



HAL
open science

Performance evaluation of DNA copy number segmentation methods

Morgane Pierre-Jean, Guillem Rigail, Pierre Neuvial

► **To cite this version:**

Morgane Pierre-Jean, Guillem Rigail, Pierre Neuvial. Performance evaluation of DNA copy number segmentation methods. *Briefings in Bioinformatics*, 2015, 16 (4), 10.1093/bib/bbu026 . hal-00952896v1

HAL Id: hal-00952896

<https://hal.science/hal-00952896v1>

Submitted on 27 Feb 2014 (v1), last revised 5 Nov 2015 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A performance evaluation framework of DNA copy number analysis methods in cancer studies; application to SNP array data segmentation methods

Morgane Pierre-Jean^{*1}, Guillem Rigail^{†2} and Pierre Neuvial^{‡1}

¹Laboratoire de Mathématiques et Modélisation d'Evry, Équipe Statistique et Génome, Université d'Evry val d'Essonne, UMR CNRS 8071, USC INRA

²INRA-CNRS-Université d'Evry, Unité de Recherche en Génomique Végétale (URGV)

February 27, 2014

Abstract

A number of bioinformatic or biostatistical methods are available for analyzing DNA copy number profiles measured from microarray or sequencing technologies. In the absence of rich enough gold standard data sets, the performance of these methods is generally assessed using unrealistic simulation studies, or based on small real data analyses.

We have designed and implemented a framework to generate realistic DNA copy number profiles of cancer samples with known truth. These profiles are generated by resampling real SNP microarray data from genomic regions with known copy-number state. The original real data have been extracted from dilutions series of tumor cell lines with matched blood samples at several concentrations. Therefore, the signal-to-noise ratio of the generated profiles can be controlled through the (known) percentage of tumor cells in the sample.

In this paper, we describe this framework and illustrate some of the benefits of the proposed data generation approach on a practical use case: a comparison study between methods for segmenting DNA copy number profiles from SNP microarrays. This study indicates that no single method is uniformly better than all others. It also helps identifying pros and cons for the compared methods as a function of

*morgane.pierrejean@genopole.cnrs.fr

†rigail@univ-evry.fr

‡pierre.neuvial@genopole.cnrs.fr

biologically informative parameters, such as the fraction of tumor cells in the sample and the proportion of heterozygous markers.

Availability: R package `jointSeg`: http://r-forge.r-project.org/R/?group_id=1562

Keywords: DNA copy number, segmentation, realistic data generation, performance evaluation.

1 Background

Changes in DNA copy numbers are a hallmark of cancer cells [10]. Therefore, the accurate detection and interpretation of such changes are two important steps toward improved diagnosis and treatment. The analysis of copy number profiles measured from high-throughput technologies such as array-comparative genomic hybridization (array-CGH), Single Nucleotide Polymorphism array (SNP array) or high-throughput DNA sequencing data raises a number of statistical and bioinformatic challenges.

Various methods have been proposed in the past ten years for analyzing such data. From a practitioner’s point of view, it is quite difficult to find which method is best for a given scientific question. In fact, it is likely that the overall difficulty of the problem depends on the context (technology, type of cancer, percentage of tumor cells). It is also likely that certain methods are more appropriate for certain contexts. Therefore, it is important to take this context into account when evaluating a set of methods, in order to 1) get a sense of the overall difficulty of the problem when interpreting the results and 2) choose appropriate methods for this context. The present work is motivated by the problem of comparing the performance of existing segmentation methods for identifying change-points from DNA copy number profiles from cancer patients. As any performance evaluation strategy, addressing this question requires the definition of three objects:

1. data with known “truth”;
2. methods to be compared;
3. criteria for performance assessment.

In this paper, we propose such a definition and illustrate how it may be used to compare segmentation methods. The main contributions of this work are

- a framework to generate realistic DNA copy-number profiles with known “truth”. This framework is generic and may be applied to any copy number data set;

- a framework to address the question of which SNP array data segmentation method performs best, depending on biologically relevant parameters.

These frameworks are implemented in the R package `jointSeg`. The rest of this paper is organized as follows. We start by giving some background on DNA copy number segmentation (Section 2) and describe our proposed data generation framework (Section 3). Then, we describe the pipeline we use for evaluating segmentation methods (Section 4). Finally, the result of our comparison study on two data sets are reported in Section 5.

2 DNA copy number segmentation

2.1 DNA copy number data

Normal cells have two copies of DNA, inherited from each biological parent of the individual. In tumor cells, parts of a chromosome of various sizes (from kilobases to a chromosome arm) may be deleted, or copied several times. As a result, DNA copy numbers in tumor cells are piecewise constant along the genome. Copy numbers can be measured using microarray or sequencing experiments. For illustration, Figure 1 displays an example of copy number signals that may be obtained from SNP-array data. Red vertical lines represent change points. In this particular example, the first region [0-2200] is normal, the second one [2200-6100] is a region where one of the parental chromosomes has been duplicated, and the third one [6100-10000] is a region of uniparental disomy, that is, a region where one of the parental chromosomes has been duplicated and the other one deleted. The top panel represents estimates of the total copy number (denoted by c). The bottom panel represents estimates of allelic ratios (denoted by b). We refer to [19] for an explanation of how these estimates may be obtained. In the normal region [0-2200], the total copy number is centered around two copies and allelic ratios have three modes centered at 0, 1/2 and 1. These modes correspond to homozygous SNPs AA ($b = 0$) and BB ($b = 1$), and heterozygous SNPs AB ($b = 1/2$). One important observation is that change points occur at the same position in both dimensions. This is explained by the fact that a change in only one of the parental copy numbers is reflected in both c and b . Therefore, it makes sense to analyze both dimensions of the signal jointly in order to identify change points.

In order to facilitate segmentation, allelic ratios (b) are generally transformed into unimodal signals, as originally proposed in [28]. This transformation is motivated by the fact that allelic ratios can be symmetrized (“folded”) and that SNPs that are homozygous in the germline (these SNPs are plotted in gray in Figure 1) can be discarded as they do not carry any information about copy-number changes. Following [2], we define the “decrease in heterozygosity” $d = 2|b - \frac{1}{2}|$ for SNPs that are heterozygous in the germline, which is essentially a rescaled version of the “mirrored/folded BAF” defined by [28]. After this transformation, DNA copy numbers can be considered as a bivariate, piecewise-constant signal, as illustrated by Figure 2.

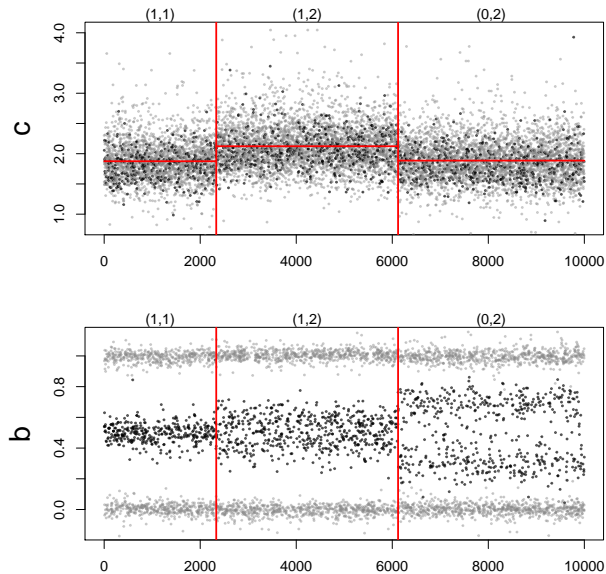


Figure 1: Example SNP array data. Total copy numbers (c), allelic ratios (b) along 10,000 genomic loci. Red vertical lines represent change points, and red horizontal lines represent mean signal levels between two change points. SNPs that are homozygous in the germline are colored in black; all of the other loci are colored in gray.

2.2 Typology of copy number segmentation methods

Many different methods have been proposed for the analysis of DNA copy number profiles. Most of them may be classified into four categories: methods based on Hidden Markov Models (HMM), multiple change-point methods, fused lasso-based methods and recursive segmentation methods.

1. HMM-based approaches rely on the idea that the recovered DNA copy number should be discrete and that these different levels can be modeled using a small number of HMM states. A typical example of such an HMM is the work of [7]. For the specific case of SNP array analysis in cancer samples, several dedicated HMM have been proposed [29, 9, 5].
2. Multiple change-point methods assume that the observed signal is affected by abrupt changes and that between these breaks the signal should be homogenous [23].
3. Methods based on a fused lasso penalty rely on the idea that, in most cases, two successive measurements should have the same estimate. This is encoded by a L_1 penalty on successive differences. The recovered signal is guaranteed to be piecewise constant. A typical example of such a fused model is the work of [30]. This class of methods can be viewed as solving a convex relaxation of the multiple change point problem.

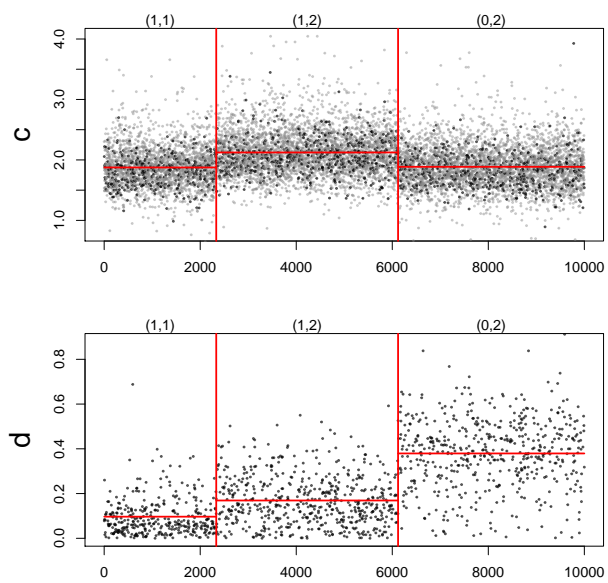


Figure 2: Example SNP array data along 10,000 genomic loci, after transformation of allelic ratios (b) into decrease in heterozygosity (d), following [2, 28]. Red vertical lines represent change points, and red horizontal lines represent mean signal levels between two change points. SNPs that are homozygous in the germline are colored in black; all of the other loci are colored in gray.

4. Recursive segmentation approaches rely on the intuitive idea that a segmentation can be recovered by recursively cutting the signal into two or more pieces. A typical example of such an recursive approach is the work of [21].

We refer to [19, 33] for a more mathematical introduction to these methods. Here, we only note that all of these methods assume that the signals are Gaussian. The above classification is by no means exhaustive (see for example [13, 1]).

3 Generating data with known “truth”

3.1 Review of existing approaches

A number of data generation mechanisms have been proposed in the context of performance evaluation of DNA copy number analysis in cancer samples, either in comparison studies [32, 17, 16, 12], or in papers describing new analysis tools. The generation of data with known “truth” can be done using either simulated or real data, both of which have opposite assets and drawbacks.

At first glance, simulated data are more appealing than real data because (i) “truth” is known with no ambiguity, (ii) the level of difficulty of the problem can be tuned as desired, and (iii) a large number of simulated data sets can be generated. As most DNA copy number segmentation methods rely on a Gaussian model (see Section 2), their performance is usually assessed using Gaussian simulations (see, for example, [23, 34]). While we do not question the usefulness of model assumptions for building statistical methods and for testing implementations, we believe that performance evaluation should as much as possible avoid to rely on a particular model. A recent study which compared several approaches for segmenting univariate DNA copy number profiles using the multiple change point approach showed that the best performing methods on Gaussian simulations performed quite poorly on real data [27, Table 3]. In the remainder of this section, we briefly review some existing approaches that have tried to take the best of both the “simulated data” and the “real data” worlds:

An automatically annotated data set [32]. The authors analyzed real data using one particular segmentation method to generate “truth”. They then used resampling to generate realistic copy-number profiles, where (Gaussian) noise was added in order to control the signal-to-noise ratio of the data set. Two drawbacks of this approach are that the notion of “truth” depends on the chosen segmentation method, and that the problem difficulty is not driven by biological considerations.

A dilution series [28]. In order to address the latter point, [28] have produced a dilution data set, where DNA from a lung cancer cell line is mixed with matched blood DNA from the same patient with varying (and known) mixture proportion (see description in Appendix A.1). Therefore, the fraction of tumor cells in the mixture controls the difficulty of the problem. The “truth” is a panel of regions whose DNA copy number status in the cell line (normal, gain, hemizygous deletion, copy-neutral LOH . . .) is known. This evaluation method has been accepted as a *de facto* standard and has been used in several subsequent papers, including [5, 20, 24].

An important drawback of this evaluation framework is that it focuses on a very limited number of regions (ten), which results in very little discrimination

between most methods in realistic settings. For example, four of the six methods compared in [20] reach maximum sensitivity in all 10 regions for tumor cell fractions greater than 25%. In practice, samples with less than 50% are rarely analyzed, in particular because the performance of most methods typically decreases severely when the fraction of tumor cells is less than 75%. We also note that sensitivity and specificity are evaluated separately in [28], and this weakness has been perpetuated in all subsequent papers based on the same evaluation framework.

A manually annotated data set [12]. The authors analyzed hundreds of neuroblastoma array-CGH profiles in order to define regions containing breakpoints (true signals), and regions not containing breakpoints (false signals). This data set is freely distributed on CRAN¹. Based on this large data set with known truth, the authors have performed a comprehensive comparison of segmentation methods for array-CGH data based on ROC curves. A drawback of this evaluation framework is that once a particular data set is chosen, it is not possible to tune the signal-to-noise ratio of the problem. Moreover, annotating a new data set is a challenging task, because it has to be large enough to contain a set of change-points that discriminate between competing segmentation methods.

A simulation model [17]. The authors designed a complex simulation model to generate “realistic” copy-number profiles. This model is implemented in the R package `CnaGen`, which is available from the authors’ web page². The simulation model depends on 24 parameters³. Some of them are directly driven by biological considerations, such as the percentage of tumor cells in the sample or intra-tumor heterogeneity. We empirically found it difficult to find a combination of parameters that yield realistic copy-number profiles. This may be due to the fact that the underlying data generation model is Gaussian. Table 1 summarizes the features of approaches reviewed above.

Reference	[32]	[28]	[12]	[17]	This paper
Based on real biological data?	☹	☹	☹	-	☹
SNR based on a biol. parameter?	-	☹	-	☹	☹
Data generation possible?	☹	-	-	☹	☹
Available as R package?	☹	-	☹	☹	☹

Table 1: Features of existing frameworks for real copy number data with known “truth”.

3.2 Proposed data generation mechanism

Based on these considerations, we propose an original data generation framework which aims at combining the advantages of all of the above-mentioned existing approaches. Two necessary and sufficient ingredients for generating a copy-number profile of length n are:

¹<http://cran.r-project.org/web/packages/neuroblastoma/>

²<http://web.bioinformatics.cicbiogune.es/cnagen/>

³`CnaGen` version 2.1.

- truth, in the form of K breakpoint positions (out of $n - 1$ intervals between two successive loci) and $K + 1$ copy-number state labels for all $K + 1$ regions between two consecutive breakpoints;
- signal, in the form of locus-level data. For SNP arrays, this is generally a $n \times 3$ matrix of total copy numbers (c), allelic ratios (b), and germline genotypes.

Our proposed approach is described below.

3.2.1 Generation of “truth”

When breakpoints and region labels are not user-supplied, we propose the following approach for generating them:

breakpoints: given a signal length n , draw K breakpoint positions uniformly out of the $n - 1$ possible intervals between successive data points (vertical red lines in Figure 3);

region labels: draw $K + 1$ region labels from a pre-defined set of copy-number state labels, such as normal, gain of one copy, hemizygous deletion, homozygous deletion, copy-neutral LOH (labels on top of each plot in Figure 3). By default, all region labels are equiprobable, but the user may provide a vector of probabilities for each desired region label. By default, successive regions are constrained in such a way that only one of the two parental copy numbers changes at the breakpoint.

3.2.2 Generation of locus-level data

Given breakpoint positions and region labels, we generate a copy-number profile as follows: for each region of size n_R between two breakpoints, we sample n_R data points from a real copy-number data corresponding to this type of region.

The data generation mechanism therefore relies on real data where the underlying region label is (assumed to be) known. We have made available two such “real data sets with known truth” in the package: each of them corresponds to a different SNP array platform (Affymetrix or Illumina), and both of them are taken from dilution series, consisting of mixtures of DNA from a tumor cell line and from blood cells originating from the same patient, with varying mixture proportions. For both data sets, we have selected several genomic regions which are representative of the diversity of copy-number states that are typically observed in tumor samples. Contrary to [32], these labels do not rely on any automatic segmentation or calling method. Both data sets are described in Appendix A.

3.3 Features of the proposed data generation mechanism

Our proposed data generation mechanism enjoys the following features:

- simplicity: small number of required parameters, all of which have a clear biological interpretation. In particular, for a given data set, the SNR is governed by the fraction of tumor cells. This is illustrated by Figure 3;
- flexibility: the user may specify breakpoint positions and region labels directly, if desired. Therefore, it is also possible to generate profiles with the same underlying “truth”, but with different SNR, as illustrated by Figure 3;

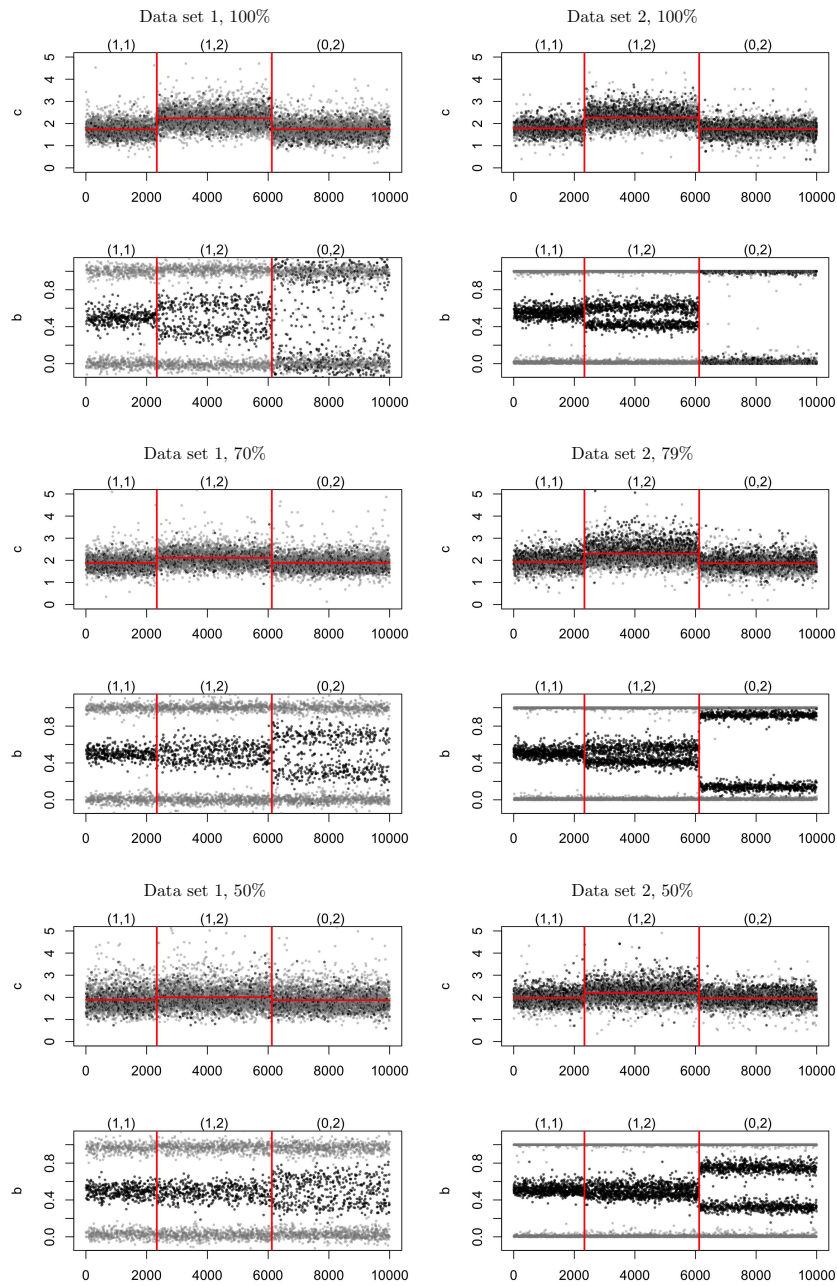


Figure 3: Illustration of the variety of copy-number profiles that can be generated from the same “truth” as in Figure 1. Each block of two plots corresponds to total copy numbers (c) and allelic ratios (b) for one particular combination of fraction of tumor cells (in rows) and data set (in columns). Red vertical lines represent change points. SNPs that are homozygous in the germline are colored in black; all of the other loci are colored in gray.

- reliability: copy-number regions were identified using the profiles with 100% tumor cells. In these profiles, the region labels may be defined manually unambiguously. Because the same tumor cell line is used for the dilutions series from a given platform, the regions identified on the profiles with 100% tumor cells can also be considered as ground truth for the profiles with less tumor cells, where direct manual identification would have been more problematic;
- versatility: the design choice of separating “truth” generation from locus-level data generation implies that it is relatively easy to:
 - annotate a new data set. Although dilution series are not publicly available for all possible platforms, it is also possible to annotate representative profiles from a given data set. Moreover, annotating a new data set is not time-expensive, as one only needs to identify a few copy-number regions.
 - extend the framework to other data types (for example array-CGH or high-throughput exome capture or whole genome sequencing) is straightforward: only a set of annotated data is required.

4 Evaluation pipeline

Now that we have a framework to generate data, we describe how to evaluate the performance of segmentation methods.

4.1 Benchmark

Synthetic copy-number profiles were generated as described in Section 3:

region-level “truth” : Each profile contains $n = 200,000$ loci in copy number signal and $K = 20$ breakpoints. We chose to impose the constraint that on average, 90% of segments are either normal (1,1), copy-neutral LOH (0,2), single copy-gain (1,2) or hemizygous deletion (0,1). The remaining 10% of regions are given less common copy-number states, such as homozygous deletion, or balanced duplication.

locus-level data: for each of $B = 50$ such “truth” profiles, corresponding locus-level data are then generated for 100%, 70% and 50% of tumor cells for data set 1, and 100%, 79% and 50% of tumor cells for data set 2.

4.2 Preprocessing

We log-transformed total copy numbers to stabilize their variance and smoothed outliers using `smooth.CNA` [21] as it improved segmentation results for all methods. Allelic ratios were converted to (unimodal) decrease in heterozygosity (d) as described in Section 2.1.

4.3 Compared segmentation methods

We evaluated different types of methods belonging to the different classes described in Section 2.2: multiple change-point, recursive, fused, and HMM-based methods. These methods are described in Table 2, where we mention which of them are

Name	R package	function	dims	Time (s)		Ref
				n=10 ⁴	n=10 ⁵	
Multiple change-point						
DP	cghseg	segmeanC0	1d	0.24	2.37	[26]
CST	cnaStruct	segment	2d	120	fail	[18]
DP	jointSeg	doDynamicProgramming	2d	140	fail	
Recursive						
CBS	DNAcopy	segment	1d	0.34	1.69	[31]
PSCBS	PSCBS	segmentByPairedPSCBS	2d	1.04	4.00	[20]
RBS	jointSeg	doRBS	2d	0.15	1.15	[8]
Fused						
GFLars	jointSeg	doGFLars	1d	0.29	3.70	[11]
GFLars	jointSeg	doGFLars	2d	0.08	0.60	[4]
HMM						
PSCN	PSCN	segmentation	2d	7.25	73	[5]

Table 2: List of DNA copy number segmentation methods evaluated.

able to process both signal dimensions (c and d) or only one of them. Not all of these methods were implemented in R. We ported from Matlab GFLseg⁴ to R the implementation of multi-dimensional dynamic programming and the group-fused LARS [4], and we implemented recursive binary segmentation [8] in R. In practice, as recommended by [4, 8, 11], both group-fused LARS and recursive binary segmentation are used to quickly identify a list of *candidate* change points, which is then pruned using dynamic programming.

4.4 Criteria for performance evaluation

Comparison studies typically assess the performance of DNA copy number analysis methods either in terms of their ability to accurately identify breakpoint locations [16, 12], copy-number states [28, 17], or both [32]. This paper focuses on the former only, because we are interested in comparing segmentation methods. The problem of evaluating strategies for calling copy-number states is left for future work.

As our proposed data generation framework provides copy number profiles with known “truth”, a natural way to evaluate the performance of a given method is to cast the problem of breakpoint detection as a binary classification problem. Specifically, for each generated copy number profile, we know where the true breakpoints are located. The number of true positives TP is the number of true breakpoints for which at least one breakpoint is detected closer than a given tolerance parameter. The number of false positives FP is defined as $FP = P - TP$, where P is the number of “positives”, that is, the total number of detected breakpoints. With this definition,

⁴Available at <http://cbio.ensmp.fr/~jvert/svn/GFLseg/html>.

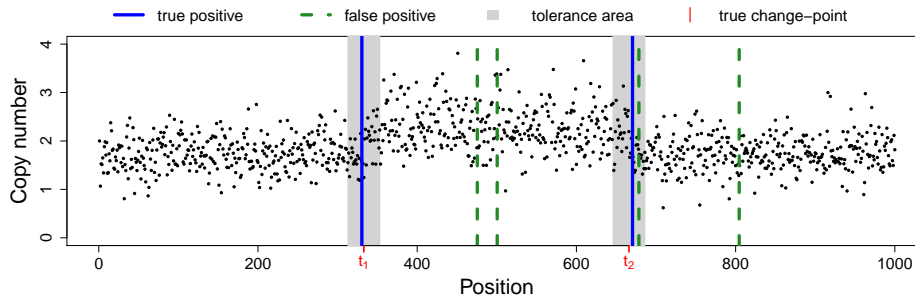


Figure 4: Definition of false positive and true positive to build performance evaluation.

whenever a method identifies two or more breakpoints within the tolerance area of a true breakpoint, one of these breakpoints counts as a true positive, while all others count as false positives. This definition of true and false positives is illustrated by Figure 4, where gray areas highlight tolerance areas around the true change-points, whose positions are identified as t_1 and t_2 on the x axis. In this example, breakpoints were detected in both shaded areas, therefore the number of true positives (solid blue lines) is two. There are four false positives (dashed green lines): one in a gray area where there is already one true positive, and three which are not within the tolerance area of any true breakpoint.

Related works. A similar definition of true and false positives is used in [32], although the authors do not mention how the above case of multiple breakpoints within the tolerance area is handled. Another related approach has been proposed in [12]. There, copy-number profiles are real, array-CGH profiles for which regions containing a breakpoint and regions containing no breakpoints have been delineated by experts. The main difference is that only a subset of the “true” and “false” breakpoints are annotated, and that the tolerance parameter cannot be tuned without the expert re-annotating the data set. Finally, a similar type of evaluation has been used by [16], at the locus level instead of the breakpoint level. This locus-level based evaluation method tends to favor segmentation methods that accurately identify large altered regions, even if they fail to detect breakpoints delineating smaller altered regions.

4.5 ROC-based evaluation

Usually, each method provides a segmentation and its associated set of breakpoints. This can be translated into a measure of sensitivity and specificity using the above definition of true and false positives. However, the methods have to be compared at the same specificity or sensitivity level in order for this comparison to be fair. Ideally, we would like to compute a Receiver Operator Characteristic (ROC) curve for each method. To do this one needs to explore a large set of possible segmentations with varying sensitivity and specificity, obtained by exploring the set of tuning parameters of each method. Such an exhaustive exploration is tedious and time consuming as soon as the number of parameters is larger than 2 or 3, and may

lead to over-optimistic results. To overcome this problem, we adopted the following strategy: for any given method m , we recovered a segmentation in k_m change points using default parameters, and we retrieved for each $k \in \{1 \dots k_m\}$ the best k subset of these k_m using dynamic programming. Another possible strategy would be to sort the k_m change points according to a measure of confidence.

One could be worried that the range of explored sensitivity/specificity is highly variable across methods. In practice, we found that the default parameters of a method tend to over-segment the data and typically, most of the true change points are found, at the cost of a more or less large number of false positives. This is in agreement with [12].

5 Results

5.1 Quantifying problem difficulty for known change points

Segmentation methods rely on a statistic to quantify the biological difference between any two regions. Based on this statistic, they aim at locating a good set of regions or equivalently, of change points. This location problem is combinatorial in nature. In this section, we try to quantify this biological difference independently of this combinatorial problem. To do this, we assume that change point positions are given *a priori* and we compare the power to call a change using total copy numbers (c) or allelic signals (d) for different types of change points. In order to perform this power study, we need to formally define the notion of power, or signal-to-noise ratio (SNR), between copy number regions. We chose a definition of SNR which is consistent with our proposed data-generation mechanism, in which DNA copy number data from a given region are sampled from a population which represents the corresponding copy-number state (see Section 3.3). Let us consider two regions and label by “0” and “1” the copy number state of two regions. For univariate signals (c or d), a natural definition of SNR is the (squared) Z statistic of the comparison between the sample means of region “0” and region “1”:

$$\text{SNR}(c) = \frac{(\bar{c}_0 - \bar{c}_1)^2}{\sigma_{c,0}^2/n_0 + \sigma_{c,1}^2/n_1} \quad (1)$$

$$\text{SNR}(d) = \frac{(\bar{d}_0 - \bar{d}_1)^2}{\sigma_{d,0}^2/n_0^* + \sigma_{d,1}^2/n_1^*}, \quad (2)$$

where n_i is the total number of loci in region i , \bar{c}_i and $\sigma_{c,i}$ are the sample mean and population standard deviation of total copy numbers in state i and $\bar{d}_i, \sigma_{d,i}$ are the sample mean and population standard deviation of the decrease in heterozygosity in state i . Note that the decrease in heterozygosity is only defined for SNPs that are heterozygous in the germline, whereas the total copy number is defined for all loci. Therefore, \bar{d}_i is calculated based on n_i^* heterozygous SNPs, while \bar{c}_i is calculated based on all n_i loci. For a given DNA sample, the fraction of heterozygous SNPs among those present on the microarray is typically close to 1/3; moreover, data set 1 contains not only SNP probes but also non-polymorphic loci, with a 1:1 ratio. As a result, the fraction n_i^*/n_i is approximately 1/6 for data set 1 and 1/3 for data set 2. A natural extension of this definition of SNR to the two-dimensional case of the statistic (c, d) is

$$\text{SNR}(c, d) = (\bar{c}_0 - \bar{c}_1, \bar{d}_0 - \bar{d}_1) (S_0 + S_1)^{-1} (\bar{c}_0 - \bar{c}_1, \bar{d}_0 - \bar{d}_1)', \quad (3)$$

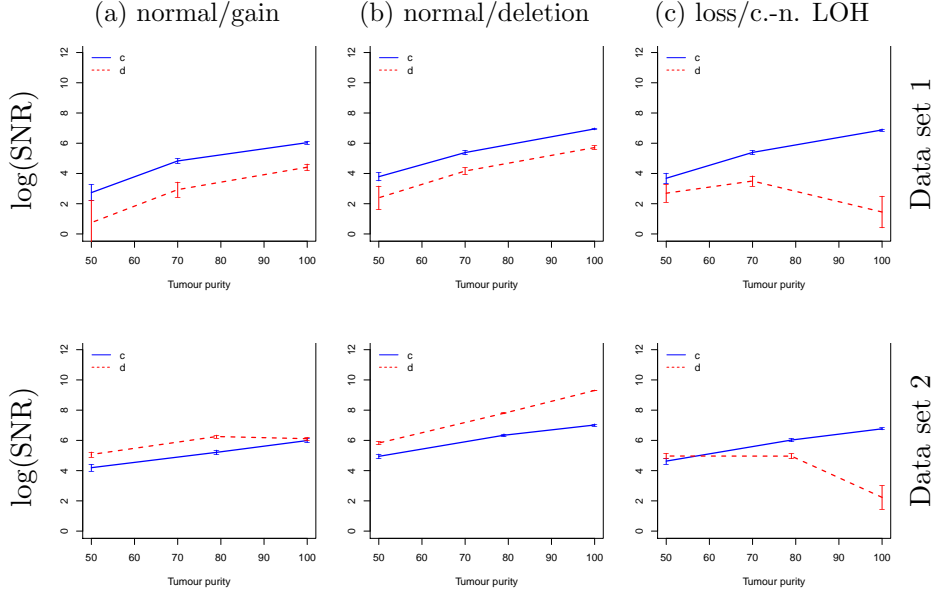


Figure 5: Average $\log(\text{SNR})$ and corresponding standard errors across 100 samples as a function of the percentage of tumor cells for total copy numbers (c , solid blue lines) and allelic ratios (d , dashed red lines). Each column corresponds to a type of copy number transition. Each row corresponds to a given data set.

where S_i is the population covariance matrix of the bivariate vector (c, d) , that is $S_i = \begin{pmatrix} \sigma_{c,i}^2/n_i & \tau_{cd,i}/n_i^* \\ \tau_{cd,i}/n_i^* & \sigma_{d,i}^2/n_i^* \end{pmatrix}$ with $\tau_{cd,i}$ the covariance between c and d in state i . In practice, the population parameters for copy-number state i (that is, $\sigma_{d,i}$, $\tau_{cd,i}$, and $\sigma_{c,i}$) are calculated from the annotated data. The sample parameters (\bar{c}_i and \bar{d}_i) are calculated from samples of n_i and n_i^* loci, respectively. Note that $\text{SNR}(c)$ and $\text{SNR}(d)$ are comparable with each other since they follow (non-centered) χ^2 distributions with 1 degree of freedom under the null hypothesis of no breakpoint between state 0 and state 1.

By definition, SNR is an increasing function of the length of each flanking segment. For $i \in \{0, 1\}$, we chose $n_i = 500$. n_i^* depends on proportion of heterozygous SNPs in the sample; as explained above, it is very close to $n_0/6$ for data set 1 and $n_0/3$ for data set 2. Therefore, the length of the flanking regions essentially acts as a constant scaling factor across all transitions and settings. Therefore, SNR only reflects differences between the underlying copy number states. Figure 5 shows the average (and standard error) of $\log(\text{SNR})$ across 100 samplings for three levels of tumor purity level, for three common types of copy number transitions for data set 1 (top panel) and data set 2 (bottom panels). Several conclusions may be drawn:

- **Difficulty generally increases with normal contamination:** SNR generally increases with the percentage of tumor cells. This is true for all types of

transitions for c . For d , the only situation in which SNR is not an increasing function of tumor purity is the case of transitions deletion and copy-neutral LOH (Figure 5c). This is expected theoretically because both of these states correspond to LOH in the tumor cells of the sample, implying that the true d in these cells is 1. In presence of normal cells, d estimates in both states are shrunk d toward 0, but in a state-specific way (see [2, Figure 4] for a detailed explanation of this phenomenon);

- **SNR levels depend on the type of copy number transition** for a given data set (that is, for a given row in Figure 5). This holds for both statistics (c or d). Note that in the case of c , this is unexpected, as all plotted transitions correspond to a one-copy gain.
- **Possibly low power.** Note that in some cases (e.g. data set 1, (a) and (c)), the computed SNR is lower than 2. Under the null hypothesis of no difference in mean levels, SNR follows a centered $\chi^2(1)$ distribution, so that this range of observed SNR correspond to p -values of the order of 1%, which is not low considering the large number of data points ($n_i = 500$).
- **Neither c or d is always the best statistic.** For a given type of transition (that is, for a given column in Figure 5) and a given statistic, the trend in SNR is comparable across data sets. However, the relative power of c with respect to d is much higher for data set 1 than for data set 2. This is directly related to the above-mentioned difference between ratios n_i^*/n_i of the number of informative loci for each statistic.

In this subsection, we assessed the intrinsic difficulty of calling a change point if the positions to test are known *a priori*. This study suggests that c and d are complementary sources of information, implying that change point detection methods should ideally take both of them into account. This study also sheds light on the fact that low percentages of tumor cells severely impacts SNR. In the remaining subsections, we assess the ability of segmentation methods to recover the true location of change points.

5.2 Robustness of the evaluation to the tolerance parameter

Our first goal was to check the influence of the tolerance parameter on the methods' performance. Our simulations were run using data generation as described in section 4.1. We computed partial areas under the ROC curves (pAUC) with a number of false positives between 0 and 10. Mean and 95% confidence intervals of pAUCs across simulation runs were calculated for each method for 5 values of the tolerance parameter (1, 2, 5, 10 and 20). For example, a tolerance of 5 means that a breakpoint is considered correct if it lies within 5 data points of the true breakpoints (see section 4.4 for more details). These results are reported in Figure 6 in the scenario without normal contamination. Similar results were observed for other scenarios.

Increasing tolerance clearly increases pAUC for all methods. This is the case even in the arguably "simple" scenario where no normal cells are present. However, in most cases the ranking of all methods is not affected by tolerance. Based on these results, we decided to report only pAUC for one particular value of tolerance: 5 loci on each side of the breakpoints.

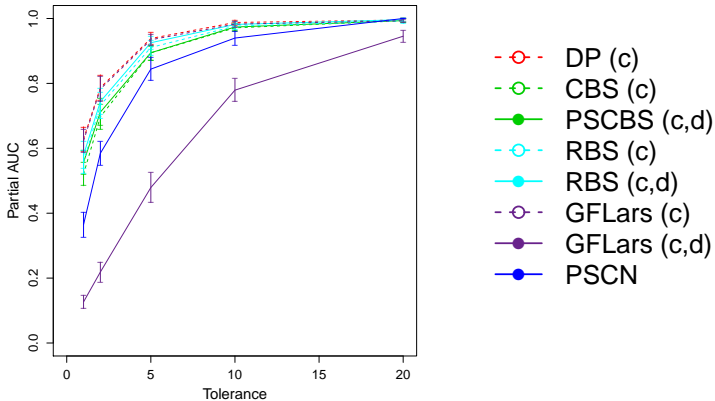


Figure 6: Method performance increase with the tolerance parameter for both data sets. Partial AUC for $FP \leq 10$ for data set 1 and 100% tumor cells.

5.3 Joint segmentation generally increases performance

This section aims at comparing the quality of segmentations obtained using total copy numbers only (c), allelic ratios only (d), and both of them (c, d) and how the quality of the segmentation is affected by the purity of the sample. As explained in section 5.1, it is typically expected that localization of the breakpoints is easier using both dimensions of the signal. To do so, we compared 6 scenarios corresponding to two data sets and three levels of purity (high, intermediate and low). Table 3 reports the pAUC of the best (c), (d) and (c, d) methods for data set 1 and 2, respectively. Detailed results for all methods are presented in Table 4.

For both data sets it is quite clear that performance in terms of pAUC deteriorates when the level of contamination increases. Interestingly, (c) methods perform better than (d) methods for both high and intermediate level of purity. For example in the case of data set 2 the minimum difference in pAUC between (c) and (d) is 26% for high level of purity and 7% for intermediate level of purity (Table 4). For a low level of purity, the results depend on the data set. For data set 1, (c) outperforms (d) but overall the pAUC are low. For data set 2, (d) outperforms (c) with a minimum pAUC difference of 7%. In that case the pAUC are quite good. These observations are in agreement with the results of Section 5.1. This difference between data set 1 and 2 can be explained by the fact that the proportion of informative probes for (d) (heterozygous SNPs) is twice less for data set 1 than for data set 2.

Interestingly, (c, d) methods do not always outperform (c)-only and (d)-only methods. However, as can be seen in Table 3, there are always several (c, d) approaches among top performers.

5.4 Choosing the appropriate method for a given context

In practice, when analyzing SNP array data, biostatisticians and bioinformaticians will choose one particular method to perform data segmentation. This choice is often *ad hoc* and based on personal experience. Our purpose here is not to make a comparison of all existing segmentation methods, but to compare relevant candidates in different classes of approaches. In the settings that we have considered it seems that RBS (c, d) performs very well. However, the point of our framework is not to select once and for all a best segmentation tool, but rather to justify the use of one method for one particular type of scenario (cancer type, cellularity, data set). In particular, we make no claim about the performance of RBS for other data sets.

Statistic	Data set 1			Data set 2		
	100%	70%	50%	100%	79%	50%
(c, d)	0.93	0.63	0.22	0.97	0.95	0.75
(c)	0.94	0.64	0.18	0.96	0.89	0.49
(d)	0.22	0.15	0.07	0.64	0.77	0.63

Table 3: Best pAUC across methods for each combination of statistic, data set and percentage of tumor cells.

Statistic	Method	Data set 1			Data set 2		
		100%	70%	50%	100%	79%	50%
(c, d)	PSCBS	0.89	0.60	0.16	0.97	0.88	0.51
	GFLars	0.46	0.35	0.13	0.90	0.87	0.61
	RBS	0.93	0.63	0.22	0.97	0.95	0.75
(c)	CBS	0.92	0.59	0.16	0.91	0.84	0.45
	GFLars	0.94	0.64	0.18	0.96	0.89	0.49
	RBS	0.91	0.62	0.17	0.90	0.84	0.48
	cghseg	0.93	0.61	0.18	0.95	0.88	0.49
(d)	CBS	0.22	0.14	0.09	0.61	0.70	0.56
	GFLars	0.22	0.15	0.08	0.64	0.77	0.62
	RBS	0.22	0.14	0.07	0.63	0.76	0.62
	cghseg	0.22	0.15	0.08	0.63	0.77	0.63

Table 4: pAUC by for each combination of method, statistic, data set and percentage of tumor cells.

5.5 Heterogeneity of breakpoint detection difficulty

An important question when using a biostatistical or bioinformatic tool is to assess its ability to recover events and to know which events they are likely to find and which of them are harder to detect. In Table 3 it can be seen that the pAUC is never at 100%. This is not necessarily surprising as the signal is quite noisy and in

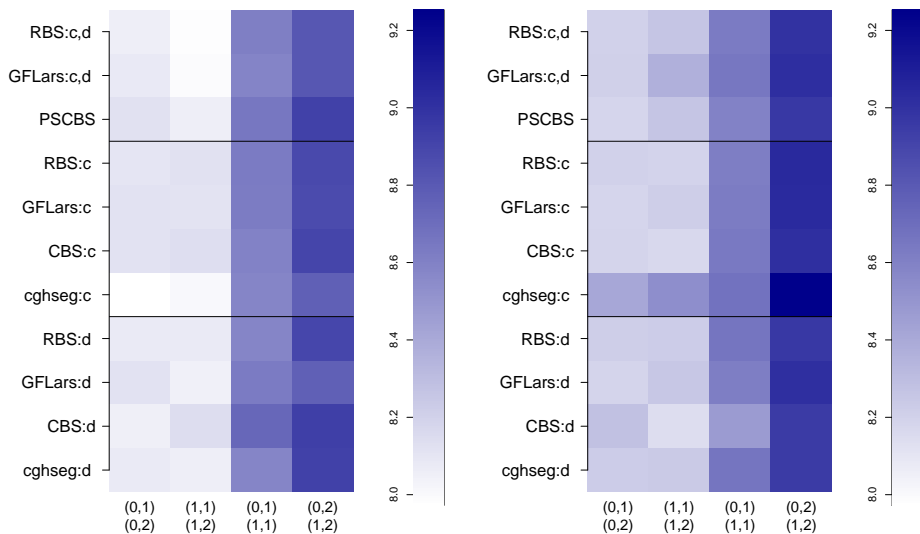


Figure 7: $\log(\text{SNR})$ for missed (left) and caught (right) breakpoints for four types of breakpoints on data set 2 with 50% normal cell contamination.

fact considering noise level the pAUC is quite high. Figure 7 demonstrates that (as could be expected) missed change-points are those for which we have a low signal to noise ratio (the right panel is darker than the left panel). However, the signal to noise ratio substantially depends on the type of change-point. Typically, in Figure 7 the column corresponding to the (0,2)-(1,2) transition is much darker than that of the (1,1)-(1,2) transition. This is confirmed by Table 5, which indicates that for a high level of normal contamination in data set 1, the proportion of missed (1,1)-(1,2) change-points is greater than $1/2$.

6 Summary and discussion

We have developed a framework to assess the performance of various DNA copy number segmentation methods. A critical aspect of this framework is that it generates realistic copy-number profiles by resampling real SNP array data. This allows to study a large number of scenarios without relying on a particular statistical model. It is our opinion that this framework is simple to use as it depends on few parameters, all of which have a straightforward biological interpretation. An R package is available and we believe that our proposed data generation scheme can be used readily as well as applied to other data sets and technologies. We illustrate the usage of this framework on two data sets.

We were able to identify which technological and biological parameters drive the performance of segmentation methods. First, it appears that the percentage of tumor cells in the sample plays a critical role: for a percentage lower than 70%, it is probably hopeless to recover the whole set of breakpoints with a high precision. We emphasize the relevance of the considered range of cellularity for applications: we expect that tumor cell lines should be well represented by the 100% setting,

Statistic	Method	(0,1)-(0,2)	(1,1)-(1,2)	(0,1)-(1,1)	(0,2)-(1,2)
(c, d)	RBS	0.40	0.47	0.32	0.31
	GFLars	0.49	0.65	0.41	0.38
	PSCBS	0.55	0.63	0.51	0.47
(c)	RBS	0.57	0.69	0.52	0.63
	GFLars	0.54	0.70	0.45	0.58
	CBS	0.59	0.71	0.52	0.62
	cghseg	0.66	0.79	0.55	0.69
(d)	RBS	0.54	0.58	0.46	0.33
	GFLars	0.56	0.54	0.42	0.29
	CBS	0.63	0.58	0.55	0.49
	cghseg	0.55	0.55	0.47	0.29

Table 5: Proportion of missed breakpoints by method, statistic and type of copy-number transition (data set 2).

while the 50% is not unusual for clinical practice. Second, it seems that different microarray technologies might lead to different performances. Specifically, the ratio between the number of informative allelic probes to the total number of probes is a crucial aspect, particularly for a high level of normal contamination. Finally, not all methods achieve similar performance across the scenarios that we have considered. Interestingly, we show that methods that take advantage of both signal dimensions are generally but not always better than those using only one of them. This variability between methods may be attributed to some extent to the biological and technological contexts, in the sense that some methods might be more adapted to certain scenarios.

Our framework provides a way to critically evaluate the performance of segmentation methods, and therefore to rationally select one or several of them for a particular data set. Such a quantitative assessment is also useful for interpretation. For example, we showed that even in favorable scenarios, performances are not perfect. Furthermore, perhaps unexpectedly, we showed that copy number transitions involving the gain or loss of a single DNA copy are not equally easy to recover, meaning that the proportion of different types of copy number transitions recovered by a particular segmentation method may not be directly interpretable.

7 Key Points

We propose a framework for generating realistic copy number profiles, and a framework to assess the performance of any segmentation method on such data. As our data generation framework resamples real data, it allows the comparison of different methods across a large number of different realistic scenarios. Our work highlights that the performance of segmentation methods is mainly driven by biological parameters such as the proportion of tumor cells in the sample. This demonstrates the importance of a correct and reproducible assessment of the performance prior to any biological interpretation of the data. The R package `jointSeg` implements this framework and allows anyone to reproduce our comparison study, either on the

data sets provided or on other data sets.

References

- [1] E. Ben-Yaacov and Y. C. Eldar. A fast and flexible method for the segmentation of aCGH data. *Bioinformatics*, 24(16):i139—i145, Aug. 2008.
- [2] H. Bengtsson, P. Neuvial, and T. P. Speed. TumorBoost: Normalization of allele-specific tumor copy numbers from a single pair of tumor-normal genotyping microarrays. *BMC bioinformatics*, 11(1):245, May 2010.
- [3] H. Bengtsson, P. Wirapati, and T. P. Speed. A single-array preprocessing method for estimating full-resolution raw copy numbers from all Affymetrix arrays including GenomeWideSNP 5 & 6. *Bioinformatics*, 27(17):2149–2156, 2009.
- [4] K. Bleakley and J.-P. Vert. The group fused Lasso for multiple change-point detection. Technical report, <http://hal.archives-ouvertes.fr/hal-00602121/en>, June 2011.
- [5] H. Chen, H. Xing, and N. R. Zhang. Estimation of parent specific DNA copy number in tumors using high-density genotyping arrays. *PLoS Computational Biology*, 7(1):e1001060, 2011.
- [6] R. Edgar, M. Domrachev, and A. Lash. Gene Expression Omnibus: NCBI gene expression and hybridization array data repository. *Nucleic acids research*, 30(1):207–210, 2002.
- [7] J. Fridlyand. Hidden Markov models approach to the analysis of array-CGH data. *Journal of Multivariate Analysis*, 90(1):132–153, July 2004.
- [8] S. Gey and E. Lebarbier. Using CART to detect multiple change points in the mean for large sample. Technical report, Statistics for Systems Biology research group, 2008.
- [9] C. D. Greenman, G. Bignell, A. Butler, S. Edkins, J. Hinton, D. Beare, S. Swamy, T. Santarius, L. Chen, S. Widaa, P. A. Futreal, and M. R. Stratton. PICNIC: an algorithm to predict absolute allelic copy number variation with microarray cancer data. *Biostatistics*, 11(1):164–175, 2010.
- [10] D. Hanahan and R. A. Weinberg. Hallmarks of cancer: the next generation. *Cell*, 144(5):646–674, 2011.
- [11] Z. Harchaoui and C. Lévy-Leduc. Catching change-points with lasso. *Advances in Neural Information Processing Systems*, 20:161–168, 2008.
- [12] T. Hocking, G. Schleiermacher, I. Janoueix-Lerosey, V. Boeva, J. Cappo, O. Delattre, F. Bach, and J.-P. Vert. Learning smoothing models of copy number profiles using breakpoint annotations. *BMC Bioinformatics*, 14(1):164, 2013.
- [13] P. Hupé, N. Stransky, J.-P. Thiery, F. Radvanyi, and E. Barillot. Analysis of array-CGH data: from signal ratio to gain and loss of DNA regions. *Bioinformatics*, 20(18):3413–3422, 2004.
- [14] Illumina, inc. Beadstudio genotyping module v3.2. User Guide.
- [15] Illumina, inc. Illumina’s Genotyping Data Normalization Methods. White paper, 2006.

- [16] W. R. Lai, M. D. Johnson, R. Kucherlapati, and P. J. Park. Comparative analysis of algorithms for identifying amplifications and deletions in array-CGH data. *Bioinformatics (Oxford, England)*, 21(19):3763–70, Oct. 2005.
- [17] D. Mosén-Ansorena, A. Aransay, and N. Rodríguez-Ezpeleta. Comparison of methods to detect copy number alterations in cancer using simulated and real genotyping data. *BMC bioinformatics*, 13(1):192, 2012.
- [18] D. Mosen-Ansorena and A. M. Aransay. Bivariate segmentation of SNP-array data for allele-specific copy number analysis in tumour samples. *BMC Bioinformatics*, 14(1):84, 2013.
- [19] P. Neuvial, H. Bengtsson, and T. P. Speed. Statistical analysis of single nucleotide polymorphism microarrays in cancer studies. In *Handbook of Statistical Bioinformatics*, Springer Handbooks of Computational Statistics. Springer, 1st edition, Mar. 2011.
- [20] A. B. Olshen, H. Bengtsson, P. Neuvial, P. T. Spellman, R. A. Olshen, and V. E. Seshan. Parent-specific copy number in paired tumor-normal studies using circular binary segmentation. *Bioinformatics*, 27(15):2038–2046, Aug 2011.
- [21] A. B. Olshen, E. S. Venkatraman, R. Lucito, and M. Wigler. Circular binary segmentation for the analysis of array-based DNA copy number data. *Biostatistics*, 5(4):557–572, 2004.
- [22] D. A. Peiffer, J. M. Le, F. J. Steemers, W. Chang, T. Jenniges, F. Garcia, K. Haden, J. Li, C. A. Shaw, J. Belmont, S. W. Cheung, R. M. Shen, D. L. Barker, and K. L. Gunderson. High-resolution genomic profiling of chromosomal aberrations using Infinium whole-genome genotyping. *Genome Res.*, 16(9):1136–1148, Sept. 2006.
- [23] F. Picard, S. Robin, M. Lavielle, C. Vaisse, and J.-J. Daudin. A statistical approach for array-CGH data analysis. *BMC bioinformatics*, 6:27, Jan. 2005.
- [24] P. M. Rancoita, M. Hutter, F. Bertoni, and I. Kwee. An integrated Bayesian analysis of LOH and copy number data. *BMC bioinformatics*, 11(1):321, 2010.
- [25] M. Rasmussen, M. Sundström, H. Göransson Kultima, J. Botling, and et al. Allele-specific copy number analysis of tumor samples with aneuploidy and tumor heterogeneity. *Genome Biol*, 12(10):R108, Oct. 2011.
- [26] G. Rigai. Pruned dynamic programming for optimal multiple change-point detection. Technical report, <http://arXiv.org/abs/1004.0887>, 2010.
- [27] G. Rigai, T. D. Hocking, F. Bach, and J.-P. Vert. Learning sparse penalties for change-point detection using max margin interval regression. *Proceedings of the 30th International Conference on Machine Learning*, JMLR W&CP, 28(3):172–180, 2013.
- [28] J. Staaf, D. Lindgren, J. Vallon-Christersson, A. Isaksson, and et al. Segmentation-based detection of allelic imbalance and loss-of-heterozygosity in cancer cells using whole genome SNP arrays. *Genome Biol*, 9(9):R136, Oct. 2008.

- [29] W. Sun, F. A. Wright, Z. Tang, S. H. Nordgard, P. V. Loo, T. Yu, V. N. Kristensen, and C. M. Perou. Integrated study of copy number states and genotype calls using high-density SNP arrays. *Nucl. Acids Res.*, 37(16):5365–5377, Sept. 2009.
- [30] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society. Series B, Statistical Methodology*, pages 91–108, 2005.
- [31] E. S. Venkatraman and A. B. Olshen. A faster circular binary segmentation algorithm for the analysis of array-CGH data. *Bioinformatics*, 23(6):657–663, Mar. 2007.
- [32] H. Willenbrock and J. Fridlyand. A comparison study: applying segmentation to array-CGH data for downstream analyses. *Bioinformatics*, 21(22):4084–91, Nov 2005.
- [33] N. Zhang. DNA copy number profiling in normal and tumor genomes. In J. Feng, W. Fu, and F. Sun, editors, *Frontiers in Computational and Systems Biology*, pages 259–281. Springer-Verlag, 2010.
- [34] N. R. Zhang and D. O. Siegmund. A modified Bayes Information Criterion with applications to the analysis of comparative genomic hybridization data. *Biometrics*, 63(1):22–32, 2007.

Appendix

A SNP array data sets

A.1 Data set 1

We have worked with a lung cancer data [25], for which raw data is accessible at NCBI GEO database [6], accession GSE29172. DNA from patient-matched lung cancer and blood cell lines NCI-H1395 and NCI-BL1395 were mixed to simulate tumor tissue with 30, 50, 70, 100% cancer cells. DNA was analyzed on Affymetrix SNP6.0 microarray. Data were normalized using ASCRMAv2 [3] followed by TumorBoost [2]. For the sake of reproducibility, the R scripts that were written to normalize this data set are distributed in the `jointSeg` package, together with the normalized data itself.

A.2 Data set 2

We have also worked with a breast cancer data [28], for which raw data is accessible at NCBI GEO database [6], accession GSE11976. DNA from patient-match breast cancer cell line (HCC1395) and its match normal HCC1395BL were mixed to simulate tumor tissue with 14, 34, 50, 79, 100% cancer cells. DNA was analyzed on Illumina HumanCNV370-Duov1 microarrays. We obtained the BAF-normalized and summarized data as calculated by the Illumina BeadStudio software [15, 14, 22]

A.3 Description of annotated copy-number regions

The list below describes the different copy number states available for data generation. They are labeled as a pair (c_1, c_2) , where c_1 corresponds to the minor copy number (the smallest of the two parental copy numbers), and c_2 corresponds to the major copy number (the largest of the two) [19].

- (1,1):** normal (one copy from each parent)
- (0,1):** hemizygous deletion (loss of one parental copy)
- (0,0):** homozygous deletion (loss of both parental copies)
- (0,2):** copy-neutral LOH (loss of one parental copy and gain of the other)
- (0,3):** loss of one parental copy and gain of two copies from the other parent)
- (1,2):** single copy gain
- (1,3):** unbalanced two-copy gain (gain of two copies from the same parent)
- (2,2):** balanced two-copy gain (gain of one copy from each parent)
- (2,3):** three-copy gain (gain of one copy from each parent, and two copies from the other parent)

CN state	(0,1)	(0,2)	(0,3)	(1,1)	(1,2)	(1,3)	(2,2)	(2,3)	(0,0)
Data set 1	22615	24135	25405	21539	19048	20903	27924	31098	0
Data set 2	2492	5484	6545	3196	2746	0	3044	0	838

Table 6: Size of annotated copy-number regions for each of the 2 data sets.

B Reproducing the figures in this paper

```
library(jointSeg)
path <- system.file("figures", package="jointSeg")
filenames <- list.files(path, pattern="*.R$")
for (filename in filenames) {
  print(filename)
  pathname <- file.path(path, filename)
  source(pathname, local=TRUE)
}
```

C Session information

```
> sessionInfo()
R version 3.0.2 (2013-09-25)
Platform: x86_64-apple-darwin10.8.0 (64-bit)

locale:
[1] fr_FR.UTF-8/fr_FR.UTF-8/fr_FR.UTF-8/C/fr_FR.UTF-8/fr_FR.UTF-8
```

attached base packages:

```
[1] parallel stats graphics grDevices utils datasets methods
[8] base
```

other attached packages:

```
[1] PSCBS_0.39.1      DNACopy_1.36.0    PSCN_1.0.1        MASS_7.3-29
[5] changepoint_1.1  zoo_1.7-10        cghseg_1.0.1      jointSeg_0.4.3
[9] acnr_0.1.4        R.utils_1.27.5    R.oo_1.15.8       R.methodsS3_1.5.2
[13] matrixStats_0.8.12
```

loaded via a namespace (and not attached):

```
[1] grid_3.0.2        lattice_0.20-24 R.cache_0.9.0
```