



A Super Scalable Algorithm for Short Segment Detection

Ning Hao¹ · Yue Selena Niu¹ · Feifei Xiao² · Heping Zhang³

Received: 30 August 2019 / Revised: 23 February 2020 / Accepted: 29 March 2020
© International Chinese Statistical Association 2020

Abstract

In many applications such as copy number variant (CNV) detection, the goal is to identify short segments on which the observations have different means or medians from the background. Those segments are usually short and hidden in a long sequence and hence are very challenging to find. We study a super scalable short segment (4S) detection algorithm in this paper. This nonparametric method clusters the locations where the observations exceed a threshold for segment detection. It is computationally efficient and does not rely on Gaussian noise assumption. Moreover, we develop a framework to assign significance levels for detected segments. We demonstrate the advantages of our proposed method by theoretical, simulation, and real data studies.

Keywords Copy number variation · Inference · Nonparametric method · Signal detection

1 Introduction

Chromosome copy number variant (CNV) is a type of structural variation with abnormal copy number changes involving DNA fragments [5, 6]. CNVs result in gains or losses of the genome, therefore interfering downstream functions of the DNA contents. Accounting for a substantial amount of genetic variation, CNVs are considered to be a risk factor for human diseases. Over the past decade, advances in genomic technologies have revealed that CNVs underlie many human diseases,

✉ Heping Zhang
heping.zhang@yale.edu

Ning Hao
nha@math.arizona.edu

¹ Department of Mathematics, University of Arizona, Tucson, AZ 85721, USA

² Department of Epidemiology and Biostatistics, University of South Carolina, Columbia, SC 29201, USA

³ Department of Biostatistics, Yale School of Public Health, New Haven, CT 06510, USA

including autism [15], cancer [4], schizophrenia [3], and major depressive disorder [13]. It is fundamental to develop fast and accurate CNV detection tools.

A variety of statistical tools have been developed to discover structural changes in CNV data during last 20 years. Popular algorithms include circular binary segmentation [14], the fused LASSO [16], likelihood ratio selection [10], and screening and ranking algorithm [11]. Some other change-point detection tools such as wild binary segmentation [8] and simultaneous multiscale change-point estimator [7] can be also applied to CNV data. See [12] for a recent review on modern change-point analysis techniques. A majority of existing methods are based on Gaussian assumption, although quantile normalization [18] or local median transformation [2] can be used for normalization. The computational complexity is also of concern for some of the existing methods as the modern technologies produce extraordinarily big data. In spite of some fast algorithms [17], few algorithms are known to possess both computational efficiency and solid theoretical foundation. Moreover, with only a few exceptions [7, 9], the existing methods focus on detection whereas not offering statistical inference. For these reasons, we develop a fast nonparametric method for CNV detection with theoretical foundation and the opportunity of conducting statistical inference.

In this paper, we model the CNVs as short segments with nonzero height parameters, which are sparsely hidden in a long sequence. The goal is to identify those segments with high probability and, moreover, to assess the significance levels for detected segments. In particular, we propose a scalable nonparametric algorithm for short segment detection. It depends on only the ranks of the absolute values of the measurements and hence requires minimal assumptions on the noise distribution. A short segment may be present when there are a large enough number of observations exceeding a certain threshold on a short segment; for instance, 8 on a segment of 10 observations are larger than the 99th percentile of the data. Following this idea, we implement a super scalable short segment (4S) detection algorithm to cluster the points to form a segment when such a phenomenon occurs. The advantages of our method are fourfold. First, this nonparametric method requires minimal assumption on the noise distribution. Second, it is super fast as the core algorithm requires only $O(n)$ operations to analyze a sequence of n measurements. In particular, it takes less than 2 s for our R codes to analyze 272 sequences with a range of about 34,000 measurements. Third, we establish a non-asymptotic theory to ensure the detection of all signal segments. Last but not least, our method can compute the significance level for each detected segment and offer a convenient approach to statistical inference.

2 Method

2.1 Notations and the Main Idea

Let $\{X_j\}_{j=1}^n$ be a sequence of random variables such that

$$X_j = \mu_j + \varepsilon_j, \quad (1)$$

where the *height parameter* vector $\mu = (\mu_1, \dots, \mu_n)^\top$ is sparse and $\{\varepsilon_j\}_{j=1}^n$ are random noises with median 0. Moreover, we assume that the nonzero entries of μ are supported in the union of disjoint intervals $\mathcal{I} = \bigcup_{k=1}^K I_k$ so that

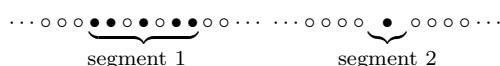
$$\mu_j = \begin{cases} v_k \neq 0, & \text{if } j \in I_k \subset \mathcal{I} \text{ for some } k; \\ 0, & \text{if } j \notin \mathcal{I}. \end{cases}$$

Here we assume that $I_k = \{\ell_k, \ell_k + 1, \dots, r_k\}$ and $r_k < \ell_{k+1} - 1$ for all k . Note that such a representation of \mathcal{I} is unique and used throughout this paper for \mathcal{I} or its estimator $\hat{\mathcal{I}}$. For convenience and without confusion, we use the interval $[\ell_k, r_k]$ to imply the set of integers $\{\ell_k, \ell_k + 1, \dots, r_k\}$ when referred to a set of locations. We call those intervals *segments*. In particular, a *signal segment* is a segment where the height parameter is a nonzero constant. Let μ_I denote a subvector of μ restricted to $I \subset \{1, \dots, n\}$. We denote by $|\mathcal{S}|$ the cardinality of a set \mathcal{S} . In particular, $|I| = r - \ell + 1$ for $I = [\ell, r]$. $\mathbf{0}$ and $\mathbf{1}$ denote vectors, $(0, \dots, 0)^\top$ and $(1, \dots, 1)^\top$, respectively.

Naturally, a primary goal for model (1) is to identify the set of signal segments $\{I_k\}_{k=1}^K$. Moreover, while rarely done, it is useful to assign a significance level for each of the detected segments. In this paper, we will study both estimation and related inference problems on segment detection. Our strategy is to cluster ‘putative locations’ using spatial information. In particular, we consider the set of positions $\mathcal{S}_c = \{j : |X_j| > c, 1 \leq j \leq n\}$, where the observations exceed a threshold $c > 0$. Intuitively, for a properly chosen c and some segment \hat{I} , if $|\hat{I} \cap \mathcal{S}_c|$ is big enough compared with $|\hat{I}|$, it is likely that $\mu_{\hat{I}} \neq \mathbf{0}$.

To illustrate our idea, we consider a game of ball painting. Suppose that we start with n white balls in a row, and paint the j th ball with black color if $|X_j| > c$. Let m be the total number of black balls, which is much smaller than n . If we observe that there are a few black balls crowded in a short segment, e.g., ‘segment 1’ illustrated in Fig. 1, it is plausible that the height parameter is not zero in the segment. Our proposed algorithm can easily identify those segments. Also it may happen that in a neighborhood there is only a single black ball, e.g., ‘segment 2’ in Fig. 1. Then, we may not have strong evidence against $\mu = 0$. To put this intuition into a sound theoretical framework, it is imperative to evaluate the significance for each pattern. In fact, given the numbers of white and black balls in a short segment, we may calculate how likely a certain pattern appears in a sequence of length n with m black balls, when white and black balls are actually randomly placed. We will develop a framework of inference based on this idea in Sect. 2.3.

Fig. 1 Segment 1 and segment 2



2.2 Algorithm

To estimate \mathcal{I} , we propose a super scalable short segment (4S) detection algorithm which is described as follows.

Step 1: thresholding. Define $\mathcal{S}_c = \{j : |X_j| > c\}$. That is, we collect the positions where the observations exceed a threshold.

Step 2: completion. Construct the completion set $\mathcal{S}_{c,d}$ by the criterion that $j \in \bar{\mathcal{S}}_c$ if and only if there exist $j_1, j_2 \in \mathcal{S}_c$, $j_1 \leq j_2 \leq j_1 + d + 1$ such that $j_1 \leq j \leq j_2$. That is, we add the whole segment $[j_1, j_2]$ into the completion set if the gap between $j_1, j_2 \in \mathcal{S}_c$ is small enough.

Write $\bar{\mathcal{S}}_{c,d} = \bigcup_{k=1}^K \tilde{I}_k$ where $\tilde{I}_k = [\tilde{\ell}_k, \tilde{r}_k]$ with $\tilde{r}_k < \tilde{\ell}_{k+1} - 1$. Note that this decomposition is unique.

Step 3: clean up. We delete \tilde{I}_k from $\bar{\mathcal{S}}_{c,d}$ if $|\tilde{I}_k| = \tilde{r}_k - \tilde{\ell}_k + 1 \leq h$, and obtain our final estimator $\hat{\mathcal{I}}_{c,d,h} = \bigcup_{k=1}^K \hat{I}_k$. That is, we ignore the segments that are too short to be considered.

The whole procedure depends on three parameters c , d , and h . The choice of c is crucial and depends on applications. d and h are relatively more flexible as we can screen false positives using significance levels defined later. We may ignore the subscripts and simply refer to \mathcal{S} , $\bar{\mathcal{S}}$ and $\hat{\mathcal{I}}$ when the sets obtained from the three steps above corresponding to c , d and h are clear in the context. Figure 2 illustrates our procedure, where the location set obtained in each step is indicated by the positions of black balls.

2.3 Theory: Consistency and Inference

Our goal is to identify the set of signal segments $\mathcal{I} = \bigcup_{k=1}^K I_k$ with a false positive control. Here we say that $I_k \in \mathcal{I}$ is identified by an estimator $\hat{\mathcal{I}} = \bigcup_{k=1}^K \hat{I}_k$ if there is a unique $\hat{I}_{k'} \in \hat{\mathcal{I}}$, such that $\hat{I}_{k'} \cap I_k \neq \emptyset$, and $\hat{I}_{k'} \cap I_j = \emptyset$ for all $I_j \in \mathcal{I}$ and $j \neq k$. Such an $\hat{I}_{k'}$ is a *true positive*. We define that \mathcal{I} is identified by an estimator $\hat{\mathcal{I}}$ if every $I_k \in \mathcal{I}$ is identified by $\hat{\mathcal{I}}$. That is, there is a one-to-one correspondence between \mathcal{I} and a subset of $\hat{\mathcal{I}}$, and the K pairs under this correspondence are the only pairs with nonempty interaction among all segments in \mathcal{I} and $\hat{\mathcal{I}}$. See Fig. 3 for an illustration of the definition.

$\mathcal{S} :$...○○○●●○●○●○○...●○○○●●○○○...
$\bar{\mathcal{S}} :$...○○○●●●●●○○...●○○○●●○○○...
$\hat{\mathcal{I}} :$...○○○●●●●●○○...○○○○○○○○○...

Fig. 2 An illustration of three steps in the 4S algorithm with $d = h = 3$. Top row: the black balls indicate the locations where the observation has absolute value larger than a threshold. Middle row: the small gaps with length $\leq d$ between the black balls are filled in with the black balls. Bottom row: the segment of black balls with length $\leq h$ is deleted

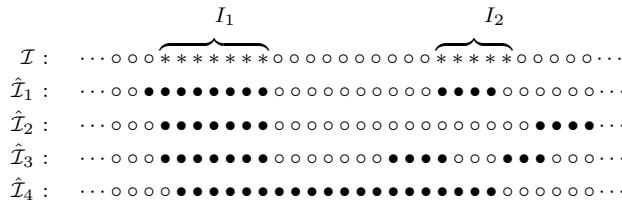


Fig. 3 An illustration of relationship between signal segments (*) and four estimators (•). \mathcal{I} consists of two signal segments I_1 and I_2 . Both I_1 and I_2 are successfully identified by $\hat{\mathcal{I}}_1$. I_1 is also identified by $\hat{\mathcal{I}}_2$ and $\hat{\mathcal{I}}_3$. $\hat{\mathcal{I}}_2$ has one true positive (left) and one false positive (right). $\hat{\mathcal{I}}_3$ has one true positive (left) and two false positives (middle and right). $\hat{\mathcal{I}}_4$ has one false positive

In our three-step procedure, the first two steps establish an estimator and the last one aims to delete obvious false positives. Our theory proceeds in two main steps. First, we characterize the non-asymptotic probability that the first two steps produce an estimator which successfully identifies \mathcal{I} . In order to identify $I_k \in \mathcal{I}$, we should ensure two conditions. Condition one is that, after step 1, the black balls are dense enough on I_k so that they do not split into two or more segments in step 2. Condition two is that, in the gap between I_k and I_{k+1} , the black balls are sparse enough so that the black balls on I_k and I_{k+1} do not connect to a big segment. Theorem 1 addresses how to bound the probabilities of these two conditions for all k .

Second, we develop a framework of inference to control false positives. As a rough control, Theorem 2 gives an upper bound for the expected number of false positives if all segments of length one are deleted in step 3. In general, after step 2, it is not optimal to decide the likelihood of a detected segment being a false positive only by its length. Therefore, for each segment in $\bar{\mathcal{S}}$, we check its original color pattern back in step 1 and calculate a p value of this pattern under null hypothesis $\mu = \mathbf{0}$. This assigns a significance level for each detected segment which helps control false positive. It is difficult to find the exact p values. Lemma 2 offers a reasonable approximation.

To facilitate theoretical analysis, we assume that, in this subsection, $\{\varepsilon_j\}_{j=1}^n$ is independent and identically distributed (IID) noises with median 0. Moreover, ε_j has a continuous density function f that is symmetric with respect to 0. Under this assumption, the black balls are randomly distributed for arbitrary threshold c when $\mu = \mathbf{0}$.

Now we investigate when a signal segment can be detected by our algorithm. Let F be the cumulative distribution function of the noise density f . We use $f_\alpha = F^{-1}(\alpha)$ to denote the α -percentile. Suppose that there is a segment I such that $|I| = L$, $\mu_I = v\mathbf{1}$ and $\mu_{I^c \cap H} = \mathbf{0}$, where H is a segment containing I such that $I^c \cap H$ is the union of two segments both of which are of length D . Without loss of generality, we assume that $v = f_\alpha > 0$ i.e., $\alpha > 0.5$. For a threshold $0 < c = f_\beta \leq v$, let us continue our game of ball painting and focus on this segment and its neighborhood. Recall that we paint the ball at position j with black color if and only if $|X_j| > c = f_\beta$. The following two events together can ensure that the segment I is identified by our method.

$$\mathcal{A} = \{\text{On } I, \text{ there does not exist a sub-segment of } \min\{d, L\} \text{ consecutive white balls}\}$$

$$\mathcal{B}_D = \{\text{On both sides of } I, \text{ within distance } D, \text{ there are } d \text{ consecutive white balls}\}$$

Event \mathcal{A} ensures that this segment can be detected as a whole segment while event \mathcal{B}_D controls the total length of the detected segment and makes sure that the detected segments are separated from each other. \mathcal{A} and \mathcal{B}_D together guarantee that our algorithm identifies a segment \hat{I} such that $\hat{I} \cap I \neq \emptyset$, and $\hat{I} \cap I' = \emptyset$ for any other signal segment I' . The following lemma gives non-asymptotic bounds for $P(\mathcal{A})$ and $P(\mathcal{B}_D)$.

Lemma 1 *Let $H \supset I$ be two segments such that $\mu_I = v\mathbf{1}$ and $\mu_{I \cap H} = \mathbf{0}$, $|I| = L$. $I^c \cap H$ is the union of two segments both of which are of length D . Let $\beta' = 2\beta - 1$. For $0 < c = f_\beta \leq v$ and $d > 0$, after the thresholding and completion steps, we have*

$$P(\mathcal{A}) \geq \begin{cases} 1 - (L - d + 2)2^{-d-1}, & \text{if } d < L \\ 1 - 2^{-L}, & \text{if } d \geq L \end{cases}$$

and

$$P(\mathcal{B}_D) \geq 1 - 2(1 - \beta'd)^{\lfloor \frac{D}{d} \rfloor}.$$

Let $v_{\min} = \min_k |v_k|$ be the minimal signal strength among all I_k 's, $L_{\min} = \min_k |I_k|$, $L_{\max} = \max_k |I_k|$ be the minimal and maximal lengths of signal segments, respectively, and $D_{\min} = \min_k (\ell_{k+1} - r_k - 1)$ be the minimal gap between two signal segments. Define $\beta_{\min} = F(v_{\min})$ so $v_{\min} = f_{\beta_{\min}}$. Let $\beta'_{\min} = 2\beta_{\min} - 1$. Taking into account all signal segments in \mathcal{I} , the theorem below gives a lower probability bound for identifying \mathcal{I} after first two steps.

Theorem 1 *With $c = v_{\min} = f_{\beta_{\min}}$, $d > 0$ and $h = 0$, $\hat{\mathcal{I}}_{c,d,h}$ can identify all signal segments in \mathcal{I} with probability at least*

$$1 - K \max \left\{ \frac{1}{2}(L_{\max} - d + 2), 1 \right\} 2^{-\min\{d, L_{\min}\}} - (K - 1)(1 - \beta'_{\min}d)^{\left\lfloor \frac{D_{\min}}{d} \right\rfloor}. \quad (2)$$

Corollary 1 *The probability (2) goes to 1 asymptotically if $\log K + \log L_{\max} \ll \min\{d, L_{\min}\} \rightarrow \infty$ and $\log K \ll D_{\min}/d \rightarrow \infty$ as $n \rightarrow \infty$.*

Although Theorem 1 gives a theoretical guarantee to recover all signal segments with a large probability, there are some false positives. As an ad-hoc way, we may take $h = 2$ or 3 to eliminate some obvious false positives. This clean up step is simple and helpful to delete isolated black balls. The Theorem below gives an upper bound on the number of false positives with a conservative choice $h = 1$.

Theorem 2 *Assume $\mu = 0$ and $|\mathcal{S}_c| = m$. Then $E|\hat{\mathcal{I}}_{c,d,h}| \leq m(1 - \prod_{k=1}^d \frac{n-m+1-k}{n-k})$ with $h = 1$.*

The expected number of false positive segments can be well controlled if both m/n and d are small. Next, we illustrate how to access the significance levels for the

detected segments by our method, which is helpful to control false positives. Recall our estimator $\hat{\mathcal{I}} = \bigcup_{k=1}^K \hat{I}_k$ where $\hat{I}_k = [\hat{\ell}_k, \hat{r}_k]$. For each \hat{I}_k , let $s_k = |\hat{I}_k|$, $t_k = |\hat{I}_k \cap \mathcal{S}_c|$, and $m = |\mathcal{S}_c|$. Now consider n balls in a row with m black and $n - m$ white balls. Let $\mathcal{A}_{n,m,s_k,t_k}$ be an event that there exists a segment of length s_k where at least t_k balls are black, in a sequence of n balls with m blacks ones. The p value of \hat{I}_k can be defined as the probability of $\mathcal{A}_{n,m,s_k,t_k}$ if the balls are randomly placed. This p value can effectively control the false positives. However, it is challenging to find the exact formula to calculate the p value. The lemma below gives an upper bound of $P(\mathcal{A}_{n,m,s_k,t_k})$.

Lemma 2 $P(\mathcal{A}_{n,m,s_k,t_k}) \leq \hat{P}(\mathcal{A}_{n,m,s_k,t_k}) = mP(Y \geq t_k - 1)$ where Y follows a hypergeometric distribution with total population size $n - 1$, number of success states $m - 1$, and number of draws $s_k - 1$.

This approximated p value is useful to eliminate false positives.

2.4 Implementation

Our proposed method is nonparametric and depends on only the rank of absolute measurements $\{|X_j|\}_{j=1}^n$. For a fixed triplet (c, d, h) , it typically needs less than $3n$ operations to determine $\hat{\mathcal{I}}$ when $|\mathcal{S}_c|/n$ is small, say, less than 0.1. We need $2n$ operations to compare each measurement with the threshold to determine \mathcal{S}_c . Let $\mathbf{w} = (w_1, \dots, w_m)^\top$ be a vector of locations in \mathcal{S}_c in an ascending order. In the completion step, we compare $\Delta\mathbf{w} = (w_2 - w_1, \dots, w_m - w_{m-1})^\top$ to a threshold. In particular, we declare that w_i and w_{i+1} belong to different segments if and only if $w_{i+1} - w_i > d$. Let $i_1, \dots, i_{\bar{K}}$ be those indices such as $w_{i_k+1} - w_{i_k} > d$. $\bar{\mathcal{S}}$ consists of segments $[w_1, w_{i_1}], \dots, [w_{i_{\bar{K}}}, w_m]$. We record the start and end points of each segment only. In the deletion step, we delete $[w_{i_k}, w_{i_{k+1}-1}]$ if its length is not greater than h . The total operations can be controlled within $2n + 10m$.

The choice of threshold c is crucial to the 4S algorithm and may need to be determined on a case-by-case basis. Here, we offer a general guideline for parameter selection. Recall that for the Gaussian model, the signal strength of a segment with length L and height $\nu = \delta\sigma$ is usually measured by $S = \delta^2 L$; see, e.g., Table 1 in [12]. If there are two segments with the same overall signal strength S , however, one with large δ and small L (say, type A), and another one with small δ and large L (say, type B), then it is usually not equally easy to detect both of them by an algorithm of complex $O(n)$. Indeed, for many segment detection algorithms, it is tricky to balance the powers to detect these two types of segments. Intuitively, the threshold parameter c controls this tradeoff in our methods. A higher threshold may be more powerful in detecting type A segments but less powerful in detecting type B segments and vice versa. In practice, we may choose the threshold as a certain sample percentile of the absolute values of the observations based on a pre-specified preference. For example, if we know the signal segments have relatively large height parameters but can be as short as 5 data points, then with a fixed n , we can find largest m such as $\hat{P}(\mathcal{A}_{n,m,5,5}) \leq 0.05$ and the percentile is chosen as $\alpha = 1 - \frac{m}{n}$. That is, we want to guarantee that a segment of 5 consecutive black balls is significant enough to

stand out. In another scenario, our preference might be longer segments with possibly lower heights. Then we may choose a threshold to include segments of length 10 with at least 6 black balls. In general, such m (or α) can be easily determined given n , s , t and p , by solving $\hat{P}(\mathcal{A}_{n,m,s,t}) \leq p$. We illustrate in Fig. 4 the relationship between $\log n$ and selected percentile α for $(s, t) = (5, 5), (10, 6)$, $p = 0.05$ and 0.1 . Because our main goal in this paper is to identify short segments, we prefer a large threshold such as 95th sample percentile. An even larger threshold can be used to identify shorter segments, and a smaller threshold can be used for detecting longer segments with lower heights.

3 Numerical Studies

3.1 Simulated Data

We use simulation studies to evaluate the performance of our method in terms of the average number of true positives (TP) and false positives (FP) for identifying signal segments. Recall that in our definition, a detected segment $\hat{I} \in \hat{\mathcal{I}}$ is a true positive, if it interacts with only one signal segment $I \in \mathcal{I}$, and it is the only one in $\hat{\mathcal{I}}$ that interacts I .

In Example 1, we show the effectiveness of our inference framework on the false positive control of the 4S algorithm by a null model. As suggested by Fig. 4, we choose the 95th percentile of absolute values of the observations as the threshold c . We set $d = 9$ and $h = 3$. We use various p value thresholds for false positive control and compare them with a vanilla version of 4S, which is the one without p value control.

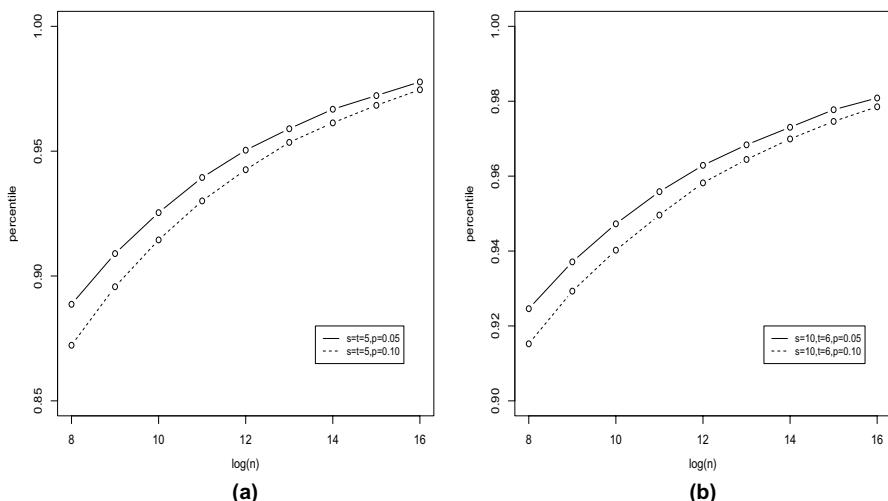


Fig. 4 Selected percentile versus $\log n$ for **a** $s = t = 5$, $p = 0.05$ and 0.10 ; **b** $s = 10$, $t = 6$, $p = 0.05$ and 0.10

Example 1 (Null Model) We generate a sequence based on model (1) with $n = 10,000$ and $\mu = \mathbf{0}$. We consider three scenarios for the error distributions. In the first two scenarios, we consider $\{\epsilon_i\}_{i=1}^n$ which are IID from $N(0, 1)$ and t_3 , respectively. In the last scenario, we consider $\{\epsilon_i\}_{i=1}^n$ which are marginally $N(0, 1)$ and jointly from an autoregressive (AR) model with autocorrelation 0.2.

As $\mathcal{I} = \emptyset$ in this example, all detected segments are FPs. We report the average FPs for three versions of 4S (Vanilla, $p = 0.05$ and $p = 0.1$) based on 100 replicates in Table 1. We see that our inference framework can effectively control the number of FPs.

In Example 2, we compare 4S with three algorithms CBS [14], LRS [10] and WBS [8]. The CBS and WBS methods, implemented by R packages DNAcopy and wbs respectively, give a segmentation of the sequence which consists of a set of all segments rather than only the signal segments. In order to include their results for comparison, we ignore the long segments (with length greater than 100) detected by CBS or WBS, which decreases their false positives. For LRS, we set the maximum length of signal segments as 50.

Example 2 We generate a sequence based on model (1) with $n = 10,000$. There are 5 signal segments with lengths 8, 16, 24, 32, and 40 respectively. We use the same error distributions as in Example 1. We consider two levels of height parameter for different signal strengths. In particular, we set height v as the 99-, and 97th percentiles of the marginal error distribution in two scenarios, labeled by S1 and S2. For the standard normal error, the height values are 2.326 and 1.881, respectively.

The threshold c we used for 4S is the 95th sample percentile of absolute values, that is around the 97.5th percentile of the error distribution, e.g., around 1.96 for the Gaussian case. Therefore, the true height is greater than c in S1, but lower in S2. Average numbers of TPs and FPs are reported in Tables 2 and 3.

Overall, CBS performs the best for the IID Gaussian case, but suffers from a low power in the heavy-tail case, and high FPs in the correlated case. LRS and WBS perform reasonable well with slightly high FPs in the heavy-tail case. 4S methods are more robust against the error type. When the noise is Gaussian and the signal strength is weak, it is slightly less powerful than the methods based on the Gaussian assumption. In terms of computation time (Table 4), 4S is about 100 times faster than all other methods.

Table 1 Average number of FPs for the null model

	4S (Vanilla)	4S ($p = 0.05$)	4S ($p = 0.1$)
$N(0,1)$	102.38	0.03	0.13
t_3	101.68	0.12	0.26
AR(1)	100.39	0.10	0.33

Table 2 Average number of TPs

	4S ($p = 0.05$)	4S ($p = 0.1$)	4S ($p = 0.5$)	CBS	LRS	WBS
S1						
N(0,1)	4.41	4.58	4.73	4.89	4.41	4.53
t_3	4.95	4.98	4.99	2.16	4.41	4.52
AR(1)	4.40	4.53	4.65	4.68	4.41	4.53
S2						
N(0,1)	3.77	3.94	4.20	4.59	3.76	3.94
t_3	3.34	3.47	3.73	0.51	3.75	3.93
AR(1)	3.75	3.94	4.09	4.43	3.76	3.95

Table 3 Average number of FPs

	4S ($p = 0.05$)	4S ($p = 0.1$)	4S ($p = 0.5$)	CBS	LRS	WBS
S1						
N(0,1)	0.02	0.05	0.29	0.10	0.05	0.14
t_3	0.04	0.09	0.36	0.13	0.45	0.36
AR(1)	0.05	0.14	0.44	1.99	0.05	0.14
S2						
N(0,1)	0.02	0.08	0.40	0.16	0.09	0.22
t_3	0.10	0.18	0.63	0.07	0.49	0.44
AR(1)	0.09	0.22	0.58	1.84	0.09	0.23

Table 4 Computation time (in s) to complete 300 sequences in S1 for each method

Method	4S	CBS	LRS	WBS
Time	0.83	115.52	108.72	86.17

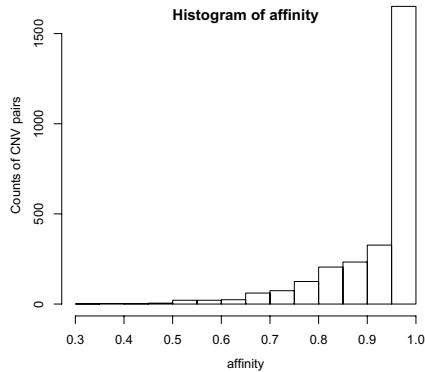
3.2 Real Data Example

We applied the 4S method to the 272 individuals from HapMap project. In particular, we tried 4S (with p value thresholds 0.05 and 0.5) to the LRR sequence of chromosome 1, which consists of 33991 measurements for each subject. We compared 4S with CBS, which has been a benchmark method in CNV detection. Note that CBS produces a segmentation of the sequence rather than the CNV segments directly. Therefore, we focused on only the short (less than 100 data points) segments detected by CBS because the long segments had means close to zero and are not likely to be CNVs. We found that most of these short segments are separated. But a very small portion of them are connected as CBS sometimes tends to over segment the sequence. Therefore, we merged two short segments detected by CBS if they are next to each other.

Table 5 Real data results: total number of detected CNVs, average length of CNVs, and computation time for all methods

	4S ($p = 0.05$)	4S ($p = 0.5$)	CBS
Number of CNVs	2832	3141	2962
Average length	28.60	27.70	23.15
Computation time	1.86	1.93	230.18

Fig. 5 Histogram of affinity, a similarity measure defined in (3), among 2753 pairs of commonly detected CNVs by 4S and CBS. Affinity equals to 1 if two detected CNVs are identical, and equals to 0 if two detected CNVs do not overlap



The 4S algorithm is extremely fast. It took less than 2 s (on a desktop with CPU 3.6 GHz Intel Core i7 and 16 GB memory) to complete 272 sequences with p values calculated for all detected segments. CBS algorithm is reasonably fast, but much slower than our algorithm. In Table 5, we list the total number of detected CNVs, average length of CNVs, and computation time for all methods.

Overall, the segment detection results were very similar. We further compared the segments detected by two algorithms, i.e., 4S with threshold $p = 0.05$ and CBS. We found that 2753 segments are in common. Here by a common segment we mean a pair of segments, one detected from each algorithm, such that they overlap to each other but do not overlap with other detected segments. Among these common segments, we calculated a similarity measure, called affinity in [1], defined as follows.

$$\rho(I, I') = \frac{|I \cap I'|}{\sqrt{|I| \cdot |I'|}}. \quad (3)$$

$\rho(I, I') = 1$ if two segments I and I' are the same and $\rho(I, I') = 0$ if they do not overlap. We found that the average value of this similarity measure is 0.9290 among 2753 pairs.

Figure 5 presents the histogram of affinity among 2753 pairs of commonly detected segments by 4S and CBS. We can see that 87.76% of those pairs have affinity values larger than 0.8. We further divided the detected segments into three groups: those detected by both methods (group 1); those detected by only 4S method with $p = 0.05$ (group 2); those detected by only CBS (group 3). For each detected segment, we calculated its length and the sample mean of the measurements on the segment. Figure 6 displays the scatter plots of sample means versus lengths for all

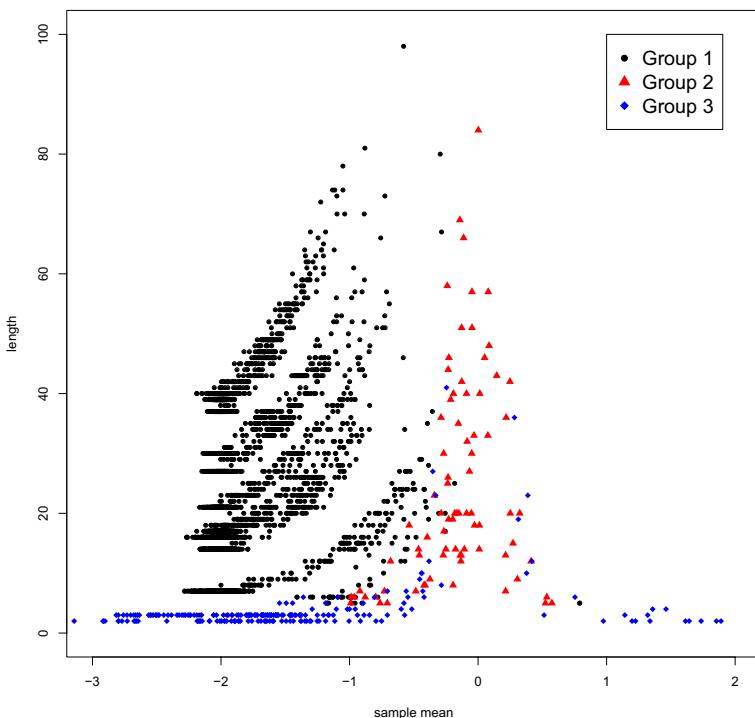


Fig. 6 Scatter plots of sample means versus lengths for detected segments in three groups. Group 1: segments detected by both methods; Group 2: segments detected by only 4S method with $p = 0.05$; Group 3: segments detected by only CBS

the segments in three groups. Most segments in group 1 carries relatively strong signals. So it is not surprised that they were detected by both algorithms. The groups 2 and 3 have much smaller sizes than group 1. In particular, we found that most segments in group 3 (i.e., those detected by only CBS) are very short, consisting of only 2 or 3 data points. Those segments are not significant in our inference framework unless we set a very high threshold c in step 1. Some segments in group 2 (i.e., those detected by only the 4S method) have relatively small sample mean values, which explains why they were not detected by CBS. Some of these segments might be true positives with the sample mean affected by outliers. Overall, the 4S and CBS methods gave similar results. The segments detected by only one method may be prone to false positives, or true positives have weak signal strengths.

4 Discussion

We proposed a scalable nonparametric algorithm for segment detection, and applied it to real data for CNV detection. Two main advantages of the 4S algorithm are its computational efficiency and independence of the normal error

assumption. We introduced an inference framework to assign significance levels to all detected segments. Our numerical studies demonstrated that our algorithm was much faster than CBS and performed similarly to CBS under the normality assumption and better when the normality assumption was violated. Although our inference framework depended on the assumption of IID noise, our numerical experiments suggested that our algorithm worked well under weakly correlated noises. Hence, the proposed method is faster and more robust against non-normal noises than CBS. Overall, the 4S algorithm is a safe and fast alternative to CBS, which has been a benchmark method in CNV studies.

In the literature, there are two popular classes of change-point models used to study CNV related problems. The first one assumes only a piecewise constant median/mean structure. The second one assumes, in addition, a baseline, which reflects the background information or normal status of the data. Quite often, it assumes that the abnormal part, called signal segments in our paper, are sparse. For the first approach, the goal is to identify the change points. In contrast, the second approach emphasizes more on segment detection rather than change-point detection. The difference is subtle for estimation but might become remarkable for inference. For example, it is technically difficult to define ‘true positive’ in the context of change-point detection [9]. But it is easier to define related concepts for segment detection as we did in this paper. Roughly speaking, the first approach is more general, and the second one is more specific and suitable to model certain CNV data, e.g., SNP array data. In particular, the 4S algorithm aims to solve change-point models in the second class. It can be applied to any data sequence when there is a baseline. When the baseline mean/median is unknown, we suggest that the data should be centered first by the estimated mean/median. Our method can not be applied to data when a baseline does not exist. Besides change-point models, there are other approaches to study CNV such as hidden Markov model [17]. Due to the space limit, we restricted our comparison to the methods based on change-point models and implemented by R packages.

Most segment detection algorithms involve one or more tuning parameters, whose values are critical to the results. In the study of segment detection, there are two trade-offs that researchers should consider in choosing algorithms as well as their parameters. The first one is the usual type I/type II errors trade-off, which might be tricky sometimes but well-known. The second one is more delicate and quite unique. For a signal segment, both its height and length determine the signal strength. Therefore, segments with weak but detectable signals can be roughly divided into two categories, the ones with small length (say, type A) and the ones with small height (say, type B). Typically, a method may detect type A segments more powerfully, but type B segments less powerfully, than the other method. For the proposed 4S algorithm, a choice of a larger threshold parameter in step 1 makes the algorithm more powerful in detecting short and high signal segments (type A), and vice versa. The 4S algorithm can be easily tuned to maximize the power in detecting of a certain type of the signal segments. We may also try different thresholding levels in data analysis in order to detect different types of segments. In general, the choice of the parameters depends on the research goals and balance of two trade-offs mentioned above.

There are various platforms and technologies which produce data for CNV detection. Besides the SNP array data studied in this work, read depth data from next-generation sequencing (NGS) technologies are often used in CNV studies. As one referee pointed out, the speed of 4S algorithm would be an advantage when applied to read depth data from whole genome sequencing. This is a wonderful research direction that we will investigate next.

An R package `SSSS` implementing our proposed method can be downloaded via <https://publichealth.yale.edu/c2s2/software>.

Acknowledgements The authors are grateful to two referees, an associate editor, and the editor for their helpful comments.

Funding The authors are partially supported by NSF grants DMS-1722691 (Niu and Hao), CCF-1740858 (Hao), DMS-1722562 (Xiao), DMS-1722544 (Zhang), Simons Foundation 524432 (Hao), University of Arizona faculty seed grant (Niu), and NIH grants R01DA016750, 1R01HG010171, and R01MH116527 (Zhang).

Appendix

Proof of Lemma 1 Let I be the interval of integers $[\ell, r]$ with $L = r + 1 - \ell$. For each $X_i, i \in I$, the probability that the ball at i is white is $\pi = P(|X_i| \leq c) \leq P(X_i \leq c) \leq \frac{1}{2}$ as $v \geq c$ and f is symmetric. It is trivial to bound $P(\mathcal{A})$ for the case $d \geq L$ as $P(\mathcal{A}^c) = \pi^L \leq 2^{-L}$. Now let us consider the case $d < L$. Let $\mathcal{E}_i, i \in I$ be the event that the first segment of d consecutive white balls starts from position i . Then

$$P(\mathcal{E}_i) = \begin{cases} \pi^d, & \text{if } i = \ell; \\ \leq (1 - \pi)\pi^d, & \text{if } \ell < i \leq r + 1 - d \\ = 0, & \text{if } i \geq r + 2 - d. \end{cases}$$

Therefore, $P(\mathcal{A}^c) = \sum_{i \in I} P(\mathcal{E}_i) \leq \pi^d + (L - d)(1 - \pi)\pi^d \leq (1 + \frac{1}{2}(L - d))(\frac{1}{2})^d$ and $P(\mathcal{A}) \geq 1 - (L - d + 2)2^{-d-1}$. Note that π is a constant depending on f , v and c , so a sharper bound than $\frac{1}{2}$ for π may be used to bound $P(\mathcal{A})$ if more information is available.

Let us consider the segment $[r + 1, r + D]$ on the right hand side of I . \mathcal{B}_D^c implies that there is at least one black ball in each of the segments $[r + 1, r + d]$, $[r + d + 1, r + 2d]$, etc. Note that $X_i \sim F$ on these segments so the probability of white ball at i is $P(|X_i| \leq c = f_\beta) = \beta' = 2\beta - 1$. Consider all $\lfloor \frac{D}{d} \rfloor$ segments of length d on the right side of I . The probability that all these segments contain at least one black ball is $(1 - \beta'^d)^{\lfloor \frac{D}{d} \rfloor}$. Therefore, $P(\mathcal{B}_D) \geq 1 - 2(1 - \beta'^d)^{\lfloor \frac{D}{d} \rfloor}$. \square

Proof of Theorem 1 For $\mathcal{I} = \bigcup_{k=1}^K I_k$, let $L_k = |I_k|$ and D_k be the gap between I_k and I_{k+1} . Define

$\mathcal{A}_k = \{\text{on } I_k, \text{ there does not exist a sub-segment of } \min\{d, L_k\} \text{ consecutive white balls}\}$
 $\mathcal{B}_k = \{\text{there are } d \text{ consecutive white balls on } D_k\}.$

Note that all segments in \mathcal{I} are identified under event $\left(\bigcap_{k=1}^K \mathcal{A}_k\right) \cap \left(\bigcap_{k=1}^{K-1} \mathcal{B}_k\right)$. By Lemma 1, $P(\mathcal{A}_k^c) \leq 2^{-L_k}$ or $\frac{1}{2}(L_k - d + 2)2^{-d}$, which can be bounded by $\max\{\frac{1}{2}(L_{\max} - d + 2), 1\}2^{-\min\{d, L_{\min}\}}$. Moreover, $P(\mathcal{B}_k^c) \leq (1 - \beta_{\min}^{td})^{\lfloor \frac{D_k}{d} \rfloor} \leq (1 - \beta_{\min}^{td})^{\lfloor \frac{D_{\min}}{d} \rfloor}$. The conclusion follows Bonferroni inequality. \square

Proof of Theorem 2 Let $\mathcal{S}_c = \{j_1, \dots, j_m\}$ be the locations of black balls after step 1. Note that j_i and j_{i+1} will be connected in step 2 if and only if $j_{i+1} - j_i \leq d$. We aims to count the number of segments with at least 2 consecutive black balls after step 2, as all isolated black balls will be eliminated in step 3. Such a segment starts at j_i only if $j_{i+1} - j_i \leq d$. So the total number of such segments is at most $|\{i : j_{i+1} - j_i \leq d\}|$. Let Z_i follow Bernoulli distribution with $Z_i = 1$ if and only if $j_{i+1} - j_i \leq d$ for $i = 1, \dots, m$. When $\mu = 0$, all black balls are randomly distributed. $P(Z_i = 0)$, i.e., the probability that all balls are white in next d positions following j_i is $\binom{n-1-d}{m-1} / \binom{n-1}{m-1} = \prod_{k=1}^d \frac{n-m+1-k}{n-k}$. Therefore, $E|\hat{\mathcal{I}}_{c,d,h}| \leq \sum_{i=1}^m Z_i \leq m(1 - \prod_{k=1}^d \frac{n-m+1-k}{n-k})$. \square

Proof of Lemma 2 We drop the subscript k in $P(\mathcal{A}_{n,m,s_k,t_k})$ as it is irrelevant in our derivation below. Under the assumption that m black balls are randomly assigned in n position, at a position of black ball, we calculate the probability that there are at least $t - 1$ black balls in next $s - 1$ positions. Let Y be the count of black balls in those $s - 1$ positions. Y follows a hypergeometric distribution with total population size $n - 1$, number of success states $m - 1$, and number of draws $s - 1$. Therefore, $P(\mathcal{A}_{n,m,s_k,t_k}) \leq mP(Y \geq t - 1)$ as there are m black balls. \square

References

1. Arias-Castro E, Donoho DL, Huo X (2005) Near-optimal detection of geometric objects by fast multiscale methods. *IEEE Trans Inf theory* 51(7):2402–2425
2. Cai TT, Jeng XJ, Li H (2012) Robust detection and identification of sparse segments in ultrahigh dimensional data analysis. *J R Stat Soc Ser B (Stat Methodol)* 74(5):773–797
3. Castellani CA, Awamleh Z, Melka MG, O'Reilly RL, Singh SM (2014) Copy number variation distribution in six monozygotic twin pairs discordant for schizophrenia. *Twin Res Hum Genet* 17(02):108–120
4. Fanale D, Iovanna JL, Calvo EL, Berthezen P, Belleau P, Dagorn JC, Ancona C, Catania G, D'alia P, Galvano A et al (2013) Analysis of germline gene copy number variants of patients with sporadic pancreatic adenocarcinoma reveals specific variations. *Oncology* 85(5):306–311
5. Feuk L, Carson AR, Scherer SW (2006) Structural variation in the human genome. *Nat Rev Genet* 7(2):85–97
6. Freeman JL, Perry GH, Feuk L, Redon R, McCarroll SA, Altshuler DM, Aburatani H, Jones KW, Tyler-Smith C, Hurles ME et al (2006) Copy number variation: new insights in genome diversity. *Genome Res* 16(8):949–961

7. Frick K, Munk A, Sieling H (2014) Multiscale change point inference. *J R Stat Soc Ser B (Stat Methodol)* 76(3):495–580
8. Fryzlewicz P (2014) Wild binary segmentation for multiple change-point detection. *Ann Stat* 42(6):2243–2281
9. Hao N, Niu YS, Zhang H (2013) Multiple change-point detection via a screening and ranking algorithm. *Stat Sin* 23:1553–1572
10. Jeng XJ, Cai TT, Li H (2010) Optimal sparse segment identification with application in copy number variation analysis. *J Am Stat Assoc* 105(491):1156–1166
11. Niu YS, Zhang H (2012) The screening and ranking algorithm to detect DNA copy number variations. *Ann Appl Stat* 6(3):1306–1326
12. Niu YS, Hao N, Zhang H (2016) Multiple change-point detection: a selective overview. *Stat Sci* 31(4):611–623. <https://doi.org/10.1214/16-STS587>
13. O'Dushlaine C, Ripke S, Ruderfer DM, Hamilton SP, Fava M, Iosifescu DV, Kohane IS, Churchill SE, Castro VM, Clements CC et al (2014) Rare copy number variation in treatment-resistant major depressive disorder. *Biol Psychiatry* 76(7):536–541
14. Olszen AB, Venkatraman E, Lucito R, Wigler M (2004) Circular binary segmentation for the analysis of array-based DNA copy number data. *Biostatistics* 5(4):557–572
15. Pinto D, Pagnamenta AT, Klei L, Anney R, Merico D, Regan R, Connroy J, Magalhaes TR, Correia C, Abrahams BS et al (2010) Functional impact of global rare copy number variation in autism spectrum disorders. *Nature* 466(7304):368–372
16. Tibshirani R, Wang P (2008) Spatial smoothing and hot spot detection for CGH data using the fused lasso. *Biostatistics* 9(1):18–29
17. Wang K, Li M, Hadley D, Liu R, Glessner J, Grant SF, Hakonarson H, Bucan M (2007) PennCNV: an integrated hidden Markov model designed for high-resolution copy number variation detection in whole-genome SNP genotyping data. *Genome Res* 17(11):1665–1674
18. Xiao F, Min X, Zhang H (2014) Modified screening and ranking algorithm for copy number variation detection. *Bioinformatics* 31(9):1341–8