:10.1101/281006

18. Pugh TJ, Kelly MA, Gowrisankar S, Hynes E, Seidman MA, Baxter SM, et al. The landscape of genetic variation in dilated cardiomyopathy as surveyed by clinical DNA sequencing. Genet Med. 2014;16: 601–608.

19. Alfares AA, Kelly MA, McDermott G, Funke BH, Lebo MS, Baxter SB, et al. Results of clinical genetic testing of 2,912 probands with hypertrophic cardiomyopathy: expanded panels offer limited additional sensitivity. Genet Med. 2015;17: 880–888.

20. Poplin R, Ruano-Rubio V, DePristo MA, Fennell TJ, Carneiro MO, Van der Auwera GA, et al. Scaling accurate genetic variant discovery to tens of thousands of samples [Internet]. bioRxiv. 2017. p. 201178. doi:10.1101/201178

21. McKenna A, Hanna M, Banks E, Sivachenko A, Cibulskis K, Kernytsky A, et al. The Genome Analysis Toolkit: a MapReduce framework for analyzing next-generation DNA sequencing data. Genome Res. 2010;20: 1297–1303.

22. Krusche P, Trigg L, Boutros PC, Mason CE, De La Vega FM, Moore BL, et al. Best Practices for Benchmarking Germline Small Variant Calls in Human Genomes [Internet]. bioRxiv. 2018. p. 270157. doi:10.1101/270157

23. Cleary JG, Braithwaite R, Gaastra K, Hilbush BS, Inglis S, Irvine SA, et al. Comparing Variant Call Files for Performance Benchmarking of Next-Generation Sequencing Variant Calling Pipelines [Internet]. bioRxiv. 2015. p. 023754. doi:10.1101/023754

24. Roy S, Coldren C, Karunamurthy A, Kip NS, Klee EW, Lincoln SE, et al. Standards and Guidelines for Validating Next-Generation Sequencing Bioinformatics Pipelines: A Joint Recommendation of the Association for Molecular Pathology and the College of American Pathologists. J Mol Diagn. 2018;20: 4–27.

25. Jennings LJ, Arcila ME, Corless C, Kamel-Reid S, Lubin IM, Pfeifer J, et al. Guidelines for Validation of Next-Generation Sequencing-Based Oncology Panels: A Joint Consensus Recommendation of the Association for Molecular Pathology and College of American Pathologists. J Mol Diagn. 2017;19: 341–365.

Fig 1A

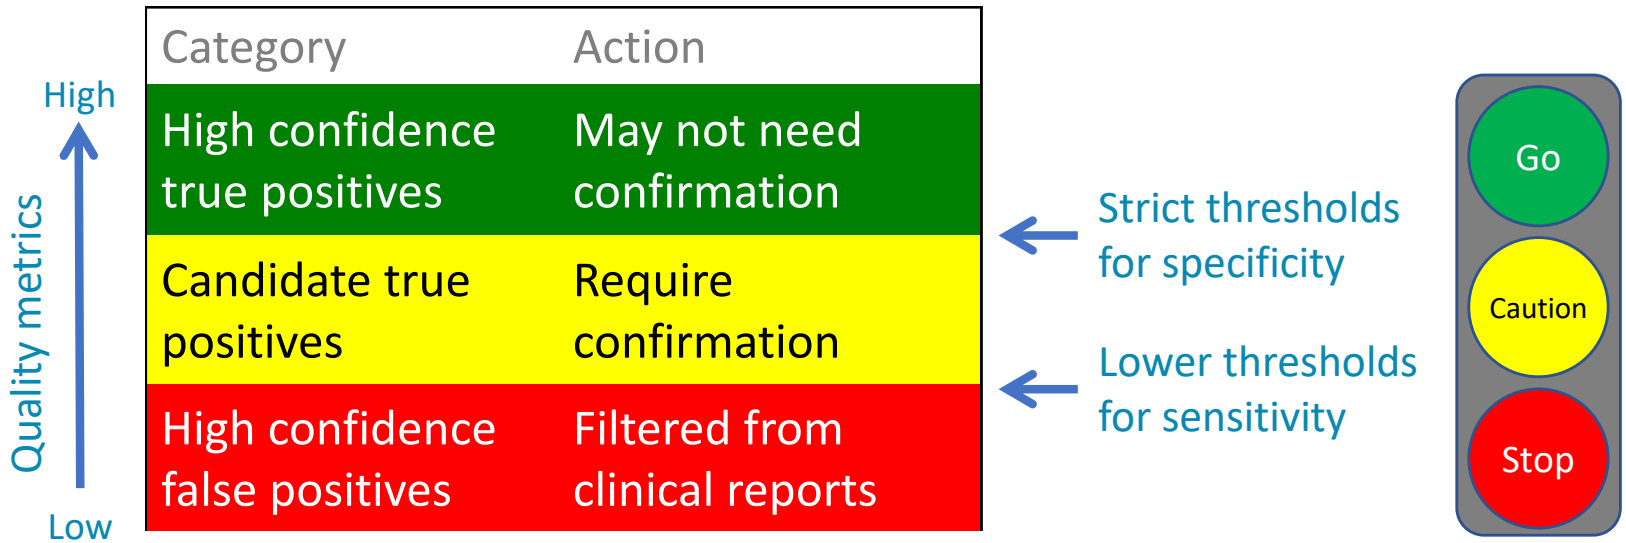
Quality metrics
High
Low
Category
Action
High confidence true positives
May not need confirmation
Candidate true positives
Require confirmation
High confidence false positives
Filtered from clinical reports
Strict thresholds for specificity
Lower thresholds for sensitivity
Go
Caution
Stop

Fig 1B

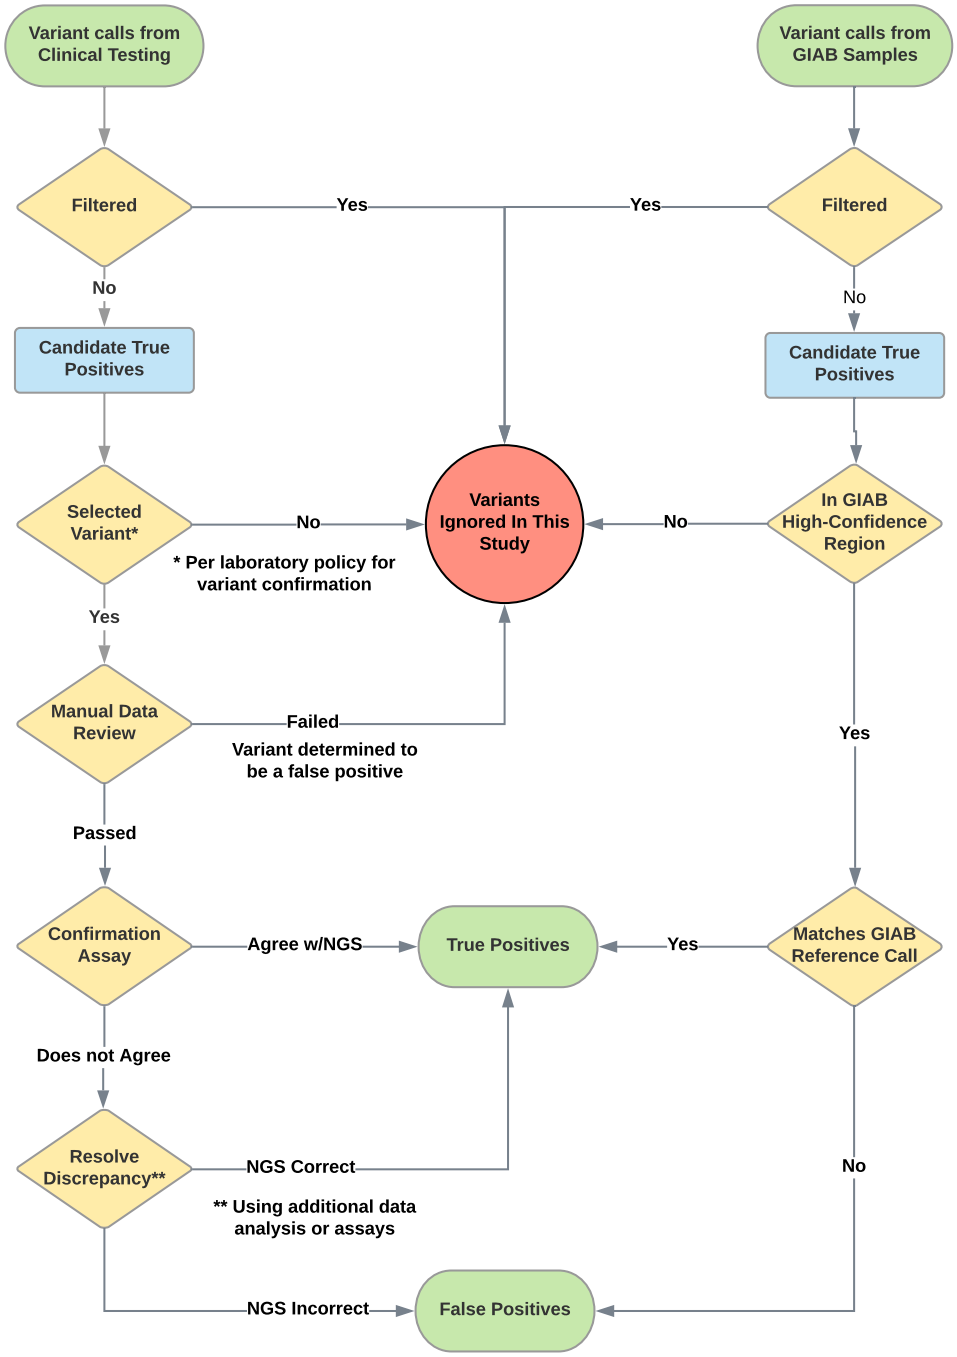
Variant calls from Clinical Testing
Variant calls from GIAB Samples
Filtered
Yes
Yes
Filtered
No
No
Candidate True Positives
Candidate True Positives
Selected Variant*
No
Variants Ignored In This Study
No
In GIAB High-Confidence Region
* Per laboratory policy for variant confirmation
Yes
Manual Data Review
Failed
Variant determined to be a false positive
Yes
Passed
Confirmation Assay
Agree w/NGS
True Positives
Yes
Matches GIAB Reference Call
Does not Agree
Resolve Discrepancy**
NGS Correct
No
** Using additional data analysis or assays
NGS Incorrect
False Positives

Fig 2 A-B

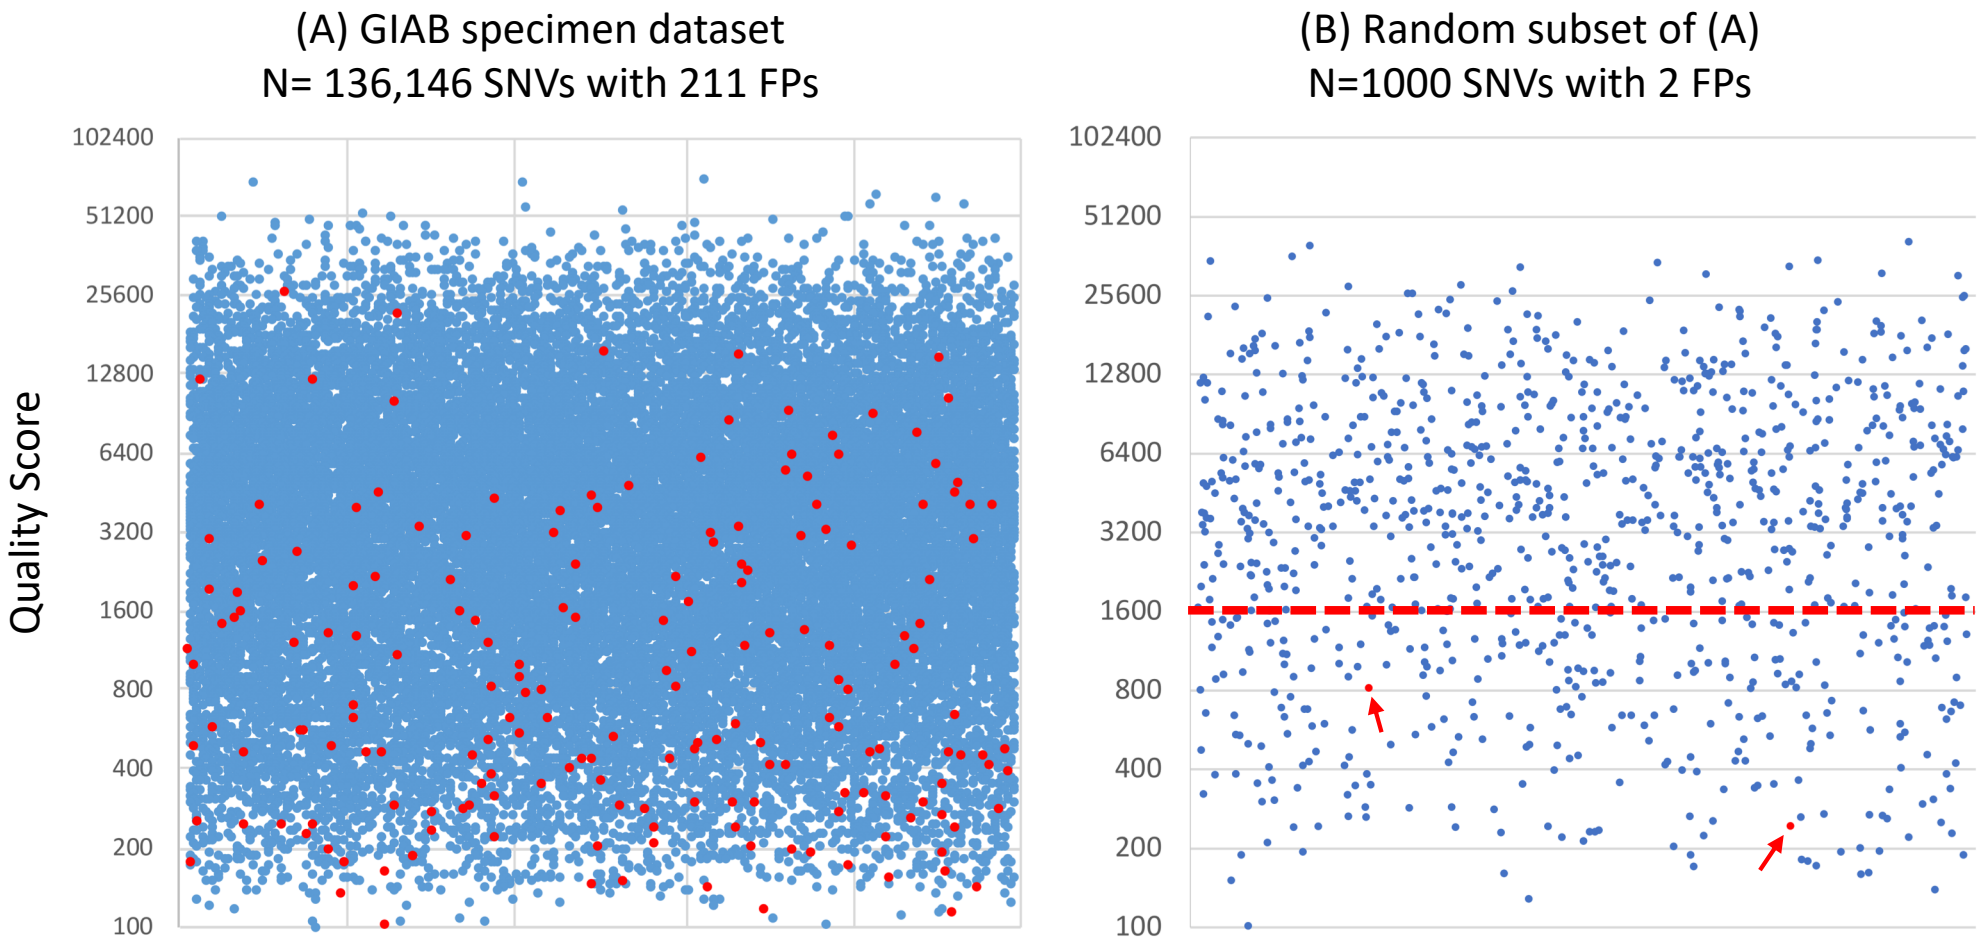
(A) GIAB specimen dataset
N= 136,146 SNVs with 211 FPs
(B) Random subset of (A)
N=1000 SNVs with 2 FPs
Quality Score
102400
51200
25600
12800
6400
3200
1600
800
400
200
100

Fig 2 C-D

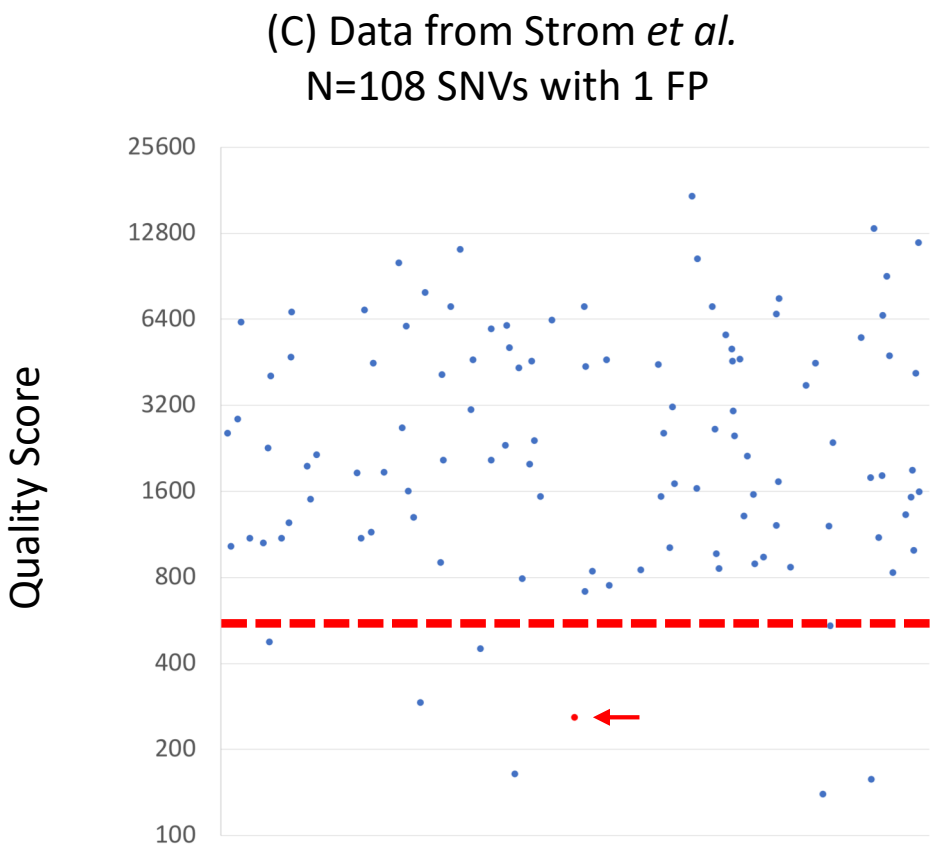
(C) Data from Strom *et al.*
N=108 SNVs with 1 FP
Quality Score
25600
12800
6400
3200
1600
800
400
200
100

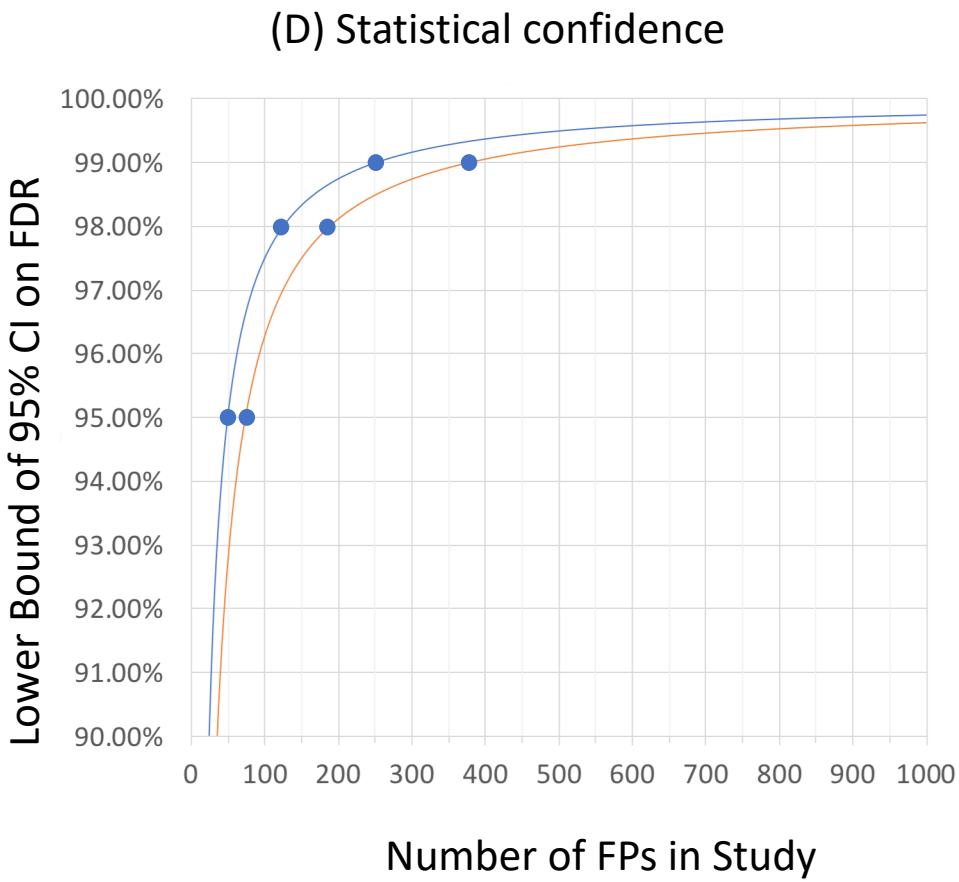
(D) Statistical confidence
Lower Bound of 95% CI on FDR
100.00%
99.00%
98.00%
97.00%
96.00%
95.00%
94.00%
93.00%
92.00%
91.00%
90.00%
0
100
200
300
400
500
600
700
800
900
1000
Number of FPs in Study

Fig 3 A-B

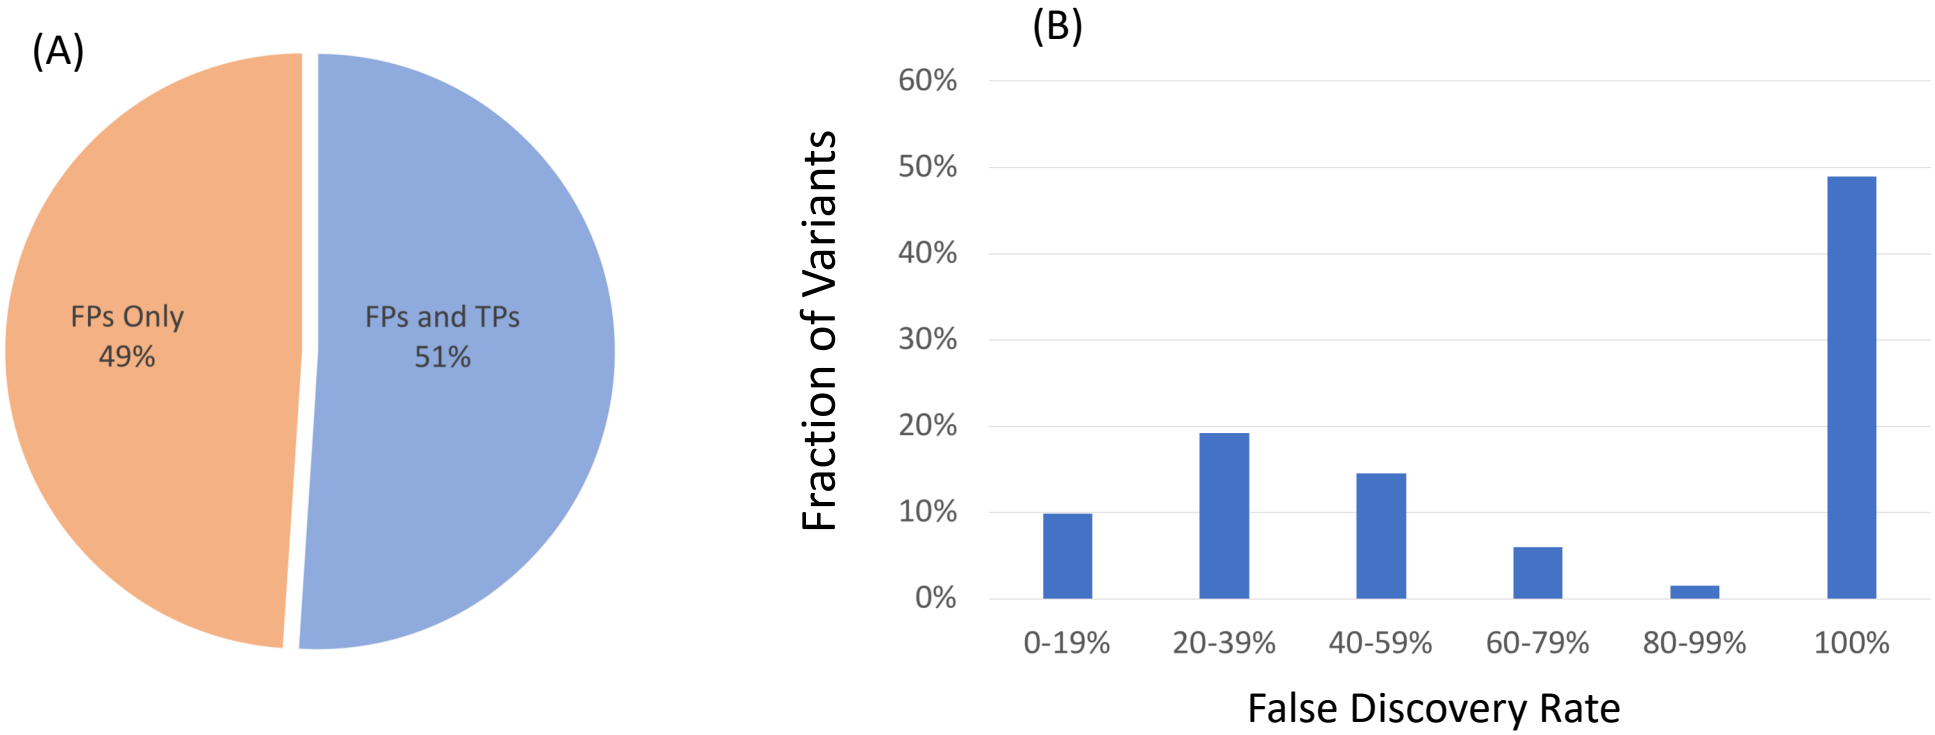
(A)
FPs Only
49%
FPs and TPs
51%
(B)
Fraction of Variants
60%
50%
40%
30%
20%
10%
0%
0-19%
20-39%
40-59%
60-79%
80-99%
100%
False Discovery Rate

Fig 4 A-B

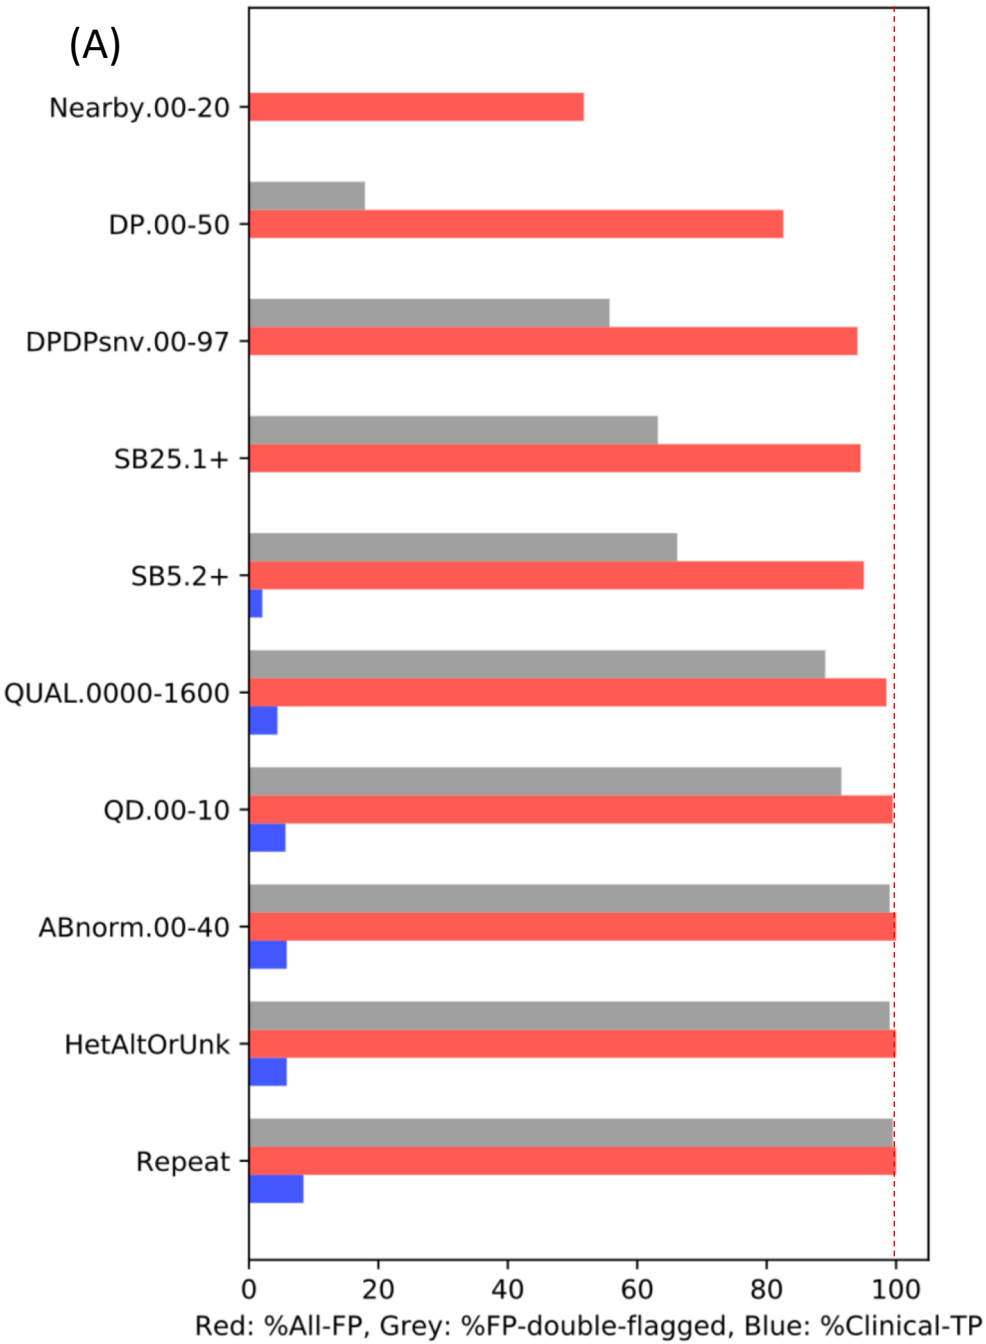
(A)
Nearby.00-20
DP.00-50
DPDPsnv.00-97
SB25.1+
SB5.2+
QUAL.0000-1600
QD.00-10
ABnorm.00-40
HetAltOrUnk
Repeat
0
20
40
60
80
100
Red: %All-FP, Grey: %FP-double-flagged, Blue: %Clinical-TP

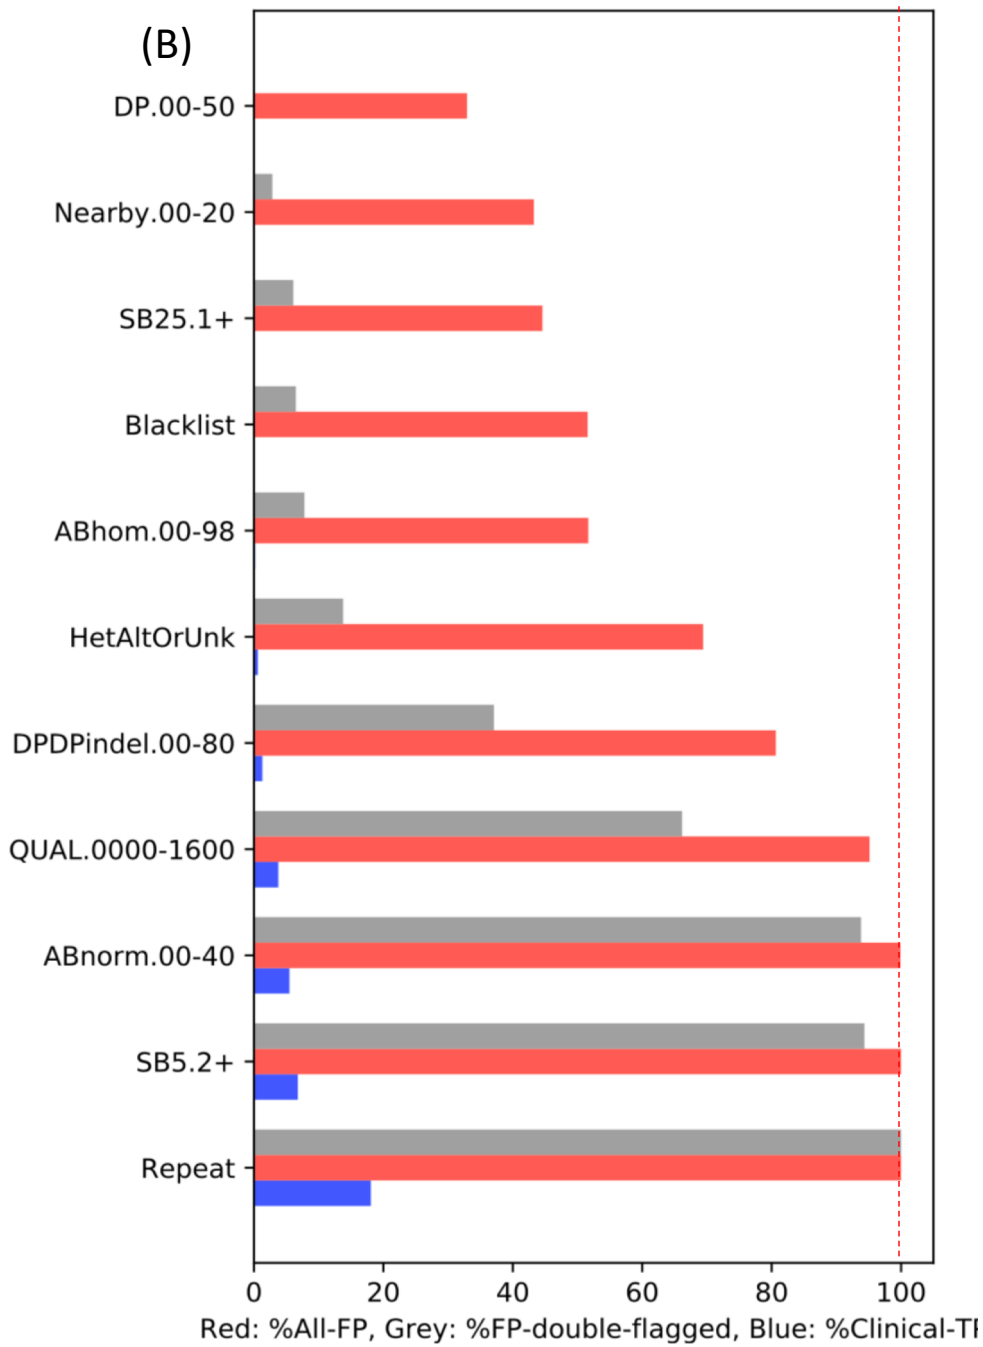
(B)
DP.00-50
Nearby.00-20
SB25.1+
Blacklist
ABhom.00-98
HetAltOrUnk
DPDPindel.00-80
QUAL.0000-1600
ABnorm.00-40
SB5.2+
Repeat
0
20
40
60
80
100
Red: %All-FP, Grey: %FP-double-flagged, Blue: %Clinical-TP

Main Table 1 **Study Data**

Source	Lab	Variant Type	Samples	Unique Variants	Variant Calls with Orthogonal Data	TP Calls	FP Calls	FDR	Total Calls	Total FPs	FPsens	CI lower bound	Notes
This study	Lab1	SNVs	GIAB	27,202	136,146	135,945	201	0.15%					
			Patients	2,840	3,699	3,689	10	0.27%	139,845	211	100%	98.9%	A,B
		Indels	GIAB	3,715	15,574	14,594	980	6.29%					
			Patients	1,749	2,274	2,262	12	0.53%	17,848	992	100%	99.8%	A,B
	Lab2	SNVs	GIAB	5,816	29,148	29,110	38	0.13%					
			Patients	4,359	4,934	4804	130	2.63%	34,082	168	100%	98.5%	A,B
		Indels	GIAB	1,185	3,617	3,343	274	7.58%					
			Patients	267	389	372	17	4.37%	4,006	291	100%	99.1%	A,B
Strom et al.	Only one	SNVs	Patients		108	107	1	0.93%	108	1	100%	5.1%	B,D
Mu et al.	Only one	SNVs	Patients		6,912	6,818	94	1.36%	6,912	94	100%	97.4%	B,D
		Indels	Patients		933	928	5	0.54%	933	5	100%	62.1%	B,D
Baudhuin et al.	Only one	SNVs	Patients	380	797	797	0	0%					
			1KG	736	736	736	0	0%	1,533	0	N/A	N/A	E,F
		Indels	Patients	63	122	122	0	0%					G
			1KG	26	26	26	0	0%	148	0	N/A	N/A	E,F

Notes:

A	Clinical and GIAB data were combined in the analysis.
B	Clinical data sets typically used manual review which can remove false positives before confirmation, reducing the FDR observed here. We did not do so in the GIAB data.
C	A large fraction of the FPs were in specific challenging regions of the *OTOA* and *CFTR* genes. Without these, FDRs wopuld be much lower.
D	We computed CIs based on data in the publication. No such statistics were provided by the study authors.
E	Baudhuin et al. analyzed only unique variants in their 1KG analysis.
F	Although discordances between NGS and reference calls were initially observed, improved 1KG reference data eliminated these.
G	Most of the indels in patients were intronic and homopolymer associated.

Abbreviations and definitions:

SNV	Single nucleotide variant
Indel	Insertion or deletion
Unique variant	A particular alteration which may be present in one or more speimens
FP	False positives from NGS. All of these false positives were detected by confirmation or reference data. In clinical practice these variants are removed from reports.
FDR	False discovery rate: Number of FPs divided by FPs plus TPs (true positives)
FPsens	The percentage of FPs captured using the study's proposed criteria
CI	95% confidence interval on FPsens
GIAB	Genome in a Bottle
1KG	1000 Genomes Project