

RESEARCH

Open Access

Comparing copy-number profiles under multi-copy amplifications and deletions



Garance Cordonnier¹ and Manuel Lafond^{2*}

From 17th RECOMB Satellite Conference on Comparative Genomics
Montpellier, France. 1-4 October 2019

Abstract

Background: During cancer progression, malignant cells accumulate somatic mutations that can lead to genetic aberrations. In particular, evolutionary events akin to segmental duplications or deletions can alter the copy-number profile (CNP) of a set of genes in a genome. Our aim is to compute the evolutionary distance between two cells for which only CNPs are known. This asks for the minimum number of segmental amplifications and deletions to turn one CNP into another. This was recently formalized into a model where each event is assumed to alter a copy-number by 1 or -1 , even though these events can affect large portions of a chromosome.

Results: We propose a general cost framework where an event can modify the copy-number of a gene by larger amounts. We show that any cost scheme that allows segmental deletions of arbitrary length makes computing the distance strongly NP-hard. We then devise a factor 2 approximation algorithm for the problem when copy-numbers are non-zero and provide an implementation called `cnp2cnp`. We evaluate our approach experimentally by reconstructing simulated cancer phylogenies from the pairwise distances inferred by `cnp2cnp` and compare it against two other alternatives, namely the MEDICC distance and the Euclidean distance.

Conclusions: The experimental results show that our distance yields more accurate phylogenies on average than these alternatives if the given CNPs are error-free, but that the MEDICC distance is slightly more robust against error in the data. In all cases, our experiments show that either our approach or the MEDICC approach should be preferred over the Euclidean distance.

Keywords: Copy-number evolution, Algorithms, Cancer phylogenies, NP-hardness

Background

Cancer is widely recognized as an evolutionary process during which cells within a population accumulate aberrant somatic mutations and replicate indefinitely [1]. These cells are divided in subpopulations, called *clones*, that share common mutation traits and form tumors. A natural problem that arises is to reconstruct the evolution of a set of clones within a tumor. This question has

recently led to the development of several phylogenetic algorithms tailored for cancer evolution. Most of them use either information of single nucleotide variants obtained from bulk [2–5] or single-cell [6–8] sequencing data, or copy-number alterations [9–13] (usually in the context of single-cell data). We refer the reader to [14] for a survey of these methods.

In this work, we are interested in the problem of inferring the minimum number of copy-number alteration events that explain how a cell evolved into another. In tumors, several events can make the copy-number of a gene different from the normal diploid two-copy state,

*Correspondence: manuel.lafond@USherbrooke.ca

²Department of Computer Science, Université de Sherbrooke, Sherbrooke, Canada

Full list of author information is available at the end of the article



© The Author(s). 2020 **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

thereby creating *copy-number aberrations*. As an example, the *breakage–fusion–bridge* (BFB) phenomenon [15] occurs when a region including a telomere breaks off a chromosome. During replication, two sister chromatids have unterminated ends and they fuse, leading to what is essentially a chromosome portion concatenated with a reversed copy of itself (see [15, 16] for a more thorough explanation). Afterwards, the centromeres of the fused chromatids get pulled in opposite directions, leading to another breakage. This BFB cycle repeats until the chromatids receive a telomere (usually after translocation). Each BFB event potentially doubles the copy-number of a gene, and since these events are known to occur in cycles, a gene copy-number may become significantly higher than normal (i.e. more than double) in a short evolutionary time span. Other examples of events include focal deletions [17, 18] or missegregation of chromosomes [19].

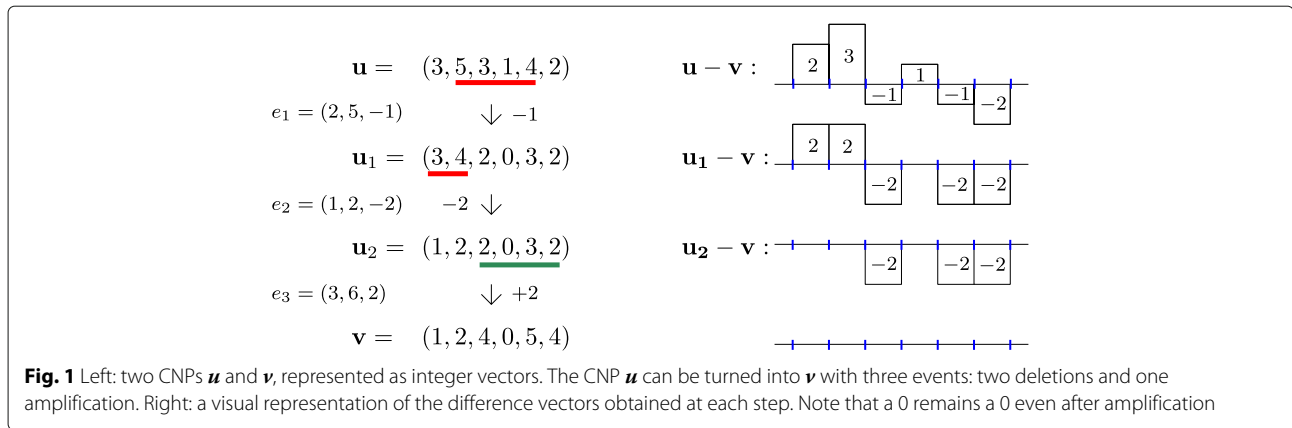
Desper et al. [20] were among the first to consider copy-number aberrations for phylogenetic reconstructions, using comparative genomic hybridization (CGH) data to reconstruct a mutation hierarchy. In [21], Liu et al. propose a distance-based approach based on CGH data to infer multi-cancer phylogenies. Single-cell phylogenetics then gained widespread attention in an influential paper of Navin et al. [9]. The authors applied single-nucleus sequencing on a breast cancer tumor, obtained the *copy-number profile* (CNP) of several cells, each represented as a vector of integers, and used the Euclidean distance to compare two CNPs. Later, Schwarz et al. [22] pointed out that a single event can amplify or delete large portions of a chromosome, thereby altering the copy-number of several genes and making the Euclidean distance overestimate the true number of events.

The authors proposed the following methodology to compare two CNPs. First, assuming diploid genomes, the copy-number for the two alleles of each gene (which can differ) is inferred from sequencing data. The correspondence between the copy-numbers and the alleles is unknown, so a *phasing* step must be applied. This consists of assigning each copy-number to one of the two alleles (this is done under a minimum-evolution principle, see [22] for details). After this step, each chromosome can be represented as a pair of CNPs, and chromosomes from two cells can be compared by computing the distances between the corresponding alleles. The distance proposed is the minimum number of *segmental* amplification and deletion events required to transform a given CNP into another.

In this work, we focus on the latter step. We assume that the CNP inference and the phasing steps have been performed, and must find a most parsimonious sequence of events explaining two given CNPs. This is analogous

to classical rearrangements problems [23], but the main novelty (and difficulty) of CNP comparison is that only copy-numbers are known, not the ordering of genes. In [22], Schwarz et al. introduced the MEDICC model, which approximates segmental events on a chromosome by events that alter a subinterval of a CNP by +1 or -1. Figure 1 shows an example turning a CNP \mathbf{u} into another \mathbf{v} (under our model where any amount of change is allowed). The problem of computing the minimum number of subinterval alterations to transform one CNP into another was solved in exponential-time in [22] by modeling CNP events with a finite-state transducer. Zeira et al. [24] gave a linear time algorithm, using a clever trick for computing each row of a quadratic-size dynamic programming table in constant time (similar to the techniques used in [25]). In [11], the large phylogeny problem under this model is shown NP-hard, though solvable with an ILP. They also present the *copy-number triplet* problem, which when given two CNPs \mathbf{u} and \mathbf{v} asks for a CNP whose sum of distances to \mathbf{u} and \mathbf{v} is minimized. The problem can be solved in pseudo-polynomial time $O(n^2N^7)$, where n is the CNP size and N the maximum copy number. Other distances and phylogenetic approaches are discussed in [10–13, 26, 27]

Our results. The above CNP comparison frameworks limit events to alter copy-numbers by 1 or -1. As we exemplified with BFB, several copies of a gene can be affected by a single event. Moreover, the MEDICC software has a copy-number limit of 4, making it inappropriate for genes attaining copy-numbers in the tens, twenties or even more, as has been reported for e.g. the MYC or EGFR genes [28–30]. In this work, we address these limitations by generalizing the Copy-Number Transformation problem defined in [22, 24]. We define a distance $d_f(\mathbf{u}, \mathbf{v})$ between two CNPs \mathbf{u} and \mathbf{v} which assigns a weight of $f(c, \delta)$ to an event that alters a copy-number of c by an amount of δ . We show that computing $d_f(\mathbf{u}, \mathbf{v})$ becomes strongly NP-hard whenever we allow deletions of any amount at unit cost. In the context of our problem, “strongly” means that our hardness holds even if N , the maximum value in \mathbf{u} and \mathbf{v} , is polynomial in n , the number of elements in our CNPs. This is especially relevant, given that the MEDICC model was initially solved in time $O(nN)$ and that the copy-number triplet problem can be solved in time $O(n^2N^7)$. Our result implies that such pseudo-polynomial time algorithm are impossible in our case unless $P = NP$. We then show that if any amount of change is permitted across an interval at unit cost, then a simple linear-time factor 2 approximation algorithm can be devised. We validate our approach by reconstructing simulated phylogenies using neighbor-joining (NJ), and compare them with the MEDICC distance and Euclidean distance. We perform our experiments on error-free data and noisy data (where the true copy-numbers



are altered by a random amount). Using a variety of simulation parameters, we show that both our distance and the MEDICC distance achieve significantly better accuracy than the Euclidean distance. Our distance is slightly more accurate on error-free data, and the MEDICC distance is slightly more tolerant to error.

Results

We first provide the required preliminary notions required to state our theoretical results. We then show that computing our copy-number distance is strongly NP-hard, and present our approximation algorithm. Finally, we present our experimental results on reconstructing simulated phylogenies.

Preliminary notions

Throughout the paper, we use the interval notations $[n] = \{1, 2, \dots, n\}$ and $[s, t] = \{s, s + 1, \dots, t\}$. Given a vector $\mathbf{u} = (u_1, \dots, u_n)$ of n integers and $i \in [n]$, we will always write u_i for the value at the i -th position of \mathbf{u} . If $u_i = 0$, then i is called a *null* position. We will assume that every vector \mathbf{u} of dimension n has special values $u_0 = u_{n+1} = 0$. We denote by $\mathbf{u}^{-[i]}$ the vector obtained by removing position $i \in [n]$, i.e. $\mathbf{u}^{-[i]} = (u_1, \dots, u_{i-1}, u_{i+1}, \dots, u_n)$. If \mathbf{v} is a vector of the same dimension, then $\mathbf{u} - \mathbf{v} = (u_1 - v_1, \dots, u_n - v_n)$.

We assume that a reference chromosome is partitioned into contiguous subsequences, called *positions*, each numbered from 1 to n . A *copy-number profile* (CNP) is a vector $\mathbf{u} = (u_1, \dots, u_n)$ of non-negative integers representing the copy-number of each position in a clone. We consider *amplification* and *deletion* events, which respectively have the effect of increasing and decreasing the number of copies in a chromosome. As in [22, 24], we assume that events affect a set of positions that are contiguous in the reference chromosome.

An *event* is a triple $e = (s, t, \delta)$ where $1 \leq s \leq t \leq n$ and $\delta \in \mathbb{Z} \setminus \{0\}$. Here the $[s, t]$ interval depicts the set of affected positions, and δ is the amount of

change. The event e is an amplification when $\delta > 0$ and a deletion when $\delta < 0$. A copy-number cannot drop below 0 and cannot increase from a 0 to another value (e.g. new genes cannot be created once completely lost). Applying event $e = (s, t, \delta)$ on a CNP \mathbf{u} yields another CNP $\mathbf{u}' = (u'_1, \dots, u'_n)$ with, for $i \in [n]$,

$$u'_i = \begin{cases} \max(u_i + \delta, 0) & \text{if } i \in [s, t] \text{ and } u_i > 0 \\ u_i & \text{otherwise} \end{cases}$$

We denote by $\mathbf{u}(e)$ the CNP obtained by applying event e on a CNP \mathbf{u} . More generally, if $E = (e_1, \dots, e_k)$ is an ordered sequence of events, we write $\mathbf{u}(E) = \mathbf{u}(e_1)(e_2) \dots (e_k)$ to denote the CNP obtained by applying each event of E in order. We may also write $\mathbf{u}(e_1 \dots e_k)$ instead of $\mathbf{u}((e_1, \dots, e_k))$. Given two CNPs \mathbf{u} and \mathbf{v} of dimension n , we say that E *transforms* \mathbf{u} into \mathbf{v} if $\mathbf{u}(E) = \mathbf{v}$.

We will often use the *difference vector* of \mathbf{u} and \mathbf{v} , and usually denote $\mathbf{w} := \mathbf{u} - \mathbf{v}$. The representation of \mathbf{w} as in Fig. 1 on the right provides the following intuition: if $\mathbf{u}(E) = \mathbf{v}$, then the events of E need to “squish” the values of \mathbf{w} to 0 to make \mathbf{u} equal to \mathbf{v} (ensuring that no value u_i of \mathbf{u} drops to 0 in the process unless $v_i = 0$).

Minimum cost transformations

Given two CNPs \mathbf{u} and \mathbf{v} , our goal is to find a minimum-cost sequence E that transforms \mathbf{u} into \mathbf{v} . In [22, 24], the cost of an event (s, t, δ) is $|\delta|$. Here, we propose a generalization by defining a cost function $f : \mathbb{N} \times \mathbb{Z} \rightarrow \mathbb{N}^{>0}$ that assigns a positive cost to altering a copy-number c by an amount of δ . That is, if we apply (s, t, δ) on \mathbf{u} , each position $i \in [s, t]$ has its own corresponding cost $f(u_i, \delta)$, which could be interpreted as the plausibility of going from copy-number u_i to $\max(u_i + \delta, 0)$. We then define the cost $cost_f(\mathbf{u}, e)$ with respect to f of applying $e = (s, t, \delta)$ on \mathbf{u} as the maximum cost within $[s, t]$, i.e.

$$cost_f(\mathbf{u}, e) = \max_{i \in [s, t]} f(u_i, \delta)$$

The events proposed in the MEDICC algorithm of Schwarz et al. can be decomposed into δ events of unit cost. This can be modeled under our framework with a function mdc defined as $mdc(u_i, \delta) = 1$ if $\delta \in \{-1, 1\}$ and $mdc(u_i, \delta) = \infty$ otherwise. Alternatively, one could state that a position with copy-number u_i can hardly more than double in a single event (assuming that amplifications are duplications), but that deletions can suppress any number of copies. We call this the doubling function dbl , defined as $dbl(u_i, \delta) = 1$ if $u_i + \delta \leq 2u_i$, and $dbl(u_i, \delta) = \infty$ otherwise.

Finally, the most permissive cost function *any* allows any movement without constraint: simply define $cost(u_i, \delta) = 1$ for any $\delta \in \mathbb{Z}$. This can, for instance, be used to model succession of events that can potentially amplify copy-numbers above their double in a short time span — an example of this being BFB cycles.

In this paper, we mostly analyze the *any* function for its simplicity, but will sometimes use the dbl function for its relevance. Given two CNPs \mathbf{u} and \mathbf{v} and a cost function f , the *cost* of a sequence of events $E = (e_1, \dots, e_k)$ satisfying $\mathbf{u}(E) = \mathbf{v}$ is equal to the sum of the cost of applying successive events of E on \mathbf{u} , i.e.

$$cost_f(\mathbf{u}, E) = cost_f(\mathbf{u}, e_1) + cost_f(\mathbf{u}(e_1), e_2) + \dots + cost_f(\mathbf{u}(e_1, \dots, e_{k-1}), e_k)$$

If $cost_f(\mathbf{u}, E) \leq cost_f(\mathbf{u}, E')$ for any other sequence E' satisfying $\mathbf{u}(E') = \mathbf{v}$, then E is called *optimal*. The f -distance between \mathbf{u} and \mathbf{v} , denoted $d_f(\mathbf{u}, \mathbf{v})$, is the cost of an optimal sequence of events transforming \mathbf{u} into \mathbf{v} . Observe that this “distance” is not symmetric (hence the use of double-quotes). For instance, if $\mathbf{u} = (1, 1)$ and $\mathbf{v} = (0, 0)$, then $d_{mdc}(\mathbf{u}, \mathbf{v}) = 1$ but $d_{mdc}(\mathbf{v}, \mathbf{u})$ is undefined since \mathbf{v} cannot be transformed into \mathbf{u} . We will therefore usually assume that \mathbf{u} does not have any null position. We note here that the *median* distance, defined as $\min_{\mathbf{w} \in V} (d_f(\mathbf{w}, \mathbf{u}) + d_f(\mathbf{w}, \mathbf{v}))$ (where V ranges over \mathbb{Z}^n), is symmetric for all the functions mentioned above. However, no efficient algorithm is known for any median distance. Our problem is the following.

The CNP-transformation problem:

Given: a source CNP \mathbf{u} , a target CNP \mathbf{v} , a cost function f and an integer k ;

Question: is $d_f(\mathbf{u}, \mathbf{v}) \leq k$?

We say that f is a *unit-cost* function if $f(c, \delta) \in \{1, \infty\}$ for any c and δ (e.g. the functions mdc , dbl and *any*). A cost function f is called *deletion-permissive* if $cost_f(u_i, \delta) = 1$ for any u_i and any $\delta < 0$, i.e. there is no particular constraint on deletions. We will mainly focus deletion-permissive functions, the rationale being that unlike duplications, deletions could suppress an arbitrary number of copies.

General properties

Before proceeding with our results on computing f -distances, we present some results of general interest that will be useful later on.

Proposition 1 For any two CNPs \mathbf{u} and \mathbf{v} of the same dimension, any position $i \in [n]$ and any cost function f , $d_f(\mathbf{u}, \mathbf{v}) \geq d_f(\mathbf{u}^{-\{i\}}, \mathbf{v}^{-\{i\}})$

We omit the proof details. The idea is that given a sequence of events E transforming \mathbf{u} into \mathbf{v} , we can apply E on $\mathbf{u}^{-\{i\}}$ by ignoring position i when it is affected. A sequence of events E is called *amp-first* if all amplifications appear before all deletions. An *amp-first reordering* of a sequence E is an amp-first sequence E' that contains the same events as E . Notice that if E has a amplifications and d deletions, then there are $a! d!$ amp-first reorderings of E .

Proposition 2 Let \mathbf{u} and \mathbf{v} be two CNPs with no null positions. If a sequence E satisfies $\mathbf{u}(E) = \mathbf{v}$, then any amp-first reordering E' of E satisfies $\mathbf{u}(E') = \mathbf{v}$.

Proof Denote $E = ((s_1, t_1, \delta_1), \dots, (s_k, t_k, \delta_k))$. For any position i , the sum $\sum_{j=1}^k \delta_k$ does not change even if we reorder the events in E , so u_i should still become v_i after reordering the event and applying them on \mathbf{u} . The only danger is that a position drops to 0 since \mathbf{v} has no null position, but this cannot happen if all amplifications are moved in front of E . \square

Given a CNP \mathbf{w} of length n , an interval $[a, b]$ is a *staircase* of \mathbf{w} if $0 < w_a < w_{a+1} < \dots < w_b$. The *length* of the staircase $[a, b]$ is $b - a + 1$. Figure 2 depicts a staircase of length 4. The next lemma can be useful to obtain quick

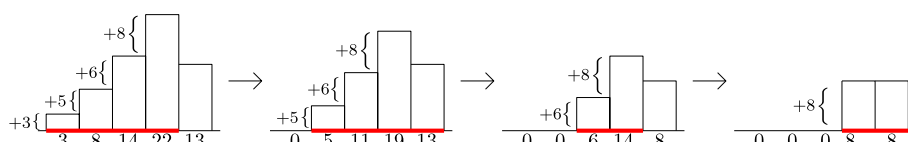


Fig. 2 A visual representation of the difference vector $\mathbf{w} = \mathbf{u} - \mathbf{v}$ leading to a staircase of length 4 in interval $[1, 4]$. For instance, setting $\mathbf{v} = (1, 1, 1, 1)$ and $\mathbf{u} = (4, 8, 15, 23, 14)$ would lead to the situation shown above. A smooth deletion sequence turning \mathbf{u} into \mathbf{v} is shown (last deletion omitted)

lower bounds on a particular instance, and plays an important role in our hardness result (proof in Supplementary material).

Lemma 1 *Let \mathbf{u}, \mathbf{v} be two CNPs with no null positions. If $\mathbf{u} - \mathbf{v}$ contains a staircase $[a, b]$ of length k , then $d_f(\mathbf{u}, \mathbf{v}) \geq k$ for any unit-cost function f .*

Strong NP-hardness

We show that the CNP-transformation problem is *strongly* NP-hard. This result holds for any deletion-permissive unit-cost function f , and even if \mathbf{u} and \mathbf{v} contain no null position (we note that in [24], null positions make the problem more complex, but not here). In particular, the hardness also holds if only deletions are allowed. We assume that we are given two CNPs \mathbf{u} and \mathbf{v} and we put $\mathbf{w} := \mathbf{u} - \mathbf{v}$.

Suppose that \mathbf{w} contains a staircase in interval $[1, k]$ for some k , and that $d_f(\mathbf{u}, \mathbf{v}) = k$. A sequence $E = (e_1, \dots, e_k)$ such that $\mathbf{u}(E) = \mathbf{v}$ is called *smooth* if, for every $i \in [k]$, $e_i = (i, b_i, w_{i-1} - w_i)$ for some $b_i \geq k$. Intuitively, E removes the first step, then the second, and so on, see Fig. 2. Observe that in a smooth deletion sequence, the positions to the right of k may or may not be affected by deletions.

Lemma 2 *Let \mathbf{u} and \mathbf{v} be two CNPs with no null positions and let f be any unit-cost function. If $\mathbf{u} - \mathbf{v}$ contains a staircase in interval $[1, k]$ and $d_f(\mathbf{u}, \mathbf{v}) = k$, then there exists a smooth sequence transforming \mathbf{u} into \mathbf{v} .*

Lemma 2 requires the most technical proof of the paper (by far), and we defer it to the Supplementary material. The reduction becomes relatively simple when given this lemma. Our reduction is from the 3-partition problem. In this problem, we are given a multi-set $S = \{s_1, \dots, s_n\}$ of $n = 3m$ positive integers. Defining $t := \frac{1}{m} \sum_{i \in [n]} s_i$ we are asked whether S can be partitioned into m subsets S_1, \dots, S_m , each of size 3, such that $\sum_{s \in S_i} s = t$ for all $i \in [m]$. This problem is known to be strongly NP-hard [31] (i.e. it is hard even if the values of S are $O(n^k)$ for some constant k). The proof can be found in the Supplementary material.

Theorem 1 *The CNP-transformation problem is strongly NP-hard for any deletion-permissive unit-cost function, even if the CNPs have no null positions.*

Approximation algorithm

In this section, we show that if \mathbf{v} does not contain any null position, then $d_{any}(\mathbf{u}, \mathbf{v})$ can be approximated within a factor of 2 in linear time. We discuss practical ways of handling null positions at the end of the section. We now assume that $f = any$ and will write $d(\mathbf{u}, \mathbf{v})$ instead of $d_{any}(\mathbf{u}, \mathbf{v})$.

As usual, \mathbf{u} and \mathbf{v} are the source and target CNPs, respectively, and $\mathbf{w} := \mathbf{u} - \mathbf{v}$. The idea of the approximation is that if two consecutive positions i and $i + 1$ have the same difference between \mathbf{u} and \mathbf{v} , i.e. $w_i = w_{i+1}$, then their value needs to change by the same amount. It might then be a good idea to treat these positions as one and always affect both with the same events. In fact, a whole interval of equal \mathbf{w} values can be treated as a single position. We show that the number of distinct equal intervals gives a good bound on $d(\mathbf{u}, \mathbf{v})$.

Approximation by flat intervals

Recall that if \mathbf{w} is a vector of n integers, it has implicit values $w_0 = w_{n+1} = 0$. We say that $[a, b]$, with $0 \leq a \leq b < n + 1$, is a *flat interval* if $w_i = w_j$ for every $a \leq i, j \leq b$. If no interval properly containing $[a, b]$ is flat, then $[a, b]$ is a *maximal flat interval*. In fact, in the remainder, we will omit the term “maximal” and always assume that discussed flat intervals are maximal. We write F_w for the set of flat intervals of \mathbf{w} . Note that this set is well-defined and that it partitions $[0, n + 1]$, by the maximality property. The intervals that contain 0 and $n + 1$ in F_w are called *extreme flat intervals*, and always have a value of 0 (also, these intervals are possibly $[0, 0]$ and/or $[n + 1, n + 1]$, but not necessarily). The key lemma says that $d_f(\mathbf{u}, \mathbf{v})$ is at least about half the number of flat intervals (see Supplementary material).

Lemma 3 *Let \mathbf{u}, \mathbf{v} be two distinct CNPs with no null positions, and let $\mathbf{w} := \mathbf{u} - \mathbf{v}$. Then for any unit-cost function f , $d_f(\mathbf{u}, \mathbf{v}) \geq \lceil (|F_w| - 1)/2 \rceil$.*

Lemma 3 yields a very simple factor 2 approximation algorithm: compute F_w , and return $|F_w| - 2$. This corresponds to a solution in which we treat each flat interval separately (ignoring the two extremities) and is guaranteed to be at most twice the optimal number of events. Computing F_w can be done in a single pass through \mathbf{w} by increasing a counter whenever we encounter a position i with $w_i \neq w_{i-1}$.

Theorem 2 *The CNP-transformation problem can be approximated within factor 2 in linear time for cost function $f = any$ when the CNPs contain no null position.*

It is open whether this could be adapted to other functions, e.g. the *dbl* function.

Improvements to the approximation algorithm

We first observe that the bound in Lemma 3 is essentially tight. This can be seen with any \mathbf{u}, \mathbf{v} such that $\mathbf{u} - \mathbf{v} = (1, 2, 3, \dots, k - 1, k, k - 1, \dots, 3, 2, 1)$ for some k . Indeed, one can decrease $|F_w|$ by two at each round. On the other hand, our naive 2-approximation is twice as bad

as optimal. We show how to improve this in a heuristic fashion by devising an algorithm that can only perform better than the naive one. We leave it as an open problem to determine the approximation guarantees of this algorithm.

Our goal is to apply events that reduce $|F_w|$ by two as many times as possible. In a greedy fashion, we apply the following strategy for our improved 2-approximation: as long as $\mathbf{u} \neq \mathbf{v}$, find an event e that reduces $|F_w|$ by 2, if one exists, and apply it to \mathbf{u} . If no such event exists, take the leftmost non-extreme flat interval $[a, b]$ of \mathbf{w} and apply the event $(a, b, -w_a)$. Repeat until $\mathbf{u} = \mathbf{v}$.

An event (i, j, δ) reduces $|F_w|$ by 2 precisely when $w_{i-1} - w_i = w_{j+1} - w_j = \delta$ ($i = j$ is possible). This way we can merge the two flat intervals at the ends of $[i, j]$. One can find a good interval by checking all the $O(n^2)$ subintervals $[i, j]$ and then, for each of them, checking whether $w_{i-1} - w_i = w_{j+1} - w_j$. Moreover, we must check whether applying the event (i, j, δ) would make a value of \mathbf{u} go below 0. Verifying every possible event can be done in time $O(n^3)$ and as there are $O(n)$ flat intervals, the algorithm takes time $O(n^4)$.

This can be improved to $O(n^2 \log n)$ by finding good events in time $O(n \log n)$. Due to space constraints, we relegate the detailed analysis of the improved heuristic to the Supplementary material. The idea is to scan \mathbf{w} from left to right and store in a *treap* data structure (see [32]) the set of flat intervals encountered so far, which allows to detect quickly whether the current flat interval could be matched with another one.

Handling null positions

Our approximation ratio is not guaranteed to hold when there are many null positions. However, we show that in many practical cases, we can simply ignore null positions and remove them. In particular, we may assume that \mathbf{v} has no two consecutive null positions (Lemma 4) and that for any null position i in \mathbf{v} , we have $w_{i-1} < w_i$ and $w_{i+1} < w_i$ (Lemma 5). Thus instances with null positions can be reduced to ones where the only null positions remaining are “sandwiched” between non-null positions with a smaller value in \mathbf{w} .

Note that our approximation can still perform badly with these two conditions. For instance, suppose that $\mathbf{u} = (15, 2, 15, 2, \dots, 15, 2)$ and $\mathbf{v} = (14, 0, 14, 0, \dots, 14, 0)$. We would solve this in about $n/2$ events. However, the two events $(1, n, -2), (1, n, 1)$ turn \mathbf{u} into \mathbf{v} . Designing a better approximation for these cases is an open problem.

Lemma 4 *Suppose that $v_i = v_{i+1} = 0$ for some position i . Then removing position i or $i + 1$, whichever is smaller in \mathbf{u} , from \mathbf{u} and \mathbf{v} preserves the distance between \mathbf{u} and \mathbf{v} . Formally, for any unit-cost function f , if $u_i \geq u_{i+1}$, then*

$d_f(\mathbf{u}, \mathbf{v}) = d_f(\mathbf{u}^{-\{i+1\}}, \mathbf{v}^{-\{i+1\}})$. Similarly if $u_{i+1} \geq u_i$, then $d_f(\mathbf{u}, \mathbf{v}) = d_f(\mathbf{u}^{-\{i\}}, \mathbf{v}^{-\{i\}})$.

Lemma 5 *Suppose $v_i = 0$ for some position i and that $w_{i-1} \geq w_i$ or $w_{i+1} \geq w_i$. Then $d_f(\mathbf{u}, \mathbf{v}) = d_f(\mathbf{u}^{-\{i\}}, \mathbf{v}^{-\{i\}})$ for any unit-cost function f .*

Experiments

We tested our flattening approximation algorithm and its improved version on simulated chromosomes that evolve along a tree through segmental tandem duplications and losses. Chromosomes were represented as strings of genes. Note that we did not simulate CNP evolution under the assumptions of our model. We evolved actual sequences as opposed to integer vectors, and the initial ordering of genes could be broken after several events. Our goal was to reconstruct phylogenies from the distances between the CNPs of the chromosomes at the leaves of the tree. We used the NJ implementation in Phylip [33, 34] and compared four distances: (1) our improved approximation; (2) our flat interval count; (3) the *mdc* cost, as in the MEDICC model; and (4) the Euclidean distance. To compute d_{mdc} , we implemented the dynamic programming algorithm of Zeira, Zehavi and Shamir [24], hereafter called the ZZS algorithm (we could not use the MEDICC software as it only handles copy-numbers up to 4). The Euclidean distance is defined as $\sqrt{\sum_{i=1}^n (u_i - v_i)^2}$, as used in [9]. For the first three distances, we took the minimum of $d(\mathbf{u}, \mathbf{v})$ or $d(\mathbf{v}, \mathbf{u})$ to get a symmetric distance, removing null positions of \mathbf{u} and filtering null positions of \mathbf{v} as in Lemmas 4 and 5.

Simulated tree generation

We now describe how the trees were generated. First, we select a rooted binary tree T on l leaves labeled $\{1, \dots, l\}$ uniformly at random. This is achieved by using the recursive splitting process described by Aldous in [35], which starts with a completely unresolved tree, splits the root in two subtrees chosen uniformly at random, and repeats on these subtrees. We then assign to the root r of T an *exemplar chromosome*, i.e. any string in which each gene occurs exactly once (note that the initial ordering of genes does not matter for our purposes).

Then for each branch uv of T from top to bottom, we select a random number of events k chosen uniformly at random in the interval $[e_{min}, e_{max}]$, where e_{min}, e_{max} are simulation parameters. To introduce some rate heterogeneity among branches, we then multiplied k by a random number chosen from a uniform distribution with mean and standard deviation 1. The chromosome string at node v is obtained by applying k random events on the chromosome string associated with its parent u . Each event is either a tandem duplication with probability Δ or a deletion with probability $1 - \Delta$. The starting position

of each event is chosen uniformly at random on the chromosome string and, to find the length t of the substring affected, we apply the following process. Start with $t = 1$, then apply the following: as long as a random number between 0 and 1 is above a given parameter r , increment t by 1 and repeat. We stop at the first random number below r . We chose to consider only values $r \leq 0.1$ since higher values resulted in copy-numbers in the hundreds, sometimes even in the thousands, a road which we did not deem necessary to explore. Setting r between 0.01 and 0.1 generally resulted in copy-numbers inside $[0, 50]$. We also note that we also experimented on a model where the event length was chosen as a random fraction of the chromosome length — this led to exponential copy-number growth and we did not investigate this model further.

We observed that this process had a tendency to produce leaf chromosomes with CNPs having between 50–60% null positions under most parameter combinations. This might be deemed unrealistic, and furthermore, our results show that no method is able to predict accurate trees under these conditions. To avoid this, we added a condition in the loop determining the length t of an event: if incrementing t implies deleting the last occurrence of a gene, we continue the procedure with probability q and stop with probability $1 - q$ (where q is another simulation parameter). This can be seen as modeling the idea that there may be resistance when attempting to remove every copy of a gene required for survival. Using q parameter values 0.25, 0.5 and 0.75, the proportion of null positions stayed in the intervals 2–5%, 6–10% and 15–25%, respectively.

Note that since each possible tree on l leaves is equally likely to be chosen, the root-to-leaf distances in a tree can be significantly different, and hence the trees are not expected to be ultrametric (for instance, a caterpillar can be selected as well as a perfectly binary tree).

Since it is difficult to determine the most realistic simulation conditions, we tested several combinations of parameters for the generation of phylogenies. The summary of the simulation parameters, along with their possible and default values, are summarized here:

- $l \in \{10, 50, 100\}$ is the number of leaves in the tree. The default is $l = 100$;
- $n \in \{10, 100, 250\}$ is the number of genes (i.e. distinct characters) in the root chromosome (n is also the number of positions in our vectors). The default is $n = 100$;
- $(e_{min}, e_{max}) \in \{(2, 4), (5, 10), (20, 40)\}$ is the range of the possible number of events on each branch. The default is $(e_{min}, e_{max}) = (5, 10)$;
- $\Delta \in \{0.25, 0.5, 0.75\}$ is the probability that an event is a duplication (and $1 - \Delta$ the probability that an event is a loss). The default is $\Delta = 0.5$;

- $r \in \{0.01, 0.05, 0.1\}$ controls the length of each event: r is the probability that we stop extending the event length. The default is $r = 0.05$;
- $q \in \{0.25, 0.5, 0.75, 1\}$ is the probability that a deletion suppresses the last copy of a gene during the length extension procedure (i.e. $1 - q$ is the probability that the extension stops if it would make a copy-number 0). The default is $q = 0.25$.

Tree reconstruction and performance measure

We generated 50 trees for each parameter combination of interest. For each tree, we took the chromosome strings at the leaves, obtained their CNPs and provided them as input to each of the four evaluated methods. We used the normalized Robinson-Foulds (RF) distance as a measure of the performance of each algorithm [36]. That is, for each inferred tree, we compare it with the “true” tree by counting the number of clades that are present in one tree but not the other, divided by $2(l - 3)$ (the maximum number of clades that can possibly differ, recalling that l is the number of leaves). This yields a number between 0 and 1: the lower the number, the better we consider the reconstruction.

Error tolerance

It should be noted that the above methodology ignores several sources of errors. Inferring exact copy-numbers from single-cell sequencing data is a non-trivial task and is still considered an open problem. The inferred CNPs are therefore expected to be noisy, especially with genes having a high copy-number. Moreover, as discussed in [22], assigning copy-numbers to their corresponding allele is also a difficult problem. Here, by only considering single-allele chromosomes, we are supposing that the aforementioned phasing step has been performed correctly, whereas copy-number assignments cannot be assumed to be error-free.

Both of the above problems have the effect of introducing incorrect copy-numbers into the CNPs. To account for this, we gave randomly altered CNPs as input to each method. More specifically, given an error-rate parameter α , for each CNP u and each position i we changed u_i to a value chosen at random from a normal distribution with mean u_i and standard deviation $\alpha \cdot u_i$ (non-integer values were rounded). We tested parameter values $\alpha \in \{0, 0.1, 0.25, 0.5, 1\}$.

Experimental results

We first ran experiments using the default values for all parameters except one in order to isolate the impact of each parameter. On error-free data, the most interesting results were obtained when varying l and n , see Fig. 3. In most situations, our CNP model slightly improves upon the MEDICC model, both of which are significantly better

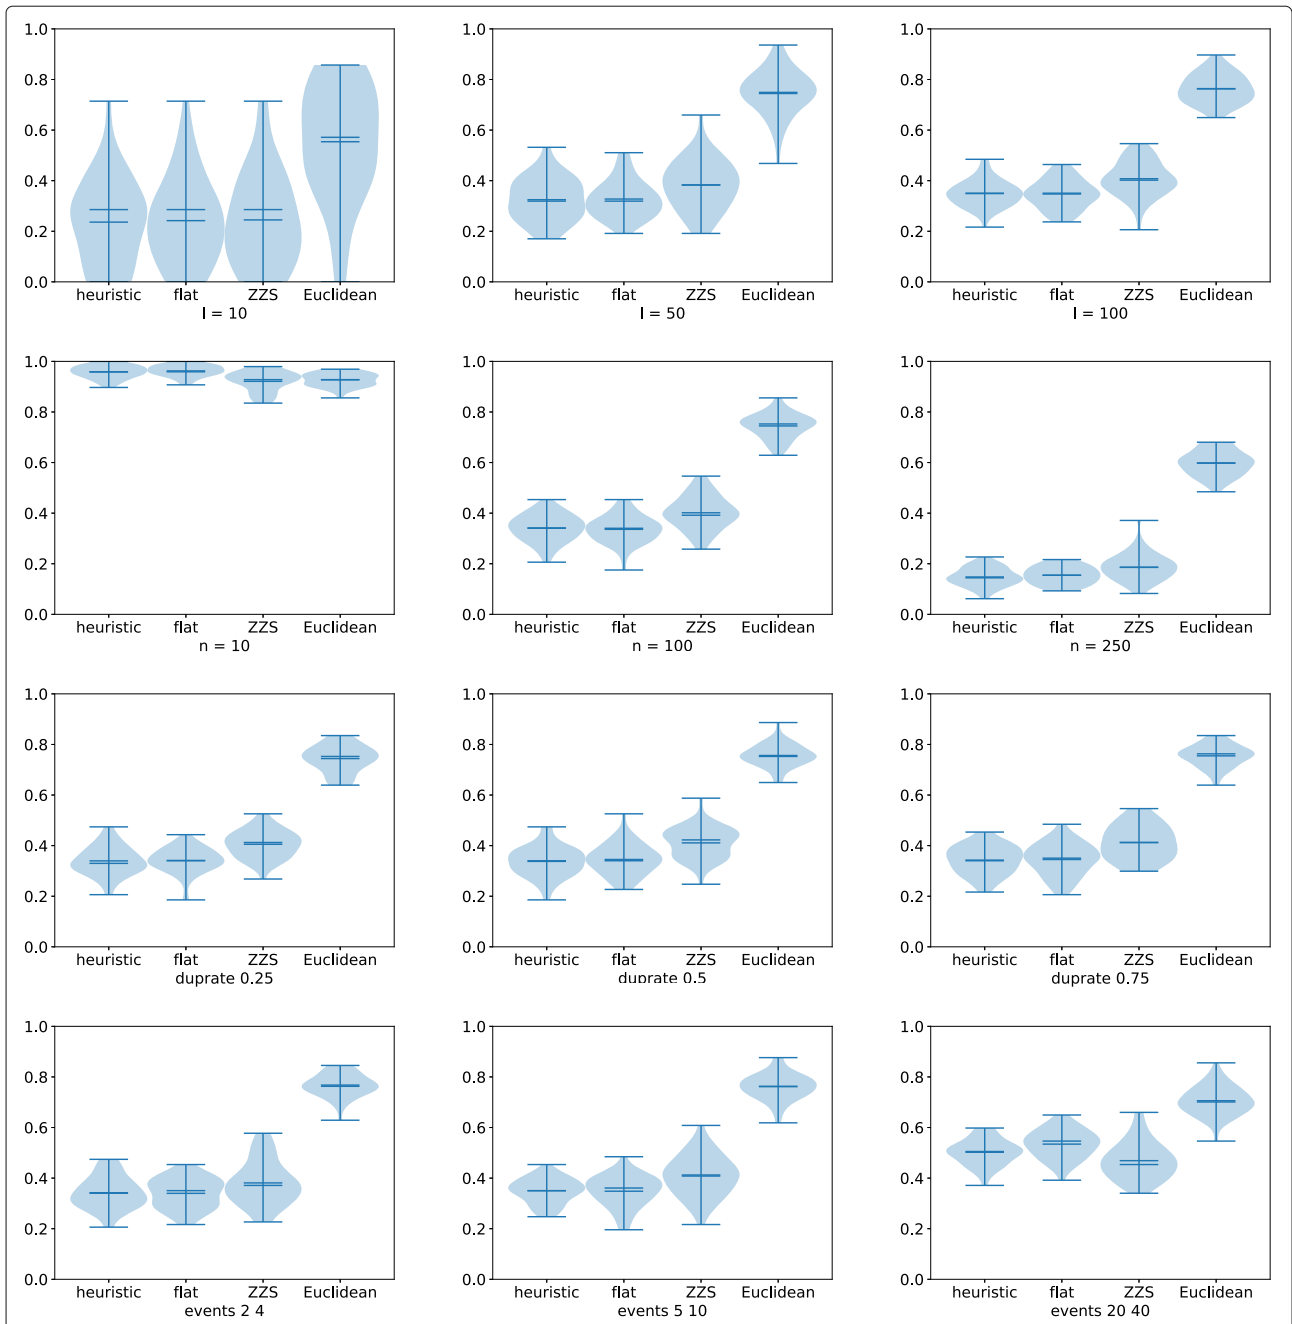


Fig. 3 Violin plots of the normalized RF distances for our improved approximation (heuristic), the algorithm that counts flat intervals (flat), the ZZS algorithm for the MEDICC model (ZZS), and the Euclidean distance on error-free data. Each plot summarizes 50 reconstructed trees with, from left to right: (top row) $l = 10, 50$ and 100 leaves; (second row) $n = 10, 100$ and 250 genes per chromosome; (third row) duplication rate $\Delta = 0.25, 0.5$ and 0.75 ; (fourth row) possible number of events per branch $(e_{min}, e_{max}) = (2, 4), (5, 10)$ and $(20, 40)$. On each row, the other parameters were set to their default as discussed in the text

than the Euclidean distance. The number n of genes is quite important: all the results are poor when each CNP has only $n = 10$ positions, but when $n = 250$, the trees more accurate. This might be because when $n = 10$, there is not enough opportunity for positions acquire a

distinct signature during the evolutionary process, making all distances very similar. This suggests that many genes or segments should be considered when analyzing copy-number variants in tumor clones. The duplication rate does not seem to affect the accuracy of the

methods, whereas accuracy tends to decrease as the number of events per branch increases. We do note that when the number of events per branch is within [5, 10], our model performs better, but when it is high, i.e. in [20, 40], the MEDICC model performs better. This tendency is confirmed under other parameterizations (see the Supplementary Material). Figure 4 show the normalized RF distances when varying parameters q and r . As mentioned before, as q gets closer to 1, the proportion of null positions is around 50–60%, making accurate distance computation difficult for all methods. As for the r parameter, the accuracy of our approach is better when $r = 0.01$ but worse when $r = 0.1$. This tendency can be observed under all parameter combinations (see the Supplementary Material). One should note that accuracy is generally better if the lengths of events are smaller. The four approaches on error-free data exhibit similar behavior on other parameter combinations — additional plots can be found in the Supplementary Materials.

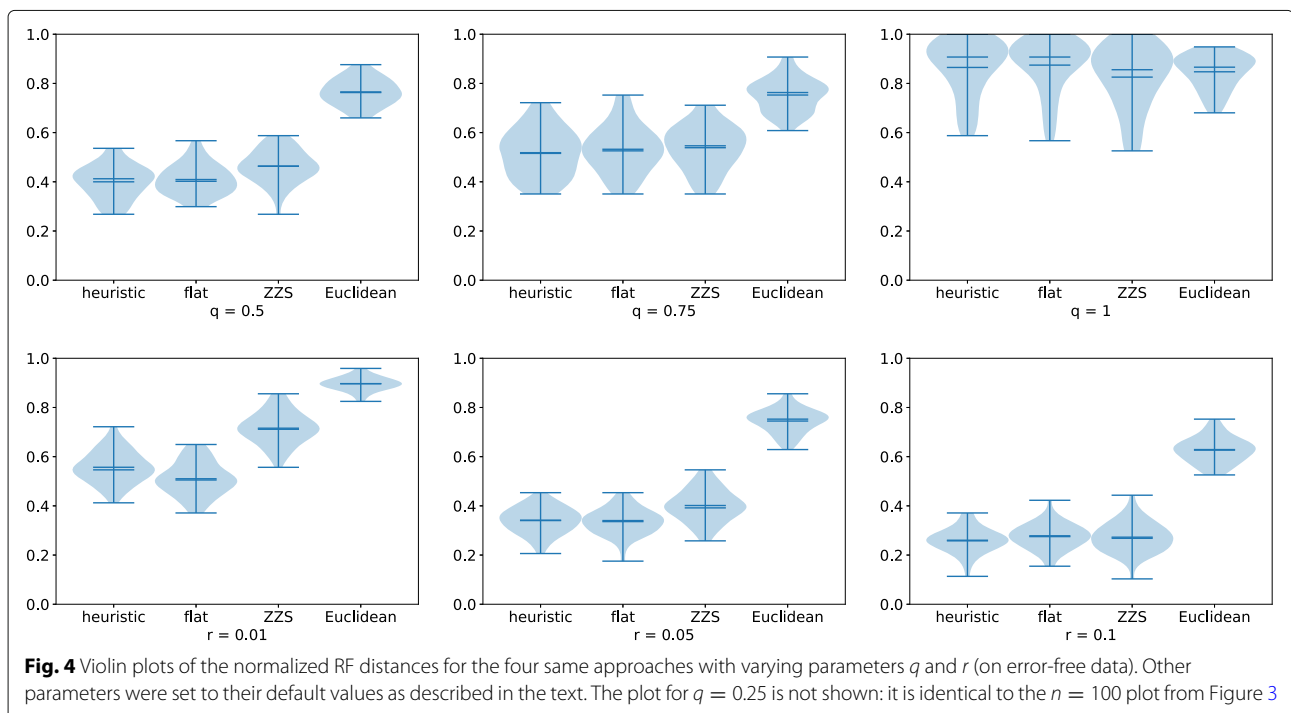
The results on CNPs containing errors are summarized in Fig. 5. We observe that the ZZS algorithm achieves slightly more accurate trees whenever the error rate is above zero. One possible explanation is that a single error in a CNP can split a flat interval into three. This can significantly alter the flat interval counts, whereas the ZZS distance is less dependant on flat intervals. The accuracy of the Euclidean distance appears to be the least affected by error rates and even performs better when $\alpha \geq 0.5$. Observe however that accuracy decays rapidly with error rates: when $\alpha \geq 0.25$, all approaches have an average

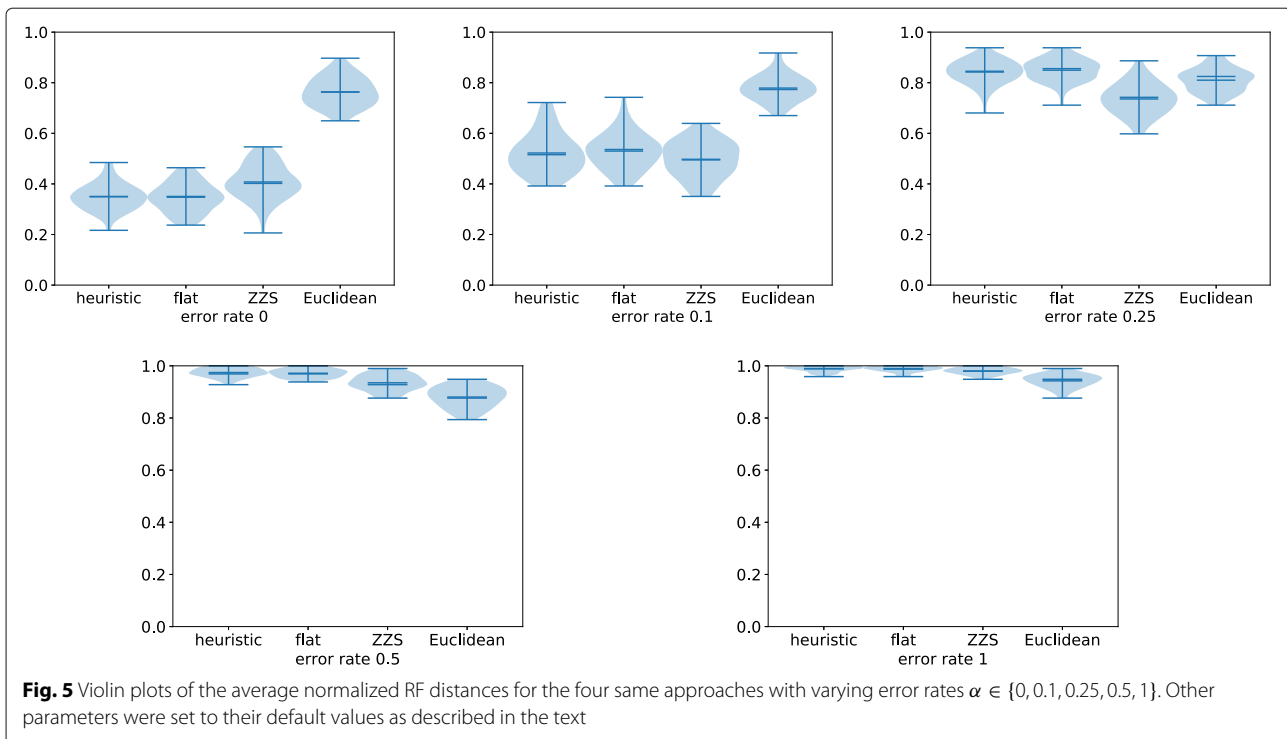
normalized distance above 0.7, casting some doubt on their practical usability in this setting. This suggests that it might be beneficial to apply a CNP error-correction procedure before comparing them (see the “Discussion” section). More results on noisy data can be accessed in the Supplementary Material.

To summarize, the heuristic and flat count algorithm always yield a lower average RF distance than the ZZS algorithm on error-free data, except when $n = 10$ (where the average is always above 0.9 anyways), and every method always outperforms the Euclidean distance. On the other hand, the ZZS approach is slightly more robust to error in the CNP counts. However, the accuracy of the heuristic, the flat count and ZZS drops quickly as error rates increase. Even though the Euclidean distance yields better trees at high error rates, their accuracy is still too poor to be able to draw meaningful conclusions from them. Whether the MEDICC model is better than ours or not, we believe that either should be preferred over the Euclidean distance when reconstructing phylogenies from distance matrices as in [9].

Discussion

The results from the “Experiments” section show that our CNP distance performs reasonably well on simulated data. The incorporation of segmental events into the model does not appear to provide a significant advantage over the unitary events of the ZZS model. However, the simulations suggest that both approaches yield better results than the traditional Euclidean distance. This demonstrates





that either our method or the ZZS algorithm should be preferred as the CNP comparison component in a single-cell phylogenetic reconstruction pipeline.

It should be noted that our algorithms only approximate the true CNP distance whereas the ZZS algorithm provides an exact solution. In order to evaluate the true performance of our segmental model, exact approaches should be developed in the future, perhaps using techniques from the field of parameterized complexity. Moreover, our approaches are very sensitive to errors, even more so than ZZS/MEDICC. One possible explanation for this is that both of our algorithms derive their distance from the number of flat intervals. A single error in a copy-number can turn one flat interval into three, and thus even moderate levels of noise can lead to highly incorrect predictions. We believe that the ZZS algorithm is less sensitive to such errors because that, if copy-number differences are large enough, a small error only increases the distance by 1, which is small in comparison to all the unit events required to handle the high difference. Therefore, even if the true event distance is overestimated, in a comparative setting the relative distances might be closer to the truth. It will be interesting to consider CNP error correction procedures based on flat intervals. For instance, when performing analysis of multiple cells, one could detect a potentially incorrect copy-number of a given segment by checking whether, after altering a predicted copy number by a small amount, several flat intervals get “fixed” when comparing the cell with others.

Another point of interest is that current approaches, including ours and ZZS/MEDICC, ignore rearrangements that change the ordering of segments. Our models assume that the set of contiguous segments remains the same in all cells during evolution. However, when duplications and deletions occur, the relative ordering of genes changes and the set of contiguous genes affected by the events will differ from that in the reference. Inversions, translocations or even chromothripsis also have the same effect. This is a difficult problem to handle if only CNPs are known, since integer vectors do not contain enough information to determine which genes are contiguous or not. One possibility is to ask the following: given two CNPs C_1 and C_2 to compare, choose a genome G_1 whose CNP is C_1 and a genome G_2 whose CNP is C_2 such that the rearrangement distance between G_1 and G_2 is minimized.

Our work also leaves several questions open. From a theoretical perspective, it remains to achieve a constant factor approximation when null positions are present in the input. Moreover, it is unknown whether the *dbl* function admits good approximation algorithms and, more generally, whether there are other biologically plausible functions that should be studied. On another note, it might be interesting to investigate the copy-number triplet problem (see the introduction) under our model, as it allows to define a symmetric distance between CNPs.

On a practical level, phylogenetic approaches that are not distance-based should be investigated. For instance, we could consider maximum parsimony as in [11], where

the objective is to minimize the number of events across branches of a tree. The recent distance-based approach of Xia et al. [12], which is based on the MEDICC model but with an extra error correction step, should also be evaluated in our setting. On another note, it remains to test our approach on real data. We have ignored the problem of calling copy-numbers and the aforementioned problem of phasing. These can introduce noise in the data and, as shown in our experiments, all the evaluated methods are sensitive to errors. This motivates the need for new methods to assign copy-numbers to alleles under our model. Also, our CNP comparison framework assumes a single-cell setting, where the CNP of each individual cell is known. Since bulk sequencing is still commonplace, it will be useful to develop methods that are able to compare genomes extracted from samples that contains multiple types of cells.

Conclusion

In this work, we provided a general framework for the comparison of CNPs depicting genomes that evolve by segmental amplifications and deletions. We have shown that if there is no bound on the number of copies that a deletion can affect, then computing the minimum number of events transforming one CNP into another is strongly NP-hard. One important implication of this result is that unless $P = NP$, one cannot use the fact that copy-numbers are not too large (e.g. under 100) to devise a practical pseudo-polynomial time algorithm, and other solutions must be explored. On the other hand, we proposed two simple and fast approximation algorithms that were shown to perform reasonably well on simulated datasets.

Supplementary information

Supplementary information accompanies this paper at <https://doi.org/10.1186/s12864-020-6611-3>.

Additional file 1: Supplementary file **S33-S1.pdf** contains all the missing proofs.

Additional file 2: Supplementary file **S33-S2.pdf** contains all the additional experimental results.

Abbreviations

BFB: Breakage-fusion-bridge; CGH: Comparative genomic hybridization; CNP: Copy-number profile; NJ: Neighbor-joining; RF: Robinson-Foulds; ZS: Zeira, Zehavi, Shamir

Figure listing

Figure 1: an example of a CNP-to-CNP transformation.
 Figure 2: a visual representation of a staircase and a smooth deletion sequence.
 Figure 3: average normalized RF distances of the four methods evaluated when varying l, n, Δ and $(\epsilon_{min}, \epsilon_{max})$.
 Figure 4: average normalized RF distances of the four methods evaluated when varying q and r .
 Figure 5: average normalized RF distances of the four methods evaluated with varying error rates.

About this supplement

This article has been published as part of *BMC Genomics Volume 21 Supplement 2, 2020: Proceedings of the 17th Annual Research in Computational Molecular Biology (RECOMB) Comparative Genomics Satellite Workshop: genomics*. The full contents of the supplement are available online at <https://bmcbgenomics.biomedcentral.com/articles/supplements/volume-21-supplement-2>.

Authors' contributions

GC and ML both participated in writing the manuscript, establishing the theoretical results, performing the experiments and implementing the algorithms. All authors have read and approved the manuscript.

Funding

Publication was funded by the Natural Sciences and Engineering Research Council (NSERC).

Availability of data and materials

The source code and data are available at: <https://github.com/AEVO-lab/cnp2cnp>.

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare that they have no competing interests.

Author details

¹Department of Computer Science, École polytechnique, Paris, France.

²Department of Computer Science, Université de Sherbrooke, Sherbrooke, Canada.

Published: 16 April 2020

References

- Nowell PC. The clonal evolution of tumor cell populations. *Science*. 1976;194(4260):23–8.
- Jiao W, Vembu S, Deshwar AG, Stein L, Morris Q. Inferring clonal evolution of tumors from single nucleotide somatic mutations. *BMC Bioinformatics*. 2014;15(1):35.
- El-Kebir M, Oesper L, Acheson-Field H, Raphael BJ. Reconstruction of clonal trees and tumor composition from multi-sample sequencing data. *Bioinformatics*. 2015;31(12):62–70.
- Malikic S, McPherson AW, Donmez N, Sahinalp CS. Clonality inference in multiple tumor samples using phylogeny. *Bioinformatics*. 2015;31(9):1349–56.
- Yuan K, Sakoparnig T, Markowitz F, Beerenwinkel N. Bitphylogeny: a probabilistic framework for reconstructing intra-tumor phylogenies. *Genome Biol*. 2015;16(1):36.
- Jahn K, Kuipers J, Beerenwinkel N. Tree inference for single-cell data. *Genome Biol*. 2016;17(1):86.
- Ross EM, Markowitz F. Oncorem: inferring tumor evolution from single-cell sequencing data. *Genome Biol*. 2016;17(1):69.
- El-Kebir M. Sphyr: tumor phylogeny estimation from single-cell sequencing data under loss and error. *Bioinformatics*. 2018;34(17):671–9.
- Navin N, Kendall J, Troge J, Andrews P, Rodgers L, McIndoo J, Cook K, Stepansky A, Levy D, Esposito D, et al. Tumour evolution inferred by single-cell sequencing. *Nature*. 2011;472(7341):90.
- Abo RP, Ducar M, Garcia EP, Thorner AR, Rojas-Rudilla V, Lin L, Sholl LM, Hahn WC, Meyerson M, Lindeman NI, et al. Breakmer: detection of structural variation in targeted massively parallel sequencing data using kmers. *Nucleic Acids Res*. 2014;43(3):19.
- El-Kebir M, Raphael BJ, Shamir R, Sharan R, Zaccaria S, Zehavi M, Zeira R. Copy-number evolution problems: complexity and algorithms. In: *International Workshop on Algorithms in Bioinformatics*. Springer; 2016. p. 137–49.
- Xia R, et al. Phylogenetic Reconstruction for Copy-Number Evolution Problems. *IEEE/ACM transactions on computational biology and bioinformatics*. 2018;16(2):694–699.

13. Zhou J, et al. Maximum parsimony analysis of gene copy number changes. In: International Workshop on Algorithms in Bioinformatics. Berlin: Springer; 2015. p. 108–20.
14. Schwartz R, Schaffer AA. The evolution of tumour phylogenetics: principles and practice. *Nat Rev Genet.* 2017;18(4):213.
15. Lo AW, Sabatier L, Fouladi B, Pottier G, Ricoul M, Mumane JP. Dna amplification by breakage/fusion/bridge cycles initiated by spontaneous telomere loss in a human cancer cell line. *Neoplasia.* 2002;4(6):531–8.
16. Marotta M, Chen X, Inoshita A, Stephens R, Budd GT, Crowe JP, Lyons J, Kondratova A, Tubbs R, Tanaka H. A common copy-number breakpoint of *erbb2* amplification in breast cancer colocalizes with a complex block of segmental duplications. *Breast Cancer Res.* 2012;14(6):150.
17. Rajaram M, Zhang J, Wang T, Li J, Kuscu C, Qi H, Kato M, Grubor V, Weil RJ, Helland A, et al. Two distinct categories of focal deletions in cancer genomes. *PLoS ONE.* 2013;8(6):66264.
18. Liu Y, Chen C, Xu Z, Scuoppo C, Rillahan CD, Gao J, Spitzer B, Bosbach B, Kastenhuber ER, Baslan T, et al. Deletions linked to *tp53* loss drive cancer through *p53*-independent mechanisms. *Nature.* 2016;531(7595):471.
19. Holland AJ, Cleveland DW. Boveri revisited: chromosomal instability, aneuploidy and tumorigenesis. *Nat Rev Mol Cell Biol.* 2009;10(7):478.
20. Desper R, Jiang F, Kallioniemi O-P, Moch H, Papadimitriou CH, Schaffer AA. Inferring tree models for oncogenesis from comparative genome hybridization data. *J Comput Biol.* 1999;6(1):37–51.
21. Liu J, Bandyopadhyay N, Ranka S, Baudis M, Kahveci T. Inferring progression models for *cgh* data. *Bioinformatics.* 2009;25(17):2208–15.
22. Schwarz RF, Trinh A, Sipos B, Brenton JD, Goldman N, Markowitz F. Phylogenetic quantification of intra-tumour heterogeneity. *PLoS Comput Biol.* 2014;10(4):1003535.
23. Fertin G, et al. Combinatorics of genome rearrangements. Cambridge: MIT press; 2009.
24. Zeira R, Zehavi M, Shamir R. A linear-time algorithm for the copy number transformation problem. *J Comput Biol.* 2017;24(12):1179–94.
25. Lafond M, Swenson KM, El-Mabrouk N. An optimal reconciliation algorithm for gene trees with polytomies. In: International Workshop on Algorithms in Bioinformatics. Berlin: Springer; 2012. p. 106–22.
26. Letouzé E, Allory Y, Bollet MA, Radvanyi F, Guyon F. Analysis of the copy number profiles of several tumor samples from the same patient reveals the successive steps in tumorigenesis. *Genome Biol.* 2010;11(7):76.
27. Paul S, Su C, Pang J, Mizera A. A Decomposition-Based Approach towards the Control of Boolean Networks. In: Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics. New York: Association for Computing Machinery; 2018. p. 11–20. <https://doi.org/10.1145/3233547.3233550>.
28. Santarius T, Shipley J, Brewer D, Stratton MR, Cooper CS. A census of amplified and overexpressed human cancer genes. *Nat Rev Cancer.* 2010;10(1):59.
29. Park HS, Jang MH, Kim EJ, Kim HJ, Lee HJ, Kim YJ, Kim JH, Kang E, Kim S-W, Kim IA, et al. High *egfr* gene copy number predicts poor outcome in triple-negative breast cancer. *Mod Pathol.* 2014;27(9):1212.
30. Campbell K, Gastier-Foster JM, Mann M, Naranjo AH, Van Ryn C, Bagatell R, Matthay KK, London WB, Irwin MS, Shimada H, et al. Association of *mycn* copy number with clinical features, tumor biology, and outcomes in neuroblastoma: A report from the children's oncology group. *Cancer.* 2017;123(21):4224–35.
31. Garey MR, Johnson DS. *Computers and Intractability*, vol 29. New York: W.H. Freeman; 2002.
32. Seidel R, Aragon CR. Randomized search trees. *Algorithmica.* 1996;16(4-5):464–97.
33. Saitou N, Nei M. The neighbor-joining method: a new method for reconstructing phylogenetic trees. *Molecular Biol Evol.* 1987;4(4):406–25.
34. Felsenstein J. PHYLIP (phylogeny Inference Package), Version 3.5 C: Joseph Felsenstein.; 1993.
35. Aldous D. Probability distributions on cladograms. In: *Random Discrete Structures*. New York: Springer; 1996. p. 1–18.
36. Robinson DF, Foulds LR. Comparison of phylogenetic trees. *Math Biosci.* 1981;53(1-2):131–47.

Publisher's Note

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

Ready to submit your research? Choose BMC and benefit from:

- fast, convenient online submission
- thorough peer review by experienced researchers in your field
- rapid publication on acceptance
- support for research data, including large and complex data types
- gold Open Access which fosters wider collaboration and increased citations
- maximum visibility for your research: over 100M website views per year

At BMC, research is always in progress.

Learn more biomedcentral.com/submissions

