



OPEN

Empirical Bayes single nucleotide variant-calling for next-generation sequencing data

Ali Karimnezhad^{1,2}✉ & Theodore J. Perkins^{3,4}

One of the fundamental computational problems in cancer genomics is the identification of single nucleotide variants (SNVs) from DNA sequencing data. Many statistical models and software implementations for SNV calling have been developed in the literature, yet, they still disagree widely on real datasets. Based on an empirical Bayesian approach, we introduce a local false discovery rate (LFDR) estimator for germline SNV calling. Our approach learns model parameters without prior information, and simultaneously accounts for information across all sites in the genomic regions of interest. We also propose another LFDR-based algorithm that reliably prioritizes a given list of mutations called by any other variant-calling algorithm. We use a suite of gold-standard cell line data to compare our LFDR approach against a collection of widely used, state of the art programs. We find that our LFDR approach approximately matches or exceeds the performance of all of these programs, despite some very large differences among them. Furthermore, when prioritizing other algorithms' calls by our LFDR score, we find that by manipulating the type I-type II tradeoff we can select subsets of variant calls with minimal loss of sensitivity but dramatic increases in precision.

Recent developments in next generation sequencing (NGS) technologies provide an insight into the task of mutation calling¹ and have made it possible to characterize the genomic alterations in a tumor in an unbiased manner². With the advancement of NGS technologies, the number of large-scale projects (especially cancer projects) dealing with point mutation in various tumor types has been increased rapidly, and many bioinformatics tools have been developed.

Several packages with different algorithms have been introduced in recent years to increase the accuracy in the mutation detection procedure. Perhaps, SAMtools³ is the most cited package for manipulating and converting alignments to different formats. It is also used for sorting and merging alignments, generating per-position information in pileup and mpileup formats, as well as calling single nucleotide variants (SNVs) and short insertion and deletions (INDELS). VarScan2⁴, a newer version of VarScan⁵, reads SAMtools's pileup or mpileup output and detects SNVs, INDELS and copy number variations using separate commands. Mutect⁶, an algorithm based on a Bayesian classifier, is claimed to be highly sensitive in the detection of very low frequency SNVs. In this paper, we use the Mutect2 tool in GATK that was developed by the Broad Institute and is used widely in the context of mutation-calling. VarDict⁷ calls different types of mutations, including SNVs, multiple-nucleotide variants, INDELS, complex and structural variants at the same time. Pisces⁸, which includes a variant-collapsing algorithm to unify variants broken up by read boundaries, basic filtering algorithms, and a simple Poisson-based variant confidence-scoring algorithm, is aimed to reduce noise or increase the likelihood of detecting true variants. In a recent study, Karimnezhad et al.⁹ compared the performance of the above-mentioned variant callers on different sets of NGS datasets sequenced on different platforms. For a list of recent mutation callers, readers may also refer to Xu¹⁰, where 46 mutation callers are reviewed.

Most variant callers in the literature rely on several pre-defined but adjustable parameters such as minimum base call quality (*BQ*), minimum mapping quality (*MQ*), strand bias threshold, etc. Remarkably, many of the parameters as well as their default values are not the same among the existing variant callers. For example, the default value for minimum allele frequency (*AF*) in Mutect2 is set to 0. See Table 1 for a list of selected parameters along with their default values in some selected algorithms. These parameters can strongly affect the output of the variant-calling algorithms. Best values for the parameters may in general depend on the type of sequencing

¹Department of Mathematics and Statistics, University of Ottawa, Ottawa K1N 9A7, Canada. ²Biostatistics and Risk Modelling Division, Bureau of Food Surveillance and Science Integration, Food Directorate, Health Products and Food Branch, Health Canada, Ottawa K1A 0K9, Canada. ³Regenerative Medicine Program, Ottawa Hospital Research Institute, Ottawa K1H 8L6, Canada. ⁴Department of Biochemistry, Microbiology and Immunology, University of Ottawa, Ottawa K1H 8M5, Canada. ✉email: a.karimnezhad@uottawa.ca

Option	Mutect2	VarScan2	SAMtools	VarDict	Piscis
Threshold for AF	0	0.01	–	0.05	0.01
Min <i>BQ</i> score	10	15	13	22.5(Illumina)	20
				15(PGM)	
<i>BQ</i> score threshold	18	–	–	–	–
Max <i>BQ</i> score	–	–	–	–	100
Min <i>MQ</i>	20	null	0	null	1
Mean min <i>MQ</i>	–	null	–	null	–
Maximum min <i>MQ</i>	null	–	–	–	–
Min coverage	–	8	–	–	10
Maximum coverage	–	–	250	–	–
Supporting reads to call a variant	–	2	–	–	–
Minimum variant quality score	–	–	–	–	15
Threshold for variant quality score filter	–	–	–	–	20
Strand bias filter	–	–	–	–	0.5
Minimum reads to strand bias	2	–	–	–	–

Table 1. Default values for options in different mutation callers. A dash means that the corresponding parameter was not defined in the caller's settings.

data, the type of mutations sought (somatic or germline), noise characteristics of the dataset, etc. As such, most users/analysts may not have enough expertise to adjust those pre-determined parameters properly. On the other hand, relying on default values for parameters may result in unreliable results and is problematic when comparing programs, because their default parameters differ.

The above approaches are either conventional Bayesian or frequentist methods, and do not take multiplicity and testing efficiency issues into account. Moreover, many programs fail to output a criterion by which SNV calls can be ranked, such as p-values, so that adjusting preference between type I and type II errors or constructing ROC curves is not naturally supported.

Some studies in the literature proposed germline variant detection in multi-sample NGS data using Bayesian methods. For example, Huang et al.¹¹ raised the question of how to quantify sequencing errors. They argue that it may not be ideal to quantify sequencing error only based on quality scores and propose an empirical Bayes method based on a Binomial model to estimate sequencing errors and identify variants. In another study, Zhao et. al¹² developed an optimal empirical Bayesian testing procedure to detect variants in NGS data, which is based on pooling a normalized amount of DNA from multiple samples. Due to a capacity issue, they assume samples may be distributed and sequenced independently in $M > 1$ pools and each pool consists of N individuals. Although having M pools with N individuals may be cost-effective, however, our available samples (similar to many other clinical labs) instruct us to base the model only on one DNA sample for one individual. Similar to Huang et al.¹¹, Ding et al.¹³ suggested that estimating sequencing errors based on quality scores may not be precise. Using single-sample NGS data, they proposed an empirical Bayes method to model the dispersion of minor AF and call germline variants.

In this paper, we propose a Binomial model and include an error rate based on a weighted average of base call and mapping quality scores as part of model parameters. We use an empirical Bayesian approach to develop a local false discovery rate (LFDR) estimator for SNV detection. In contrast to existing algorithms, our novel approach calls SNVs not just on a site-by-site basis, but by simultaneously using information across all the sites to build a probabilistic model of the data.

To the best of our knowledge, LFDR estimation has not been employed in the task of variant-calling, but it has been well developed in a variety of other contexts, and different strategies have been introduced in the literature. To name a few LFDR estimation based approaches, readers may refer to Pan et al.¹⁴, Efron et. al¹⁵, Efron¹⁶, Padilla and Bickel¹⁷, Yang et al.¹⁸, Karimnezhad and Bickel¹⁹, and Karimnezhad²⁰.

Obviously, all variant-calling algorithms determine mutations based on some observed evidence, but since evidence supporting variants (including number of reference and alternative read counts) differs from site to site, the confidence in calling a site as a variant site varies. We fill this gap by introducing an LFDR-based algorithm that meaningfully scores variants called by any variant caller and prioritizes them from most to least probable variants. This helps with significantly reducing many false positives.

The structure of this work is as follows. In the “Methods” section, we provide a detailed presentation of the model, method and the algorithm we propose for germline SNV calling. We also propose a modified version of the LFDR algorithm so that it can prioritize variants called by any desired variant caller. In the “Results” section, we use a suite of gold-standard cell line data to evaluate the performance of our proposed LFDR approach against a collection of widely used, state of the art programs: MuTect2, SAMtools, VarScan2, VarDict, and Piscis.

Methods

We consider analyzing a single DNA sample from a healthy patient, where only germline mutations are expected to be seen. We also assume that the data has been mapped to a reference genome, which specifies one of *A*, *C*, *G*, or *T* bases for every position. Each mapped read shows one of the four possible bases (*A*, *C*, *G*, or *T*) for the same position. With no loss of generality, we also suppose, for a replicate (technical or biological), that sequencing covers *p* sites (individual positions in the genome) of which *p*₀ sites are non-mutant. For each locus *i*, *i* = 1, . . . , *p*, we assume that there are *K*_{*i*} (short) known number of reads covering the locus *i*. At each locus *i*, we suppose there are four possible bases (*A*, *C*, *G*, *T*) of which *R*_{*i*} random reads carry the reference allele. We further suppose that of the remaining *K*_{*i*} − *R*_{*i*} reads, *M*_{*i*} random number of reads carry the most alternative (dominant) allele and, we assign the remaining random alternative reads to *X*_{1*i*} and *X*_{2*i*}. We briefly refer to these random variables by **X**_{*i*} = (*R*_{*i*}, *M*_{*i*}, *X*_{1*i*}, *X*_{2*i*}). Generally, the larger *M*_{*i*} is, the more support we have for an alternative allele being present at site *i* in the DNA being sequenced. Indeed, if the data had no sequencing or mapping errors, *M*_{*i*} > 0 could only arise as the result of an alternative allele being present. However, it is well-recognized that sequence data does contain errors. Although those errors can arise in many ways, the standard model error employed by the community is the error in mistakenly reporting an alternative allele as one of the other three possible alleles, and vice versa. We refer to this error by *e*_{*i*}. Obviously, if there is no error, mutant sites are expected to be those *p* − *p*₀ sites, for which *M*_{*i*} is positive.

Suppose that θ_i represents the dominant alternative AF at the *i*th site. Obviously, given the information that *i*th site has no mutation, $\theta_i = 0$ and otherwise, $\theta_i > 0$. We assume $P(\theta_i = 0) = \pi_0$ and $P(\theta_i > 0) = 1 - \pi_0$. If $\theta_i > 0$, we assume it follows a probability density function $g(\cdot)$. We assume that each read is drawn independently from either the reference or alternative allele pools. However, we also assume an independent chance *e*_{*i*} of read error in each read at each site, which randomly changes the correct allele to one of the other three options. The error term *e*_{*i*} can be calculated in different ways. For example, in Mutect, error is estimated based on Phred-scaled quality scores of mapped reads to each position. We suggest to take the error rate as a weighted average of average base call qualities and average mapping qualities. Supposed at position *i*, for an arbitrary base pair *U* ∈ {*A*, *C*, *G*, *T*}, BQ_i^U and MQ_i^U represent average base call quality and average mapping quality across all reads at that position, respectively, and let Q_i^U be the average of BQ_i^U and MQ_i^U , i.e., $Q_i^U = \frac{1}{2}(BQ_i^U + MQ_i^U)$. Depending on whether the base pair observed at position *i* is a reference (*R*), dominant (*M*) or one of the remaining alternative alleles (*X*₁ or *X*₂), we calculate Q_i^R , Q_i^M , $Q_i^{X_1}$ and $Q_i^{X_2}$, and use them to calculate the weighted average

quality $Q_i = \frac{Q_i^R \times R_i + Q_i^M \times M_i + Q_i^{X_1} \times X_{1i} + Q_i^{X_2} \times X_{2i}}{R_i + M_i + X_{1i} + X_{2i}}$. Then, we define the error rate to be $e_i = 10^{-\frac{Q_i}{10}}$. As in Mutect⁶, we define the following four probabilities. The probability that a randomly chosen read covering site *i* shows the reference allele is $p_{R_i} = \theta_i \frac{e_i}{3} + (1 - \theta_i)(1 - e_i)$. This is explained as the sum of the probabilities that the DNA actually contained the reference allele at site *i* and it was correctly read, $(1 - \theta_i)(1 - e_i)$, and probability that the DNA actually contained the alternative allele but it was misread, and by chance, it was misread as the reference allele, $\theta_i \frac{e_i}{3}$. Similarly, the probability that a random read covering site *i* shows the alternative allele is $p_{M_i} = \theta_i(1 - e_i) + (1 - \theta_i) \frac{e_i}{3}$, which arises either as correct reading of the alternative allele or misreading the reference allele as the alternative. Finally $p_{X_{1i}} = p_{X_{2i}} = \frac{e_i}{3}$ is the chance that one of the two other alleles that are not reference or alternative occurs at site *i*. Putting these together, the probability of the total data at site *i*, **X**_{*i*} is multinomial with *K*_{*i*} tries (reads) and probabilities of reference, alternative, and the other two alleles. At this point, we diverge from the model of Mutect by proposing a new model governing the parameter θ_i , which ultimately allows us to develop our empirical Bayesian LFDR approach. We assume $\mathbf{X}_i | \theta_i \sim \text{Multi}(K_i, p_{R_i}, p_{M_i}, p_{X_{1i}}, p_{X_{2i}})$ where $p_{R_i} = \theta_i \frac{e_i}{3} + (1 - \theta_i)(1 - e_i)$, $p_{M_i} = \theta_i(1 - e_i) + (1 - \theta_i) \frac{e_i}{3}$, $p_{X_{1i}} = \frac{e_i}{3}$, $p_{X_{2i}} = \frac{e_i}{3}$, where θ_i could be either 0 or positive with probabilities π_0 and $1 - \pi_0$, respectively. If $\theta_i > 0$, we assume $\theta_i \sim g(\cdot)$, where $g(\cdot)$ is unknown. In our proposed approach we do not need to know the form of $g(\cdot)$.

To discover whether site *i* is a mutant site, we propose testing the null hypothesis $H_{0i} : \theta_i = 0$ against the alternative hypothesis $H_{1i} : \theta_i > 0$, *i* = 1, . . . , *p*. To do so, we focus on estimating $\psi_i \equiv P(\theta_i | \mathbf{X}_i)$, the posterior probability that the null hypothesis H_{0i} is true. Once estimated, it is compared with a pre-specified threshold leading to either rejecting or failing to reject the null hypothesis. The quantity ψ_i is well-known in the literature as LFDR.

With the above settings, the following probability functions are derived under the null and the alternative hypotheses, respectively,

$$P(\mathbf{X}_i | \theta_i = 0) = \binom{K_i}{R_i, M_i, X_{1i}, X_{2i}} (1 - e_i)^{R_i} \left(\frac{e_i}{3}\right)^{M_i + X_{1i} + X_{2i}} = \binom{K_i}{R_i, M_i, X_{1i}, X_{2i}} (1 - e_i)^{R_i} \left(\frac{e_i}{3}\right)^{K_i - R_i} \tag{1}$$

and

$$\begin{aligned} P(\mathbf{X}_i | \theta_i > 0) &= \int P(\mathbf{X}_i | \theta_i > 0) g(\theta_i) d\theta_i \\ &= \int_0^1 \binom{K_i}{R_i, M_i, X_{1i}, X_{2i}} \left(\theta_i \frac{e_i}{3} + (1 - \theta_i)(1 - e_i)\right)^{R_i} \left(\theta_i(1 - e_i) + (1 - \theta_i) \frac{e_i}{3}\right)^{M_i} \left(\frac{e_i}{3}\right)^{X_{1i} + X_{2i}} g(\theta_i) d\theta_i. \end{aligned} \tag{2}$$

Now, using the Bayes formula, the LFDR can be expressed by

$$\psi_i = \frac{\pi_0 P(\mathbf{X}_i | \theta_i = 0)}{\pi_0 P(\mathbf{X}_i | \theta_i = 0) + (1 - \pi_0) P(\mathbf{X}_i | \theta_i > 0)}, \tag{3}$$

where $P(X_i|\theta_i = 0)$ and $P(X_i|\theta_i > 0)$ are given by (1) and (2), respectively, and π_0 is the proportion of non-mutant sites. Both the parameters π_0 and $g(\theta_i)$ (in $P(X_i|\theta_i > 0)$) are unknown and need to be estimated before making any inference.

Now, let δ_i be a binary decision rule corresponding to i th set of hypotheses H_{0i} and H_{1i} . We assume that $\delta_i = 1$ if the null hypothesis H_{0i} is rejected, and $\delta_i = 0$, otherwise. But since such binary decisions can lead to some errors, we define the following loss function when testing the null hypothesis at site i :

$$L(\theta_i, \delta_i) = \begin{cases} 0 & \theta_i = 0, \delta_i = 0 \text{ or } \theta_i > 0, \delta_i = 1, \\ l_I & \theta_i = 0, \delta_i = 1, \\ l_{II} & \theta_i > 0, \delta_i = 0, \end{cases}$$

where l_I and l_{II} are loss values incurred due to making type I and type II errors, respectively, see Karimnezhad and Bickel¹⁹. Also, see Zhao et. al¹² where a specific version of this loss with $l_I = \lambda$ and $l_{II} = 1$ is used.

To take all p sites in a sample to account, suppose that $\delta = (\delta_1, \dots, \delta_p)$ represents a vector of estimators of $\theta = (\theta_1, \dots, \theta_p)$ and measure the aggregated loss by $L(\theta, \delta) = \sum_{i=1}^p L(\theta_i, \delta_i)$, which takes both type-I and type-II errors into account. Now, to derive a Bayesian decision rule in the above-mentioned hypothesis testing problem, let $\rho_X(\theta, \delta) = E[L(\theta, \delta)|X]$ denote the posterior risk of choosing a decision vector δ , where $X = (X_1, \dots, X_p)$. Then, similar to Karimnezhad and Bickel¹⁹, it can be verified that the Bayesian decision vector $\delta^B = (\delta_1^B, \dots, \delta_p^B)$ with

$$\delta_i^B = \begin{cases} 0 & \psi_i > \frac{l_I}{l_I + l_{II}}, \\ 1 & \psi_i \leq \frac{l_I}{l_I + l_{II}}, \end{cases} \quad (4)$$

where ψ_i is given by (3), minimizes the posterior risk w.r.t. δ .

The above equation suggests that mutant sites can be determined by estimating ψ_i , $i = 1, \dots, p$, and then comparing it with the threshold $\frac{l_I}{l_I + l_{II}}$. However, estimating ψ_i is challenging due to the complicatedness of the form of $P(X_i|\theta_i > 0)$ as well as $g(\cdot)$, the distribution of AFs.

One immediate solution to estimating $g(\cdot)$ would be to assign a non-informative prior to the alternative AFs θ_i . In a similar situation, Zhao et. al¹² consider a uniform distribution on the interval $(0, a)$, and then they estimate the value of a as well as π_0 using the traditional method of moments. While simple, this assumption seems unrealistic. Figure 1 represents the distribution of AFs at known mutant as well as non-mutant sites from TST-GM12878 data (see Results for more information).

From Fig. 1, we observe that the assumption that AFs take any number between 0 and 1 with equal probability is unrealistic. Estimating π_0 may be connected to the estimation of $g(\cdot)$ and consequently, a bad estimate of $g(\cdot)$ may directly lead to a bad estimate of π_0 . This indeed is true in some datasets with ambiguous AF ranges, including our datasets. Clearly, from this figure, AFs of mutant sites is distributed around either 0.5 or 1. This highlights the importance of finding a suitable estimate of $g(\cdot)$. Dependency of estimated values of parameters of LFDR to each other can be seen in different algorithms in the literature. See for example Zhao et. al¹² where their proposed estimator of π_0 is a function of the estimated value of the upper bound a in the uniform prior assigned to the alternative AF. For other examples, see Pan et al.¹⁴ and Karimnezhad and Bickel¹⁹.

As Fig. 1 suggests, one may fit a beta distribution to model the alternative AFs, which adds extra parameters to the model. To avoid model complexity, as an innovative approach, we suggest estimating $g(\cdot)$ empirically,

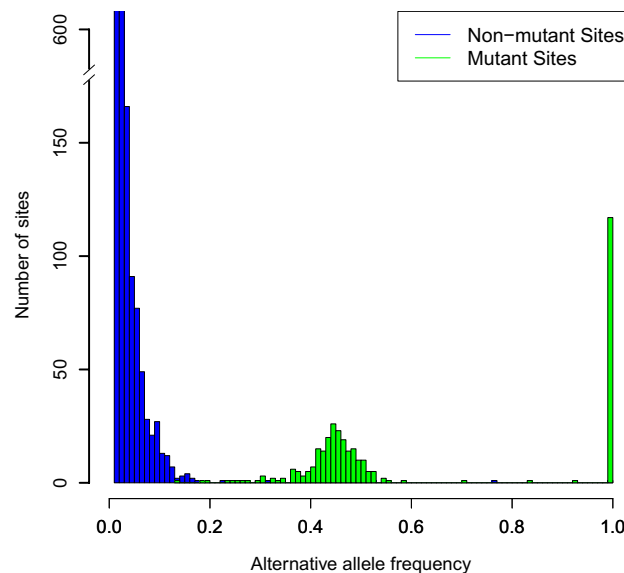


Figure 1. Histogram of distribution of alternative AFs from TST-GM12878 data.

based on the alternative AFs at a set of known or suspected mutated sites. This raises the question of which set of sites to use.

An immediate solution to the above problem would be to take all sites for which M_i counts are positive. But as many non-zero alternative AFs in NGS studies turn out to correspond to artifacts, we need to find and exclude them from our list so that a precise empirical distribution for $g(\cdot)$ can be estimated. Suppose that \mathcal{I}_s stands for the set of s indices of those non-zero θ_i corresponding to mutant sites, and let $\boldsymbol{\theta}_s$ be the vector of the corresponding empirical AFs. In our settings, s represents the number of elements of the set \mathcal{I}_s . Now, let $G(\cdot)$ represent the cumulative distribution function (CDF) that corresponds to $g(\cdot)$. Then, the corresponding empirical CDF is given by $\widehat{G}_s(t) = \frac{1}{s} \sum_{i \in \mathcal{I}_s} 1_{\theta_i \leq t}$, $t \in \mathfrak{R}$. Obviously, due to the strong law of large numbers, $\widehat{G}_s(t)$ converges to $G(t)$ as $n \rightarrow \infty$ almost surely, for every value of t , and this implies that $\widehat{G}_s(t)$ is a consistent estimator. Then, one may estimate $P(X_i | \theta_i > 0)$ by

$$P_{\boldsymbol{\theta}_s}(X_i | \theta_i > 0) \simeq \frac{1}{s} \sum_{i \in \mathcal{I}_s} \left\{ \binom{K_i}{R_i, M_i, X_{1i}, X_{2i}} \left(\theta_i \frac{e_i}{3} + (1 - \theta_i)(1 - e_i) \right)^{R_i} \left(\theta_i(1 - e_i) + (1 - \theta_i) \frac{e_i}{3} \right)^{M_i} \left(\frac{e_i}{3} \right)^{X_{1i} + X_{2i}} \right\}. \quad (5)$$

Since \mathcal{I}_s includes many artifacts, s needs to be optimally learned. We propose an algorithm that allows for repeatedly updating the set of mutant sites and leads to a precise estimate of $g(\cdot)$.

An empirical Bayes mutation detection procedure

To discover mutant sites, we propose the procedure outlined in Algorithm 1. The proposed algorithm, like many existing procedures in the literature, including Mutect2, VarScan2⁴, VarDict⁷ and Pisces⁸ is comprised of three main steps.

In the pre-processing step, we propose excluding low quality bases from the input file (either FASTQ or BAM format). Thus, observations in our model are only those bases that pass a minimum average BQ and average MQ threshold. BQ corresponds to an error rate of $e = 10^{-\frac{BQ}{10}}$, see for example Cibulskis et al.⁶ and Dunn, et al.⁸. Thus taking $BQ = 20$ corresponds to an error rate of 0.01 which can be inferred as expecting one miscalled base in reading 100 bases. MQ is related to aligning reads to a reference genome, and it usually varies between 0 and 60. Similar to Pisces, we impose a threshold of 20, and 30 to the minimum BQ and MQ , respectively. Different publicly available softwares such as SAMtools³ and bam-readcount (<https://github.com/genome/bam-readcount>) can be applied to exclude such low quality reads. Once low quality bases are excluded, the counts K_i , R_i , M_i , X_{1i} and X_{2i} need to be formed.

Although in Eq. (5) we only include all non-zero empirical AFs, one may tweak the algorithm by just assuming that AFs follow, for example, a uniform distribution on the interval (0, 1), or a subset of it. Thus, \mathcal{I}_s^j in Step 2.6 of Algorithm 1, may be alternatively taken to be a set of N (for example 1000) sites for which their AFs are randomly generated from a uniform distribution on the interval (0, 1). However, once the estimation procedure enters Step 2.10, the updated \mathcal{I}_s^j set will only depend on the actual empirical AFs. A third approach would be to estimate LFDRs by focusing only on the original uniformly sampled AFs, without updating the set of indices \mathcal{I}_s^j in step 2.10, i.e., for all j , $\mathcal{I}_s^j = \mathcal{I}_s^1$. We refer to this approach by “uniform” estimation. For comparison purposes, we use these three methods in our data analysis and will refer to them by “empirical”, “uniform/empirical” and “uniform” estimation of $g(\cdot)$, respectively. In Step 2.10, we took $\epsilon = 0.001$.

In the post-processing step, we take 0.01 and 10 as default AF threshold (AFT) and read depth threshold (DPT) values, respectively. This is not uncommon, and many variant callers impose such thresholds. For example, Pisces⁸ has the same default values for these thresholds. We remark that one may apply this filtering in Step 1. However, the final estimated π_0 will reflect the proportion of non-mutant sites w.r.t. the filtered input file rather than the original one.

Step 1. Pre-processing. Exclude reads not passing minimum average base call quality (BQ) and minimum average mapping quality (MQ) thresholds.

Step 2. Estimation.

Step 2.1 Specify the LFDR threshold $l_I/(l_I + l_{II})$.

Step 2.2 Specify the error rate e_i , $i = 1, \dots, p$.

Step 2.3 For each $i = 1, \dots, p$, calculate $P(\mathbf{X}_i | \theta_i = 0)$.

Step 2.4 Take $j = 1$. This represents an iteration number.

Step 2.5 Choose an initial value for $\pi_0 \in (0, 1)$, set $\hat{\pi}_0^j = \pi_0$.

Step 2.6 Define \mathcal{S}_s^j to be the set of all sites for which M_i is positive, and let $\hat{\theta}_s^j = \hat{\theta}_s$ represent the vector of the corresponding empirical AFs.

Step 2.7 For each $i = 1, \dots, p$, calculate $P_{\hat{\theta}_s^j}(\mathbf{X}_i | \theta_i > 0)$.

Step 2.8 Estimate LFDRs by

$$\hat{\Psi}_i^j = \frac{\hat{\pi}_0^j P(\mathbf{X}_i | \theta_i = 0)}{\hat{\pi}_0^j P(\mathbf{X}_i | \theta_i = 0) + (1 - \hat{\pi}_0^j) P_{\hat{\theta}_s^j}(\mathbf{X}_i | \theta_i > 0)}, \quad i = 1, \dots, p.$$

Step 2.9 Compute the Bayes rule

$$\delta_i^{B,j} = \begin{cases} 0 & \hat{\Psi}_i^j > \frac{l_I}{l_I + l_{II}}, \\ 1 & \hat{\Psi}_i^j \leq \frac{l_I}{l_I + l_{II}}. \end{cases}$$

Step 2.9 Step up j by one, i.e., $j = j + 1$.

Step 2.10 Reset \mathcal{S}_s^j to be the set of indices i where $\delta_i^{B,j-1} = 1$, and update $\hat{\theta}_s^j$ based on the new \mathcal{S}_s^j .

Step 2.11 Define $\hat{\pi}_0^j = 1 - \frac{1}{p} \sum_{i=1}^p \delta_i^{B,j-1}$.

Step 2.12 Repeat Steps 2.7-2.11 until the difference between the new estimate of $\hat{\pi}_0$ and its previous value does not exceed a small number ϵ , i.e., $|\hat{\pi}_0^j - \hat{\pi}_0^{j-1}| < \epsilon$. Refer to the last j by j^* .

Step 2.13 Now, any site i for which $\delta_i^{B,j^*} = 1$ is a mutant site.

Step 3. Post-Processing. Exclude sites with AF and read depth (or coverage) not passing a pre-determined threshold.

Algorithm 1. LFDR-based variant-calling algorithm.

LFDR as a variant prioritization tool

We suggest that our LFDR calculation can be applied to prioritize any list of mutations called by any mutation caller from most to least probable mutations. This can be done by modifying Algorithm 1 so that $\hat{\pi}_0$ is estimated as the number of reference sites according to the caller, out of all those assayed, and $g(\cdot)$ is estimated as the empirical distribution of variant sites. The LFDR calculation is then applied to all variant sites using these estimates of π_0 and $g(\cdot)$. Algorithm 2 outlines the prioritization steps.

- Step 1** Specify the error rate e_i for genomic site i , $i = 1, \dots, p$.
- Step 2** For each $i = 1, \dots, p$, calculate $P(\mathbf{X}_i | \theta_i = 0)$ in equation (1).
- Step 3** Set $\hat{\pi}_0 = 1 - \frac{p_0}{p}$, where p represents the total number of sites in a genomic region of interest of which p_0 sites are declared by any chosen variant caller to be non-mutant sites.
- Step 4** In equation (5), define \mathcal{S}_s to be the set of all $p - p_0$ sites deemed to be variant sites by the chosen variant caller, and let $\hat{\theta}_s$ represent the vector of the corresponding AFs.
- Step 5** For each $i = 1, \dots, p$, approximate $P_{\hat{\theta}_s}(\mathbf{X}_i | \theta_i > 0)$ in (5).
- Step 6** Estimate the corresponding LFDRs by

$$\hat{\psi}_i = \frac{\hat{\pi}_0 P(\mathbf{X}_i | \theta_i = 0)}{\hat{\pi}_0 P(\mathbf{X}_i | \theta_i = 0) + (1 - \hat{\pi}_0) P_{\hat{\theta}_s}(\mathbf{X}_i | \theta_i > 0)}, \quad i = 1, \dots, p.$$

Algorithm 2. LFDR-based prioritizing algorithm for calls made by a desired variant caller.

Results

We now focus on evaluating our proposed mutation calling algorithm on some data generated by clinical assays. We employ DNA sequencing data generated from two Coriell cell lines, GM12877 and GM12878, previously studied by Karimnezhad et al.⁹. These are immortalized healthy (non-cancer) cell lines that are expected to carry only germline variants. These cell lines have been extremely well characterized by numerous groups. Thus, we can know exactly which variants they harbor relative to the reference human genome, allowing us to accurately assess true and false positives and negatives. In particular, we rely on a list of known mutations published by Eberle et al.²² where a comprehensive and genome-wide catalog of high-confidence variants mutations for a collection of Coriell cell lines, including GM12877 and GM12878 is presented (available through <https://www.ncbi.nlm.nih.gov/gap/> under accession phs001224.v1.p1). The data we used was sequenced in two different ways: (1) on an Illumina NextSeq500 sequencer using Illumina's TruSight170 targeted sample preparation method (TST170 for short), and on an Ion Torrent PGM sequencer using an OncoPrint Focus targeted panel (OF for short).

The TST170 data spans 514761 total bases, covering parts of 170 genes, and includes six technical replicates of each of the Coriell cell lines GM12877 and GM12878. By intersecting the list of known mutations with the TST170 genomic regions, we determined that our GM12877 and GM12878 data should contain 336 and 343 mutations, respectively. The OF panel covers 29008 total bases in 47 genes, and includes three technical replicates of each of the Coriell cell lines GM12877 and GM12878. There should be 24 and 26 known mutations present in each replicate of the GM12877 and GM12878 data, respectively. For more information regarding the datasets and also sequencing platforms, readers may refer to Karimnezhad et al.⁹.

For a given single replicate, denote the total number of true positives, false positives and false negatives by TP , FP and FN respectively. We measure the performance of the algorithms by computing precision or positive predictive value $Prec = \frac{TP}{TP + FP}$ and sensitivity $Sens = \frac{TP}{TP + FN}$. A good algorithm is expected to have high $Prec$ and $Sens$ rates.

To apply the algorithm on the TST170 as well as OF replicates, we used the bam-readcount package to calculate the counts K_i , R_i , M_i , X_{1i} and X_{2i} , and then followed Steps 1-3 of Algorithm 1. We then compared the list of final detected variants with the lists of known mutations. Also, to investigate the impact of a chosen LFDR threshold in Step 2.1 of the algorithm on the detection accuracy, we picked 10^{-300} , 10^{-200} , 10^{-100} , 10^{-50} , 10^{-20} , 10^{-10} , 10^{-2} and 0.5.

We measured the performance of the proposed approaches by calculating TP , FP , FN , and consequently $Prec$ and $Sens$ values for different datasets and LFDR thresholds. As a second performance measurement, we compared the proposed algorithm with Mutect2, VarScan2, SAMtools, VarDict and PISCES. Illumina offers PISCES for the analysis of TST170 data through their customized pipeline. The application accepts paired-end fastq files as inputs, generates BAM files, and after aligning to the reference human genome (hg19) by the Isaac aligner²³, uses PISCES to generate a list of mutations. Then, the final list of mutations is generated after some internal filtering. Our work was carried out in the context of a clinical sequencing project, where use of hg19 remains widespread, prompting us to use hg19 for alignments. However, we would expect very similar results with hg38. Our data is obtained by targeted sequencing of a panel of clinically-relevant genes, and there should be little or no difference between hg19 and hg38 for these highly-studied genes. We should add that, to compare the performance of our proposed algorithm with PISCES, as well as VarScan2, Mutect2, SAMtools and VarDict, we used the same aligned BAM files to reduce some possible alignment errors. Because differences in default parameters of the five algorithms in Table 1 are potential sources of discrepancies in the list of final mutations, we set their parameters as similar as possible. We set the minimum variant AF to 0.01. We set minimum base call quality and minimum mapping quality to 20. Finally, we set the minimum coverage for a called variant at 10 reads.

Figure 2 represents the average $Prec$ and $Sens$ values over replicates based on calls made by the LFDR approach for different thresholds, as well as the other five variant callers. From this figure, we observe that for all the LFDR

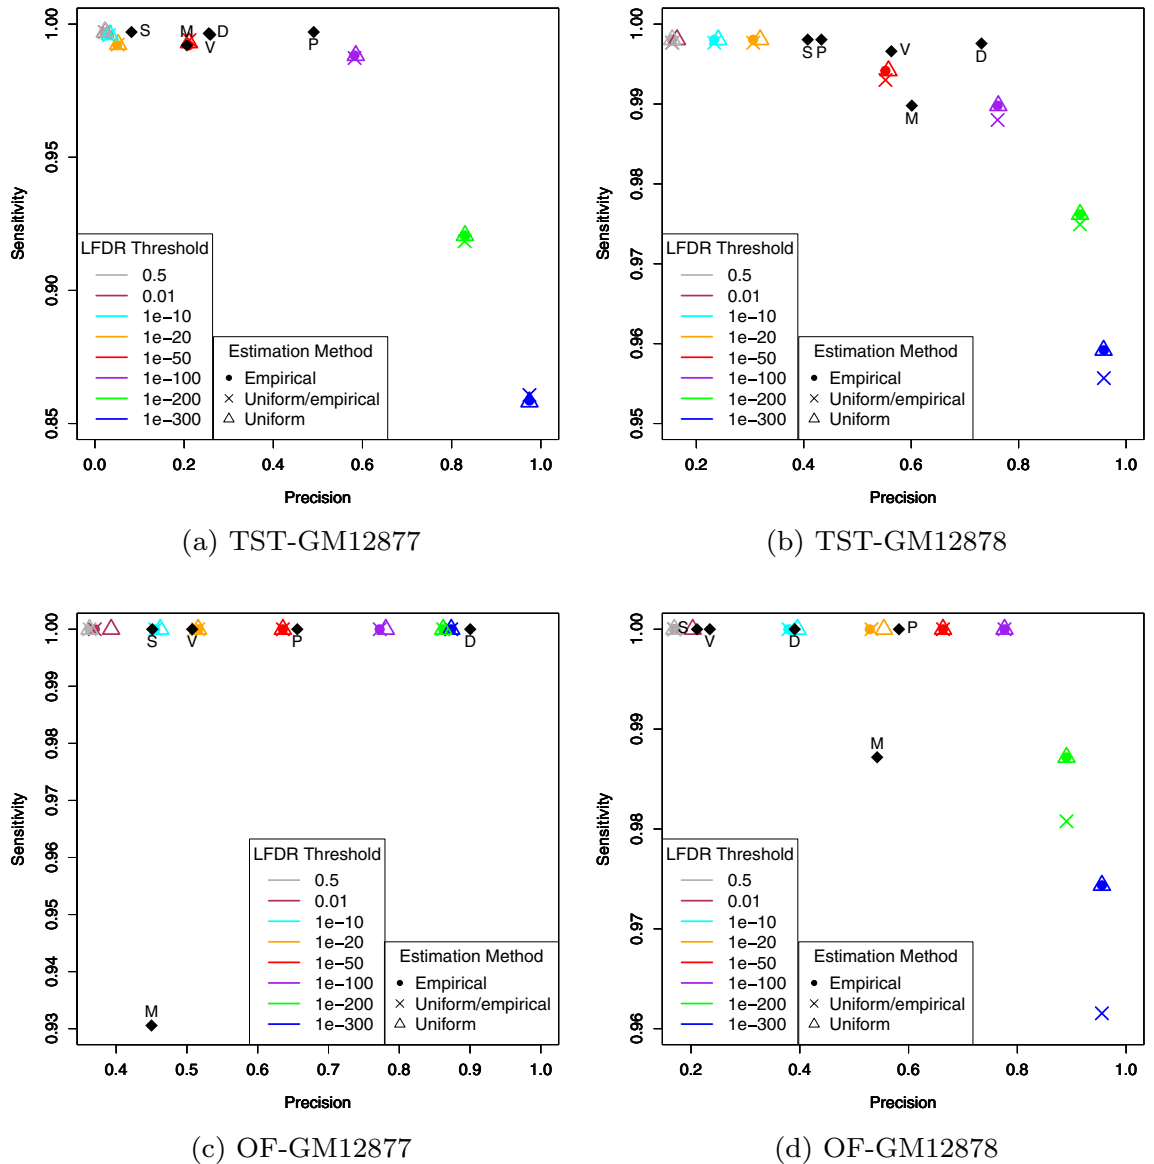


Figure 2. Average of *Prec* and *Sens* values over replicates based on calls made by the LFDR approach for different LFDR thresholds. Average of *Prec* and *Sens* values over replicates based on calls made by the other five variant callers were added for comparison purposes. M, V, S, D and P stand for MuTect2, VarScan2, SAMtools, VarDict and Pises, respectively.

thresholds, *Sens* values are mostly high while *Prec* is increased by decreasing the LFDR threshold. Indeed, when the LFDR threshold is high, say 0.5 for example, many false positives are allowed to be in the list of variants called. However, a strict LFDR threshold, say 10^{-50} for example, leads to few(er) false positives and consequently an improved *Prec* is gained. All the three estimation methods (empirical, uniform/empirical and uniform) performed nearly the same. Especially, when the LFDR threshold is below 10^{-50} , the three methods led to almost the same *Prec* and *Sens* values.

Comparing with the other variant callers, we observe that the LFDR approach has equal or better *Sens* for some LFDR thresholds compared to all other algorithms. And for other thresholds, it has equal or better *Prec*. For example, from panel (a) in Fig. 2 we observe that the LFDR approach with a threshold of 10^{-100} outperforms all the other five variant callers in terms of *Prec*. However, it reports a smaller *Sens*. When looking at other panels, slightly different but remarkable performance is observed. For example, by choosing the same LFDR threshold, we observe in panel (b) that the LFDR approach outperforms MuTect2 in terms of *Prec*, and has almost the same *Sens*, and outperforms the remaining four variant callers in terms of *Prec*. However, its *Sens* is a bit smaller than the others. But when looking at panel (d), we observe that the LFDR approach with the same LFDR threshold outperforms all the other variant callers in terms of both *Prec* and *Sens*. We observe that having a varying LFDR threshold allows for a wide range of tradeoff between *Sens* and *Prec*.

Figure 3 represents convergence of estimated π_0 in replicate 1 of each dataset for different LFDR thresholds. From this figure, we observe that the algorithm converged mostly in 3-7 iterations, and most remarkably, for

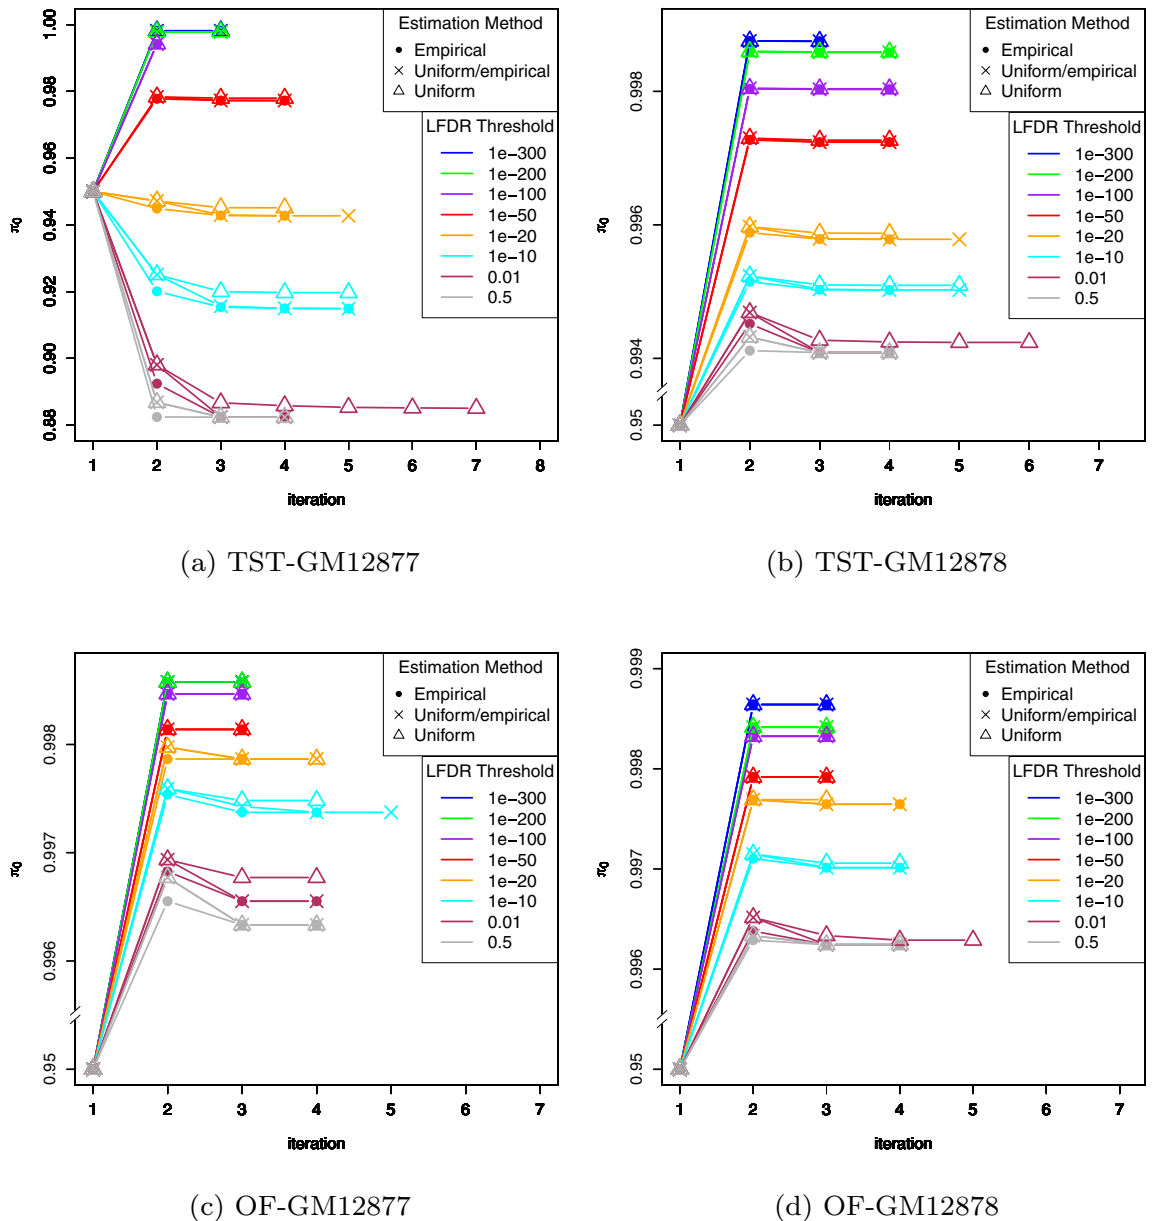


Figure 3. Estimated π_0 in replicate 1 of each dataset for different values of LFDR thresholds. The solid grey line represents the true π_0 .

stringent LFDR thresholds (10^{-100} or less), it converged in at most 4 iterations. We also note that in each panel, those stringent LFDR thresholds resulted in the greatest accuracy in estimating π_0 . In fact, the lower is the LFDR threshold, the more accurate is the corresponding estimated π_0 . Comparing with Fig. 2, we also observe that the most accurate estimated π_0 leads to the highest *Prec*. However, this is not necessarily true when looking at *Sens* values.

We also investigated the impact of each method on the magnitude of LFDR estimates. Figure 4 represents estimated LFDRs for replicate 1 of each of our datasets based on taking 10^{-300} as the LFDR threshold. This figure reflects that both the empirical and uniform/empirical approaches coincide on detecting the same variants. The corresponding estimated LFDRs are mostly either close to zero (that corresponds to mutant sites) or one (that corresponds to non-mutant sites), revealing that both approaches were able to nicely categorize the sites as either mutant or non-mutant sites. But, LFDRs estimated using the uniform approach did not follow this structure and there are many sites with medium estimated LFDRs. This concludes that both the empirical and uniform/empirical approaches outperform the uniform approach in classifying sites to non-mutant as well mutant categories.

Following the steps in Algorithm 2, we calculated the LFDR values for variants detected by the five variant callers (MuTect2, VarScan2, SAMtools, VarDict and Pises) across replicate 1 of all the datasets. Figure 5 represents the corresponding LFDRs and whether the variants are TPs or FPs, along with the number of TPs and FPs. From the figure we observe that all those TPs led to either zero or very close to zero estimated LFDR values, as expected. For example, we observe from panel (a) that of the total 1532 variants detected by MuTect2,

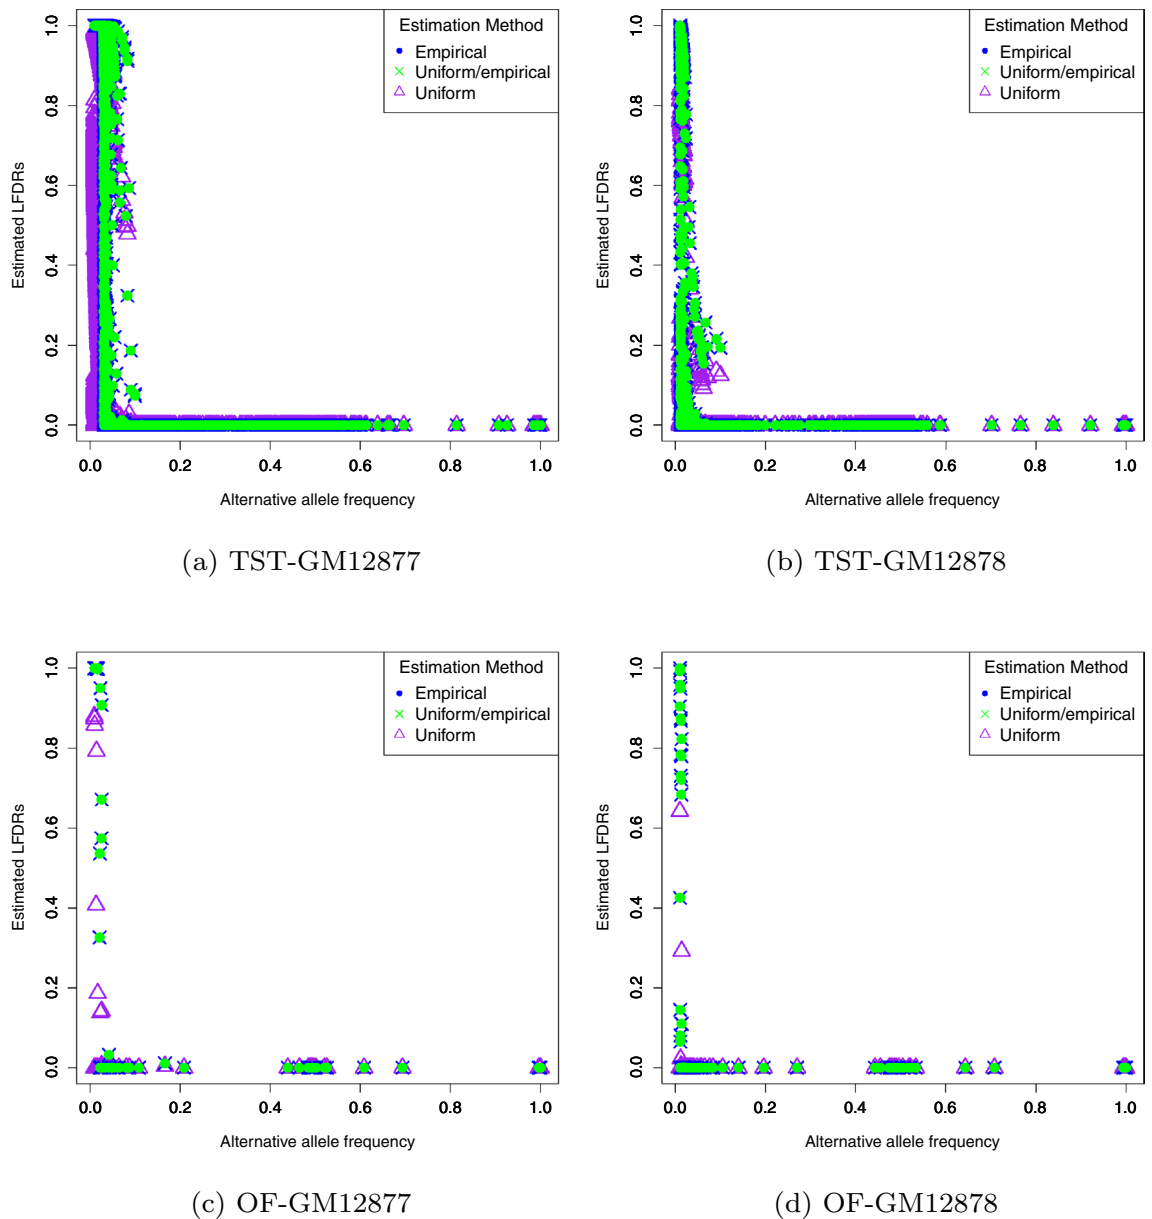


Figure 4. Estimated LFDRs for replicate 1 of each dataset based on taking 10^{-300} as the LFDR threshold.

334 variants with estimated LFDR values of zero or very close to zero are TPs, and 1198 variants with estimated LFDRs varying between zero and one are FPs. Comparing these numbers with 336, the number of known variants expected to be present in TST-GM12878 data, we realize that imposing a small LFDR threshold would dramatically reduce the number of FPs, and consequently *Prec* gets significantly improved. The lower is the LFDRs threshold, the more FPs are eliminated from the list of variants. Therefore, the proposed algorithm, with less complication compared to Algorithm 1, can prioritize results of other variant callers by simply eliminating FPs based on their LFDR values.

Naturally, prioritizing another variant caller's output, and potentially deciding that some of those calls are below threshold, cannot increase *Sens*. *Sens* increase could only be achieved by adding in missed true variants. However, prioritization and thresholding has the potential to eliminate many FPs, and thus increase *Prec*, hopefully with little or no loss to *Sens*.

We applied this prioritization approach to all calls made by the five variant callers in all replicates in the four datasets. We then calculated average *Prec* and *Sens* values for different LFDR thresholds as shown in Fig. 6. For comparison purposes, we also included *Prec* and *Sens* values calculated for each individual variant caller in Fig. 2. From Fig. 6, we observe that prioritizing variants using the LFDR method successfully leads to a significant increase in *Prec* values of the variant callers with minimal loss of *Sens* values for any threshold chosen between 0.5 and 10^{-50} . As an example, from Fig. 6a, VarDict led to *Prec* and *Sens* values of 0.256 and 0.997, respectively, and prioritizing its calls using the LFDR approach with a threshold of 10^{-20} , led to 0.801 and 0.992 as new *Prec*

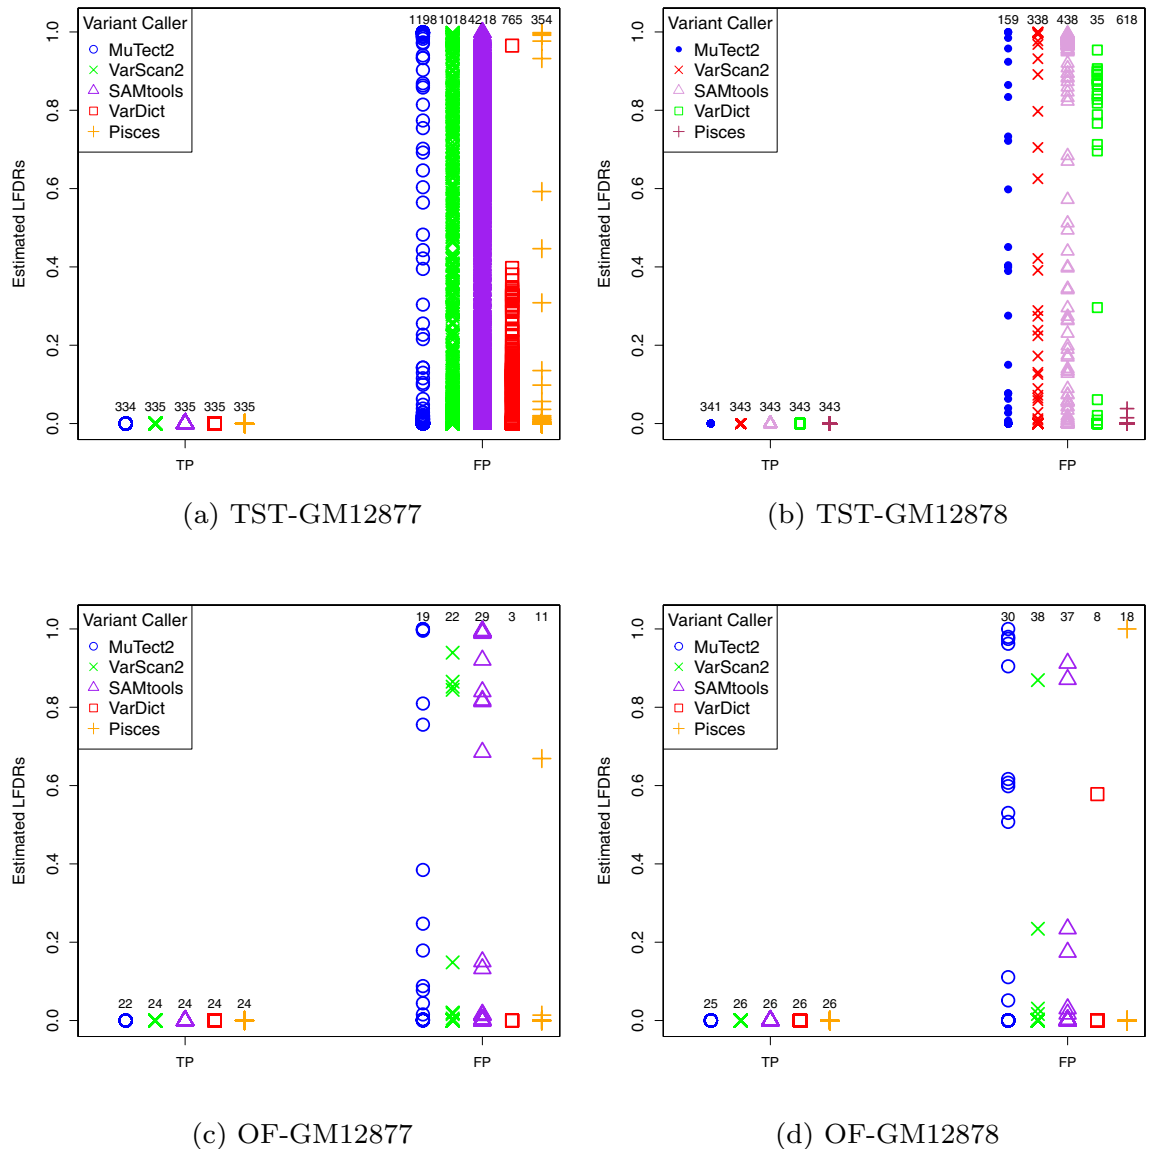


Figure 5. Estimated LFDRs for TP and FP calls made by the five variant callers in replicate 1 of each dataset.

and *Sens* values. Thus, LFDR prioritization appears to be a powerful approach to further filter the output of other variant-calling algorithms.

Discussion

In this paper, we introduced a novel LFDR-based approach that can be built into a SNV caller, or can be used to prioritize variants called by other algorithms. The algorithm requires some pre-determined numbers including the threshold $l_I / (l_I + l_{II})$ and an initial value for π_0 . Although, the algorithm allows for introducing different error rates through e_i at each genomic site to the model, users may set it to a global error rate $e = 0.01$ which corresponds to a base call quality of 20. In fact, the default error value in some well performed software packages including Pisces⁸ has been set to 0.01. Assuming a global error rate facilitates the speed of computation and our data analyses convinced us that it has an ignorable effect on the number of detected mutations.

Both of our LFDR algorithms have few parameters to adjust and thus it is easier to tune, compared to the five variant callers we used in this study. Like the five variant callers we used, our algorithm uses some pre-specified numbers for BQ , MQ and error rates e_i , but it does not require any other parameter to adjust, except the LFDR threshold. And yet as an advantage, the threshold can be controlled by users/analysts and is easy to interpret, due to the fact that the LFDR is a probability and varies only on the interval $[0, 1]$. The closer to zero (one) it is, the more confidence is gained in detecting variant (non-variant) sites. We considered 8 different LFDR thresholds, and noticed that, thresholds between 10^{-100} and 10^{-20} lead to better results in term of *Prec* and *Sens* values on the datasets we studied. To establish the generality of this observation, further testing would be needed on other datasets from different sources or of different types.

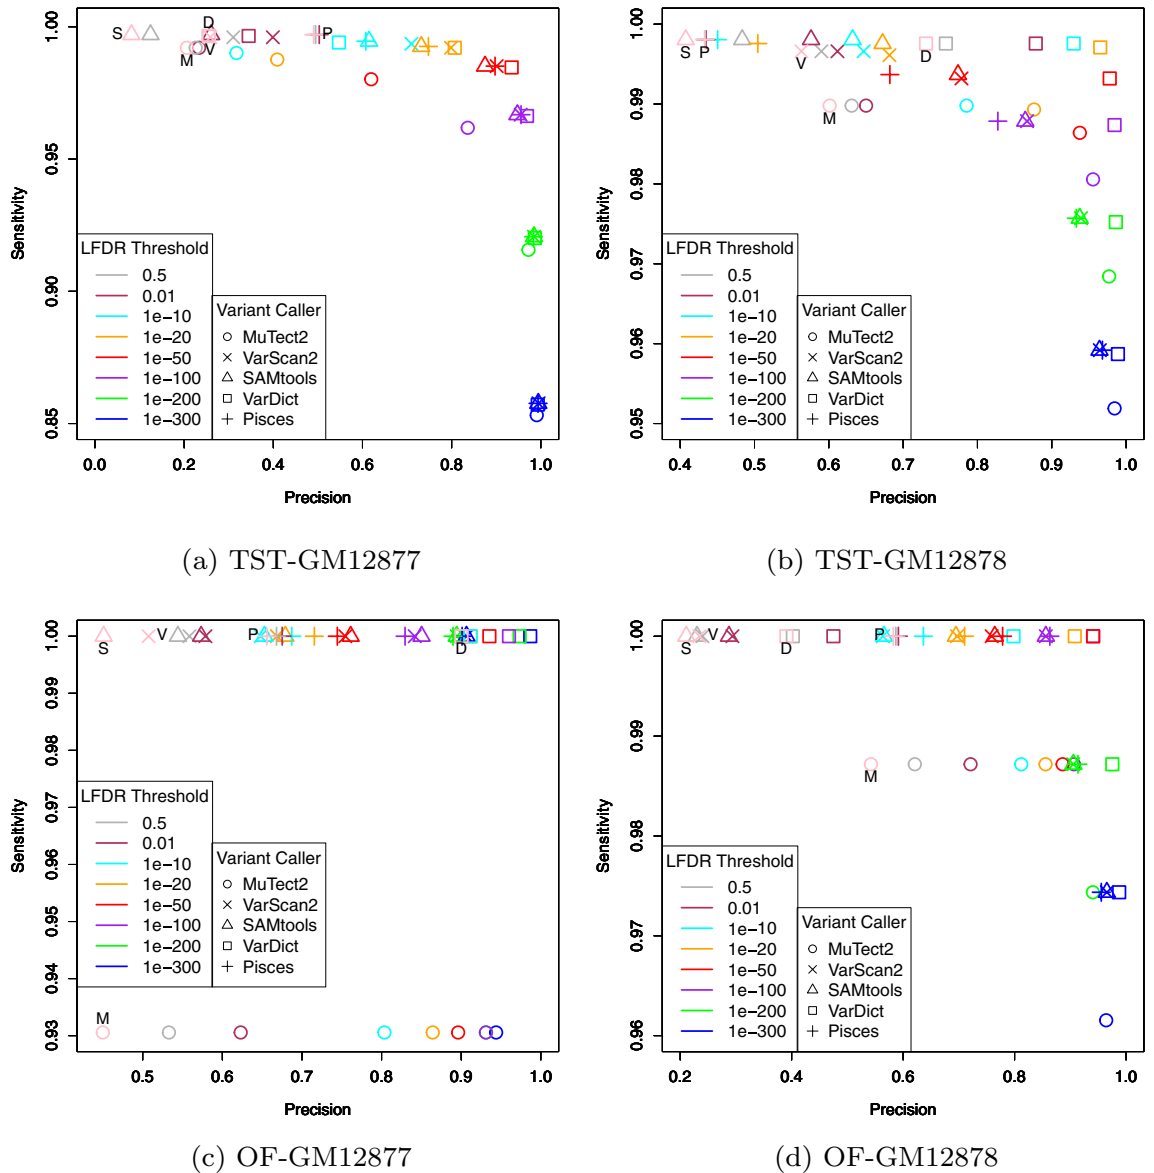


Figure 6. Average of *Prec* and *Sens* values for calls prioritized by the LFDR approach. The pink symbols refer to the *Prec* and *Sens* values of each individual variant caller calculated in Fig. 2.

We used three methods to estimate $g(\cdot)$ in Eq. (2). As expected, both the empirical and uniform/empirical approaches performed the same and identified the same variants. This is because both approaches lead to the same \mathcal{S} , set. The uniform approach also led to very similar results. This could be because we are ignoring the magnitude of LFDR values and we just compare them with the chosen thresholds. One possible explanation could also be the fact that the data we used in our project do not have any genuine mutations at low AFs, and thus, accurately estimating $g(\cdot)$ has little effect on the performance. However, as we showed in Fig. 4, the uniform approach is not able to firmly categorize sites to either mutant and non-mutants sites.

In this paper, we focused on germline SNV calling. However, as the proposed model learns the form of $g(\cdot)$ in Eq. (2), the proposed SNV calling approach may be used in a tumor-only mode, without comparison to healthy tissue from the same patient, which is often the case in clinical sequencing. The accuracy of the proposed method for tumore-only samples remains open to be investigated.

One may incorporate uncertainty about the parameter θ_i using other Bayesian frameworks. For example, one may compute the Bayes factor $BF = \frac{P(X_i|\theta_i=0)}{P(X_i|\theta_i>0)}$ and re-express the LFDR quantity defined in Eq. (3) as

$$\psi_i = \frac{\pi_0 BF}{\pi_0 BF + (1 - \pi_0)},$$

or

$$\psi_i = \frac{PO \times BF}{PO \times BF + 1},$$

where $PO = \frac{\pi_0}{1-\pi_0}$ can be interpreted as the prior odds of having no mutations, see Wakefield²¹ for more details. However, having the BF alone does not seem to be enough to identify variant sites and π_0 still needs to be somehow estimated. The beauty of our proposed approach is that it estimates both $P(X_i|\theta_i > 0)$ (and hence the Bayes factor) and π_0 without having any prior information regarding the distribution of θ_i .

We end our discussion with highlighting that in some situations, one may notice conflicts in calling a site a variant site. This may happen when multiple replicates are studied simultaneously and a variant caller detects a mutation at a specific site in one replicate and does not call that variant in the other replicate(s). One simple way to resolve this inconsistency would be to take the intersection of the variants called in all replicates. However, this may exclude some important variants (if borderline) and in general ignores the confidence any one replicate gives us. In an extreme case, if a replicate failed to have data in a region, and thus no variants were called, that lack of data would effectively overrule the positive data in other replicates. Alternatively, one may seek to take advantage of advanced decision-theoretic approaches. In a different but still applicable context, Karimnezhad and Bickel¹⁹ developed an empirical Bayesian approach that takes advantage of prior information from multiple reference classes and leads to a unique decision in conflicting situations. This could be a potential future research problem.

Data availability

All sequencing data is available from the Sequence Read Archive under project accession PRJNA614006. Human genome version hg19/GRCh37 is available from the UCSC Genome Browser website (genome.ucsc.edu).

Code availability

The R code implementing the proposed methods is available on Github <https://github.com/empiricalbayes/SNVLFDR>.

Received: 29 December 2022; Accepted: 11 January 2024

Published online: 18 January 2024

References

- Wong, K. M., Hudson, T. J. & McPherson, J. D. Unraveling the genetics of cancer: Genome sequencing and beyond. *Annu. Rev. Hum. Genet.* **12**, 407–430 (2011).
- Morgensztern, D., Devarakonda, S., Mitsudomi, T., Maher, C. & Govindan, R. Mutational events in lung cancer: Present and developing technologies. In *IASLC Thoracic Oncology* (2nd edn) 95–103 (Elsevier, 2018).
- Li, H. *et al.* The sequence alignment/map format and SAMtools. *Bioinformatics* **25**(16), 2078–2079 (2009).
- Koboldt, D. C. *et al.* VarScan 2: Somatic mutation and copy number alteration discovery in cancer by exome sequencing. *Genom. Res.* **22**(3), 568–576 (2012).
- Koboldt, D. C. *et al.* VarScan: Variant detection in massively parallel sequencing of individual and pooled samples. *Bioinformatics* **25**(17), 2283–2285 (2009).
- Cibulskis, K. *et al.* Sensitive detection of somatic point mutations in impure and heterogeneous cancer samples. *Nat. Biotechnol.* **31**, 213–219 (2013).
- Lai, Z. *et al.* VarDict: A novel and versatile variant caller for next-generation sequencing in cancer research. *Nucleic Acids Res.* **44**(11), e108. <https://doi.org/10.1093/nar/gkw227> (2016).
- Dunn, T., Berry, G., Emig-Agius, D., Jiang, Y., Iyer, A. *et al.* Pisces: An accurate and versatile single sample somatic and germline variant caller. In *Proceedings of the 8th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics* 595–595 (ACM, 2017).
- Karimnezhad, A. *et al.* Accuracy and reproducibility of somatic point mutation calling in clinical-type targeted sequencing data. *BMC Med. Genom.* **13**(1), 1–14 (2020).
- Xu, C. A review of somatic single nucleotide variant calling algorithms for next-generation sequencing data. *Comput. Struct. Biotechnol. J.* **16**, 15–24 (2018).
- Huang, G., Wang, S., Wang, X. & You, N. An empirical Bayes method for genotyping and SNP detection using multi-sample next-generation sequencing data. *Bioinformatics* **32**(21), 3240–3245 (2016).
- Zhao, Z., Wang, W. & Wei, Z. An empirical Bayes testing procedure for detecting variants in analysis of next generation sequencing data. *Ann. Appl. Stat.* 2229–2248 (2013).
- Ding, W., Kou, Q., Wang, X., Xu, Q. & You, N. Single-sample SNP detection by empirical Bayes method using next-generation sequencing data. *Stat. Interface* **8**(4), 457–462 (2015).
- Pan, W., Lin, J. & Le, C. T. A mixture model approach to detecting differentially expressed genes with microarray data. *Funct. Integr. Genom.* **3**(3), 117–124 (2003).
- Efron, B., Tibshirani, R., Storey, J. D. & Tusher, V. Empirical Bayes analysis of a microarray experiment. *J. Am. Stat. Assoc.* **96**(456), 1151–1160 (2001).
- Efron, B. *Large-scale inference: Empirical Bayes methods for estimation, testing, and prediction* (Cambridge University Press, 2012).
- Padilla, M. & Bickel, D. R. Estimators of the local false discovery rate designed for small numbers of tests. *Stat. Appl. Genet. Mol. Biol.* **11**(5), 515. <https://doi.org/10.1515/1544-6115.1807> (2012).
- Yang, Y., Aghababazadeh, F. A. & Bickel, D. R. Parametric estimation of the local false discovery rate for identifying genetic associations. *IEEE/ACM Trans. Comput. Biol. Bioinform.* **10**, 98–108 (2013).
- Karimnezhad, A. & Bickel, D. R. Incorporating prior knowledge about genetic variants into the analysis of genetic association data: An empirical Bayes approach. *IEEE/ACM Trans. Comput. Biol. Bioinform.* **17**(2), 635–646 (2020).
- Karimnezhad, A. A simple yet efficient parametric method of local false discovery rate estimation designed for genome-wide association data analysis. *Stat. Methods Appl.* **31**, 159–180 (2022).
- Wakefield, J. A Bayesian measure of the probability of false discovery in genetic epidemiology studies. *Am. J. Hum. Genet.* **81**(2), 208–227 (2007).
- Eberle, M. A. *et al.* A reference dataset of 5.4 million phased human variants validated by genetic inheritance from sequencing a three-generation 17-member pedigree. *Genome Res.* **27**(1), 157–164 (2017).
- Raczy, C. *et al.* Isaac: ultra-fast whole-genome secondary analysis on illumina sequencing platforms. *Bioinformatics* **29**(16), 2041–2043 (2013).

Acknowledgements

The authors wish to thank two anonymous reviewers and the Editor for their valuable comments and suggestions. The authors wish to acknowledge the technical support of the Ottawa Hospital Research Institute/University of Ottawa Bioinformatics Core Facility and StemCore Laboratories for their expert technical assistance. The sequencing data was generated by StemCore Laboratories at the Ottawa Hospital Research Institute. This Project is funded by the Government of Canada through Genome Canada and Ontario Genomics, and by the Government of Ontario, under Genome Canada Grant GAPP 6450. This work was supported by NSERC Discovery Grant RGPIN-2019-06604 to TJP.

Author contributions

AK and TJP designed the study and wrote the paper. AK analyzed the data. The authors reviewed and approved the final manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to A.K.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© Crown 2024