

Assessing the robustness of parsimonious predictions for gene neighborhoods from reconciled phylogenies

Ashok Rajaraman, Cedric Chauve, Yann Ponty

► **To cite this version:**

Ashok Rajaraman, Cedric Chauve, Yann Ponty. Assessing the robustness of parsimonious predictions for gene neighborhoods from reconciled phylogenies. Robert Harrison; Yaohang Li; Ion Mandoiu, ISBRA - 11th International Symposium on Bioinformatics Research and Applications - 2015, Jun 2015, Norfolk, Virginia, United States. LNCS, 9096, pp.260–271, 2015, <10.1007/978-3-319-19048-8_22>. <hal-01104587v3>

HAL Id: hal-01104587

<https://hal.inria.fr/hal-01104587v3>

Submitted on 19 Mar 2015

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.



Assessing the robustness of parsimonious predictions for gene neighborhoods from reconciled phylogenies

Ashok Rajaraman¹, Cedric Chauve¹, and Yann Ponty^{1,2,3}

¹ Department of Mathematics, Simon Fraser University, Burnaby, Canada

² Pacific Institute for Mathematical Sciences, CNRS UMI3069, Vancouver, Canada

³ CNRS/LIX, Ecole Polytechnique, Palaiseau, France

Abstract. The availability of many assembled genomes opens the way to study the evolution of syntenic character within a phylogenetic context. The DeCo algorithm, recently introduced by Bérard *et al.*, computes parsimonious evolutionary scenarios for gene adjacencies, from pairs of reconciled gene trees. Following the approach pioneered by Sturmfels and Pachter, we describe how to modify the DeCo dynamic programming algorithm to identify classes of cost schemes that generate similar parsimonious evolutionary scenarios for gene adjacencies. We also describe how to assess the robustness, again to changes of the cost scheme, of the presence or absence of specific ancestral gene adjacencies in parsimonious evolutionary scenarios. We apply our method to six thousands mammalian gene families, and show that computing the robustness to changes of cost schemes provides interesting insights on the DeCo model.

1 Introduction

Reconstructing evolutionary histories of genomic characters along a given species phylogeny is a long-standing problem in computational biology. This problem has been studied for several types of genomic characters (DNA sequences and gene content for example), for which efficient algorithms exist to compute parsimonious evolutionary scenarios. Recently, Bérard *et al.* [2] extended the corpus of such results to syntenic characters. They defined a model for the evolution of gene adjacencies within a species phylogeny, together with an efficient dynamic programming (DP) algorithm, called DeCo, to compute parsimonious evolutionary histories that minimize the total cost of gene adjacencies gain and break, for a given cost scheme associating a cost to each of these two events. Reconstructing evolutionary scenarios for syntenic characters is an important step towards more comprehensive models of genome evolution, going beyond classical sequence/ content frameworks, as it implicitly integrates genome rearrangements [5]. Application of such methods include the study of genome rearrangement rates and the reconstruction of ancestral gene order. Moreover, DeCo is the only existing tractable model that considers the evolution of gene adjacencies within a general phylogenetic framework; so far other tractable models of genome rearrangements accounting for a given species phylogeny are either limited to single-copy genes and ignore gene-specific events [3, 18], assume restrictions on the gene duplication events, such as considering only

whole-genome duplication (see [7] and references there), or require a dated species phylogeny [11].

The evolutionary events considered by DeCo, gene adjacency gain and break caused by genome rearrangement, are rare evolutionary events compared to gene-family specific events. It is then important to assess the robustness of inferences made by DeCo, whether it is of a parsimony cost or of an individual feature such as the presence of a specific ancestral adjacency. We recently explored an approach that considers the set of all possible evolutionary scenarios under a Boltzmann probability distribution for a fixed cost scheme [6]. A second approach consists of assessing how robust features of evolutionary scenarios are to changes in the cost associated to evolutionary events (the cost scheme). Such approaches have recently been considered for the gene tree reconciliation problem and have been shown to significantly improve the results obtained from purely parsimonious approaches [1, 10]. This relates to the general problem of deciding the precise cost to assign to evolutionary events in evolutionary models, a recurring question in the context of parsimony-based approaches in phylogenetics.

This motivates the precise questions tackled in this work. First, how robust is a parsimonious evolutionary scenario to a change of the costs associated to adjacency gains and breaks? Similarly, how robust is an inferred parsimonious gene adjacency to a change in these costs? We address this problem using a methodology that has been formalized into a rigorous algebraic framework by Pachter and Sturmfels [15, 14, 13], that we refer to as the *polytope approach*. Its main features, summarized in Fig. 1 for assessing the robustness of evolutionary scenarios, are (1) associating each evolutionary scenario to a *signature*, a vector of two integers (g, b) where g is the number of adjacency gains and b the number of adjacency breaks; and (2) partitioning the space of cost schemes into convex regions such that, for all the cost schemes within a region, all optimal solutions obtained with such cost schemes have the same signature. This partition can be computed by an algorithm that is a direct translation of the DP algorithm into a polytope framework. Furthermore, the same framework can be extended to assess the robustness of inferred parsimonious ancestral adjacencies.

2 Preliminary: models and problems

A *phylogeny* is a rooted tree which describes the evolutionary relationships of a set of elements (species, genes, ...) represented by its nodes: internal nodes correspond to ancestral elements, leaves to extant elements, and edges represent direct descents between parents and children. For a node v of a phylogeny, we denote by $s(v)$ the species it belongs to. For a tree T and a node x of T , we denote by $T(x)$ the subtree rooted at x . If x is an internal node, we assume it has either one child, denoted by $a(x)$, or two children, denoted by $a(x)$ and $b(x)$.

Species tree and reconciled gene trees. A species tree S is a binary tree that describes the evolution of a set of species from a common ancestor through the mechanism of *speciation*. A reconciled gene tree is a binary tree that describes the evolution of




Fig. 1. Outline of our method for assessing the robustness of an evolutionary scenario: Starting from two reconciled gene trees and a set of extant adjacencies (a.), the polytope of parsimonious signatures is computed (b.). Its normal vectors define a segmentation of the space of cost schemes into cones (c.), each associated with a signature. Here, the positive quadrant is fully covered by a single cone, meaning that the parsimonious prediction does not depend on the precise cost scheme. In general (d.), the robustness of a prediction (here, obtained using the (1, 1) scheme) to perturbations of the scheme can be measured as the smallest angle θ such that a cost scheme at angular distance θ no longer predicts the signature (a, b) .

a set of genes, called a *gene family*, within a given species tree S , through the evolutionary mechanisms of speciation, *gene duplication* and *gene loss*. Therefore, each leaf of a gene tree G represents either a gene loss or an extant gene, while each internal node represents an ancestral gene. In a reconciled gene tree, we associate every ancestral gene (an internal node g) to an evolutionary event $e(g)$ that leads to the creation of the two children $a(g)$ and $b(g)$: $e(g)$ is a speciation (denoted by Spec) if the species pair $\{s(a(g)), s(b(g))\}$ is equal to the species pair $\{s(s(g)), b(s(g))\}$, $s(a(g)) \neq s(b(g))$, or a gene duplication (GDup) if $s(a(g)) = s(b(g)) = s(g)$. If g is a leaf, then $e(g)$, as stated before, indicates either a gene loss (GLoss) or an extant gene (Extant), in which case $e(g)$ is not an evolutionary event *stricto sensu*. A *pre-speciation* ancestral gene is an internal node g such that $e(g) = \text{Spec}$. See Fig. 2 for an illustration.

Adjacency trees and forests. We consider now that we are given two reconciled gene trees G_1 and G_2 , representing two gene families evolving within a species tree S . A *gene adjacency* is a pair of genes (one from G_1 and one from G_2) that appear consecutively along a chromosome, for a given species, ancestral or extant. Gene adjacencies evolve within a species tree S through the evolutionary events of speciation, gene duplication, gene loss (these three events are modeled in the reconciled gene trees), and *adjacency duplication* (ADup), *adjacency loss* (ALoss) and *adjacency break* (ABreak), that are adjacency-specific events.

Following the model introduced in [2], we represent such an evolutionary history using an *adjacency forest*, composed of *adjacency trees*. An adjacency tree represents the evolution of an ancestral gene adjacency (located at the root of the tree) through the following events: (1) The duplication of an adjacency $\{g_1, g_2\}$, where g_1 and g_2 are respectively genes from G_1 and G_2 such that $s(g_1) = s(g_2)$, follows from the simultaneous duplication of both its genes g_1 and g_2 (so $e(g_1) = e(g_2) = \text{GDup}$), resulting in the creation of two distinct adjacencies each belonging to $\{a(g_1), b(g_1)\} \times$

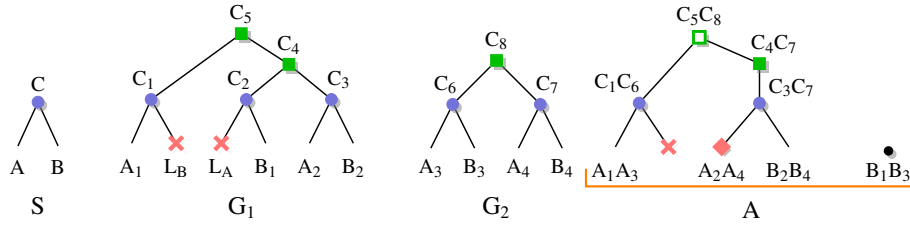


Fig. 2. A species tree S , with two extant species A and B and an ancestral species C . Two reconciled gene trees G_1 and G_2 , with four extant genes in genome A , four extant genes in genome B and three ancestral genes in genome C . The set of extant gene adjacencies is $(A_1 A_3, B_1 B_3, B_2 B_4)$. An adjacency forest A composed of two adjacency trees. Blue dots represent speciation nodes. Leaves are extant species/genes/adjacencies, except the one labeled by a red cross (gene loss) or a red diamond (adjacency breaks). Green squares are (gene or adjacency) duplication nodes. Gene labels refer to the species they belong to. Every node of the adjacency tree is labeled by a gene adjacency. Figure adapted from [2].

$\{a(g_2), b(g_2)\}$; (2) The loss of an adjacency, which can occur due to several events, such as the loss of exactly one of its genes (gene loss, GLoss), the loss of both its genes (adjacency loss, ALoss) or a genome rearrangement that breaks the contiguity between the two genes (adjacency break, ABreak); (3) The creation/gain of an adjacency (denoted by AGain), for example due to a genome rearrangement, that results in the creation of a new adjacency tree whose root is the newly created adjacency.

With this model, one can model the evolution of two gene families along a species phylogeny by a triple (G_1, G_2, A) : G_1 and G_2 are reconciled gene trees representing the evolution of these families in terms of gene-specific events and A is an adjacency forest consistent with G_1 and G_2 . Similar to species trees and reconciled gene trees, internal nodes of an adjacency tree are associated to ancestral adjacencies, while leaves are associated to extant adjacencies or lost adjacencies (due to a gene loss, adjacency loss or adjacency break), and are labeled by evolutionary events. The label $e(v)$ of an internal node v of an adjacency forest A belongs to $\{\text{Spec}, \text{GDup}, \text{ADup}\}$, while the label $e(v)$ of a leaf belongs to $\{\text{Extant}, \text{GLoss}, \text{ALoss}, \text{ABreak}\}$, as shown in Fig. 2.

Signatures, descriptors and parsimonious scenarios. The *signature* of an adjacency forest A is an ordered pair of integers $\sigma(A) = (g_A, b_A)$ where g_A (resp. b_A) is the number of adjacency gains (resp. adjacency breaks) in A . A *cost scheme* is a pair $\mathbf{x} = (x_0, x_1)$ of non-negative real numbers, where x_0 is the cost of an adjacency gain and x_1 the cost of an adjacency break. The *cost* of an adjacency forest A for a given cost scheme \mathbf{x} is the number $S(A) = x_0 \times g_A + x_1 \times b_A$. The adjacency forest A in an evolutionary scenario (G_1, G_2, A) is *parsimonious for \mathbf{x}* if there is no other evolutionary scenario (G_1, G_2, B) such that $S(B) < S(A)$. The signature the adjacency forest A in Fig. 2 is $(1, 1)$, and this adjacency forest is parsimonious for the cost scheme $(1, 1)$.

A *descriptor* of a scenario is a boolean or integer valued feature of the solution which does not contribute to the cost of the scenario, but rather represents a feature of a scenario. For instance, the presence/absence of an ancestral adjacency in a given

adjacency forest A can be described as a boolean. Given k descriptors a_1, \dots, a_k , we define an *extended signature* of a scenario A as a tuple $\sigma_{a_1, \dots, a_k}(A) = (g, b, s_{a_1}, \dots, s_{a_k})$, where g, b are the numbers of adjacency gains and breaks in A respectively, and s_{a_i} is the value of the descriptor a_i for A .

The DeCo algorithm. Bérard *et al.* [2] showed that, given a pair of reconciled gene trees G_1 and G_2 , a list L of extant gene adjacencies, and a cost scheme \mathbf{x} , one can use a DP algorithm to compute an evolutionary scenario (G_1, G_2, A) , where A is a parsimonious adjacency forest such that L is exactly the set of leaves of A labeled Extant. The DeCo algorithm computes, for every pair of nodes g_1 (from G_1) and g_2 (from G_2) such that $s(g_1) = s(g_2)$, two quantities $c_1(g_1, g_2)$ and $c_0(g_1, g_2)$, that correspond respectively to the cost of a parsimonious adjacency forest for the pairs of subtrees $G(g_1)$ and $G(g_2)$, under the hypothesis that g_1 and g_2 form (for c_1) or do not form (for c_0) an ancestral adjacency. As usual in dynamic programming along a species tree, the cost of a parsimonious adjacency forest for G_1 and G_2 is given by $\min(c_1(r_1, r_2), c_0(r_1, r_2))$ where r_1 is the root of G_1 and r_2 the root of G_2 . In [6], we recently generalized DeCo into a DP algorithm DeClone that allows one to explore the space of all possible adjacency evolutionary scenarios for a given cost scheme.

Robustness problems. The first problem we are interested in is the *signature robustness problem*. A signature $\sigma = (g, b)$ is parsimonious for a cost scheme \mathbf{x} if there exists at least one adjacency forest A that is parsimonious for \mathbf{x} and has signature $\sigma(A) = \sigma$. The robustness of the signature σ is defined as the difference between \mathbf{x} and the closest cost scheme for which σ is no longer parsimonious. To measure this difference, we rely on a geometric representation of a cost scheme. Assuming that a cost scheme $\mathbf{x} = (x_0, x_1) \in \mathbb{R}^2$ provides sufficient information to evaluate the cost of an adjacency forest, the predictions under such a model remain unchanged upon multiplying \mathbf{x} by any positive number, allowing us to assume that $\|\mathbf{x}\| = 1$ without loss of generality. So $\mathbf{x} = (x_0, x_1)$ can be summarized as an angle θ (expressed in radians), and the difference between two cost schemes is indicated by their associated angular distance.

However, signatures only provide a quantitative summary of the evolutionary events described by a parsimonious adjacency forest. In particular, signatures discard any information about predicted sets of ancestral adjacencies. We address the robustness of inferred parsimonious adjacencies through the *parsimonious adjacency robustness problem*. Let $a = (g_1, g_2)$ be an ancestral adjacency featured in a parsimonious adjacency forest for a cost scheme \mathbf{x} . We say that a is parsimonious for a cost scheme \mathbf{y} if a belongs to every adjacency forest that is parsimonious for \mathbf{y} . The robustness of a is defined as the angular distance from \mathbf{x} to the closest cost scheme \mathbf{y} for which a is no longer parsimonious.

3 Methods

If the signature for a given adjacency forest A is given by the vector $\sigma(A) = (g, b)$, and the cost scheme is given by the vector $\mathbf{x} = (x_0, x_1)$, then the parsimony cost of DeCo

can be written as the inner product $\langle \mathbf{x}, \sigma(A) \rangle = g \times x_0 + b \times x_1$. DeCo computes the following quantity for a pair of gene trees G_1 and G_2 .

$$c(G_1, G_2) = \min_{A \in \mathcal{F}(G_1, G_2)} \langle \mathbf{x}, \sigma(A) \rangle, \quad (1)$$

where $\mathcal{F}(G_1, G_2)$ denotes the set of all possible adjacency forests that can be constructed from G_1 and G_2 , irrespective of the cost scheme.

For a given adjacency forest A , we will consider a single descriptor a , indicating the presence or absence of an ancestral adjacency $a = (g_1, g_2) \in G_1 \times G_2$ in A , where $s_a = 1$ if it is present in A , and 0 otherwise. Since, by definition, a descriptor does not contribute to the cost, when considering the robustness of specific adjacencies, we will consider cost schemes of the form $\mathbf{x} = (x_0, x_1, 0)$, and DeCo will compute Eq. (1) as usual.

For a given cost scheme \mathbf{x} , two adjacency forests A_1 and A_2 such that $\sigma(A_1) = \sigma(A_2)$ will have the same associated cost. We can thus define an equivalence class in $\mathcal{F}(G_1, G_2)$ based on the signatures. However, for a given potential ancestral adjacency $a = (g_1, g_2) \in G_1 \times G_2$, the adjacency forests in this equivalence class may have different extended signatures, differing only in the last coordinate. Thus, there may be two adjacency forests A_1 and A_2 with extended signatures $(g, b, 1)$ and $(g, b, 0)$ respectively, and they will have the same cost for all cost schemes. Evolutionary scenarios with the same extended signature also naturally form an equivalence class in $\mathcal{F}(G_1, G_2)$.

Convex polytopes from signatures. Let us denote the set of signatures of all scenarios in $\mathcal{F}(G_1, G_2)$ by $\sigma(\mathcal{F}(G_1, G_2))$, and the set of extended signatures for a given adjacency a by $\sigma_a(\mathcal{F}(G_1, G_2))$. Each of these is a point in \mathbb{R}^d , where $d = 2$ for signatures and $d = 3$ for extended signatures. In order to explore the parameter space of parsimonious solutions to DeCo, we use these sets of points to construct a *convex polytope* in \mathbb{R}^d . A convex polytope is simply the set of all convex combinations of points in a given set, in this case the set of signatures or extended signatures [15]. Thus, for each pair of gene trees G_1, G_2 and a list of extant adjacencies, we can theoretically construct a convex polytope in \mathbb{R}^2 by taking the convex combinations of all signatures in $\sigma(\mathcal{F}(G_1, G_2))$. This definition generalizes to a convex polytope in \mathbb{R}^3 when extended signatures $\sigma_a(\mathcal{F}(G_1, G_2))$ are considered for some ancestral adjacency a . Viewing the set of evolutionary scenarios as a polytope allows us to deduce some useful properties:

1. Any (resp. extended) signature that is parsimonious for some cost scheme \mathbf{x} lies on the surface of the polytope;
2. If a (resp. extended) signature is parsimonious for two cost schemes \mathbf{x} and \mathbf{x}' , then it is also parsimonious for any cost scheme *in between* (i.e. for any convex combination of \mathbf{x} and \mathbf{x}').

Traditionally, a polytope is represented as a set of inequations, which is inappropriate for our intended application. Therefore, we adopt a slightly modified representation, and denote the polytope of $\mathcal{F}(G_1, G_2)$ as the list of signatures that are represented within $\mathcal{F}(G_1, G_2)$ and lie on the convex hull of the polytope.

A *vertex* in a polytope is a signature (resp. extended signature) which is parsimonious for some cost scheme. The domain of parsimony of a vertex \mathbf{v} is the set of cost schemes for which \mathbf{v} is parsimonious. From Property 2, the domain of parsimony for a vertex \mathbf{v} is a *cone* in \mathbb{R}^d , formally defined as:

$$\text{Cone}(\mathbf{v}) = \left\{ \mathbf{x} \in \mathbb{R}^d : \langle \mathbf{x}, \mathbf{v} \rangle \leq \langle \mathbf{x}, \mathbf{w} \rangle \forall \mathbf{w} \in P \right\}. \quad (2)$$

The set of cones associated with the vertices of a polytope form a partition of the cost schemes space [15], which allows us to assess the effect of perturbing the cost scheme on the optimal solution of DeCo for this cost scheme.

Computing the polytope. Building on earlier work on parametric sequence alignment [8], Pachter and Sturmfels [14, 15] described the concept of *polytope propagation*, based on the observation that the polytope of a DP (minimization) scheme can be computed through an algebraic substitution. Accordingly, any point that lies strictly within the polytope is suboptimal for any cost scheme, and can be safely discarded by a procedure that repeatedly computes the *convex hull* $H(P)$ of the (intermediates) polytopes produced by the modified DP scheme. In the context of the DeCo DP scheme, the precise modifications are:

1. Any occurrence of the $+$ operator is replaced by \oplus , the (convex) *Minkowski sum* operator, defined for P_1, P_2 two polytopes as

$$P_1 \oplus P_2 = H(\{p_1 + p_2 \mid (p_1, p_2) \in P_1 \times P_2\});$$

2. Any occurrence of the min operator is replaced by \cup , the *convex union* operator, defined for P_1, P_2 two polytopes as

$$P_1 \cup P_2 = H(P_1 \cup P_2);$$

3. Any occurrence of an *adjacency gain* cost is replaced by the vector $(1, 0)$ (resp. $(1, 0, 0)$ for extended signatures);
4. Any occurrence of an *adjacency break* cost is replaced by the vector $(0, 1)$ (resp. $(0, 1, 0)$ for extended signatures);
5. (Extended signatures only) An event that corresponds to the prediction of a fixed ancestral adjacency a in a scenario is replaced by the vector $(0, 0, 1)$;

By making this substitution, we can efficiently compute the polytope associated with two input gene trees G_1 and G_2 , having sizes n_1 and n_2 respectively, through $O(n_1 \times n_2)$ executions of the convex hull procedure. In place of the integers used by the original minimization approach, intermediate convex polytopes are now processed by individual operations, and stored in the DP tables, so the overall time and space complexities of the algorithm critically depend on the size of the polytopes, i.e. its number of vertices. Pachter and Sturmfels proved that, in general, the number of vertices on the surface of the polytope is $O(n^{d-1})$, where d is the number of dimensions, and n is the size of the DP table. In our case, the number of vertices in the 2D polytope associated with simple signatures is in $O(n_1 \times n_2)$. This upper bound also holds for extended signatures, as the third coordinate is a boolean, and the resulting 3D polytope is in fact the union of two 2D polytopes. The total cost of computing the polytope is therefore bounded by $O(n_1^2 \times n_2^2 \times \log(n_1 \times n_2))$, e.g. using Chan's convex

hull algorithm [4]. As for the computation of the cones, let us note that the cone of a vertex v in a given polytope P is fully delimited by a set of vectors, which can be computed from P as the normal vectors, pointing towards the center of mass of P , of each of the facets in which v appears. This computation can be performed as a postprocessing using simple linear algebra, and its complexity will remain largely dominated by that of the DP-fuelled polytope computation.

Assessing signature and adjacency robustness. The cones associated with the polytope of a given instance cover all the real-valued cost schemes, including those associating negative costs to events. These later cost schemes are not valid, and so, we only consider cones which contain at least one positive cost scheme. Given a fixed cost scheme \mathbf{y} , the vertex associated to the cone containing this cost scheme corresponds to the signature of all parsimonious scenarios for this cost scheme. In order to assess the robustness of this signature, we can calculate the smallest angular perturbation needed to move from \mathbf{y} to a cost scheme whose parsimonious scenarios do not have this signature. This is simply the angular distance from \mathbf{y} to the nearest boundary of the cone which contains it. Using this methods, we assign a numerical value to the robustness of the signatures of parsimonious scenarios on a number of instances for a particular cost scheme.

In the case of extended signatures $\sigma_a(\mathcal{F}(G_1, G_2))$ for an adjacency a , the polytope is 3-dimensional. The cones associated with the vertices, as defined algebraically, now partition \mathbb{R}^3 , the set of cost schemes (x_0, x_1, x_2) , where x_2 indicates the cost of a distinguished adjacency. Since the third coordinate is a descriptor, it does not contribute to the cost scheme, and we therefore restrict our analysis to the $\mathbb{R}^+ \times \mathbb{R}^+ \times \{0\}$ subset of the cost scheme space. Precisely, we take the intersection of the plane $x_2 = 0$ with each cone associated to a vertex (g, b, s_a) , and obtain the region in which the extended signature (g, b, s_a) is parsimonious. This region is a 2D cone.

However, the cost of an extended signature is independent of the entry in its last coordinate, and there may exist two different extended signatures $(g, b, 0)$ and $(g, b, 1)$, both parsimonious for all the cost schemes found in the 2D cone. It is also possible for adjacent cones to have different signatures, yet feature a given adjacency. The robustness of a given adjacency a is computed from the cones using a greedy algorithm which, starting from the cone containing \mathbf{x} , explores the adjacent cones in both directions (clockwise/counter-clockwise) until it finds one that no longer predicts a , i.e. is associated with at least one signature $(g', b', 0)$.

4 Results

We considered 5,039 reconciled gene trees and 50,389 extant gene adjacencies, forming 6,074 DeCo instances, with genes taken from 36 extant mammalian genomes from the Ensembl database in 2012. In [2], this data was analyzed with DeCo, using the cost scheme (1,1), that computed a single parsimonious adjacency forest per instance. These adjacency forests defined 96,482 ancestral adjacencies (adjacencies between two pre-speciation genes from the same ancestral species), covering 112,188 ancestral genes.

We first considered all 6,074 instances, and computed for each signature the robustness of the parsimonious signature obtained with the cost scheme (1,1). Interestingly, we observe (Fig. 3(A)) that for more than half of the instances, the parsimonious signature is robust to a change of cost scheme, as the associated cone is the complete first quadrant of the real plane. On the other hand, for 945 instances the parsimonious signature for the cost scheme (1,1) is not robust to any change in the cost scheme; these cases correspond to interesting instances where the cost scheme (1,1) lies at the border of two cones, meaning that two parsimonious signatures exist for the cost scheme (1,1), and any small change of cost scheme tips the balance towards one of these two signatures. More generally, as revealed by Fig. 3(A), we observe an extreme robustness of parsimonious signatures: there is a $\sim 80\%$ overlap between the sets of signatures that are parsimonious for any (positive) cost scheme, and for the (1,1) cost scheme. This observation supports the notion of a sparsely-populated search space for attainable signatures. In this vision, signatures are generally isolated, making it difficult to trade adjacency gains for breaks (or vice-versa) in order to challenge the (1,1)-parsimonious prediction. We hypothesize that such a phenomenon is essentially combinatorial, as extra adjacency gains typically lead, through duplications to more subsequent adjacency breaks.

Next, to evaluate the stability of the total number of evolutionary events inferred by parsimonious adjacency forests, we recorded two counts of evolutionary events for each instance: the number of syntenic events (adjacencies gains and breaks) of the parsimonious signature (called the *parsimonious syntenic events count*), and the maximum number of syntenic events taken over all signatures that are parsimonious for some cost scheme (called the *maximum syntenic events count*). We observe that the average parsimonious (resp. maximum) syntenic events count is 1.25 (resp. 1.66). This shows a strong robustness of the (low) number of syntenic events to changes in the cost scheme.

We then considered the robustness of individual ancestral adjacencies. Using the variant DeClone of DeCo that explores the set of all evolutionary scenarios [6], we extracted, for each instance, the set of ancestral adjacencies that belong to all parsimonious solutions for the cost scheme (1,1), and computed their robustness as defined in the previous sections. This set of ancestral adjacencies contains 87,019 adjacencies covering 106,903 ancestral genes. The robustness of these adjacencies is summarized in Fig. 3(B, left and center columns). It is interesting to observe that few adjacencies have a low robustness, while, conversely, a large majority of the universally parsimonious adjacencies are completely robust to a change of cost scheme (97,593 out of 106,639). This suggests that the DeCo model of parsimonious adjacency forests is robust, and infers highly supported ancestral adjacencies, which is reasonable given the relative sparsity of genome rearrangements in evolution compared to smaller scale evolutionary events.

Besides the notions of robustness, an indirect validation criterion used to assess the quality of an adjacency forest is the limited presence of *syntenic conflicts*. An ancestral gene is said to participate in a syntenic conflict if it belongs to three or more ancestral adjacencies, as a gene can only be adjacent to at most two neighboring genes along a chromosome. An ancestral adjacency participates in a syntenic con-

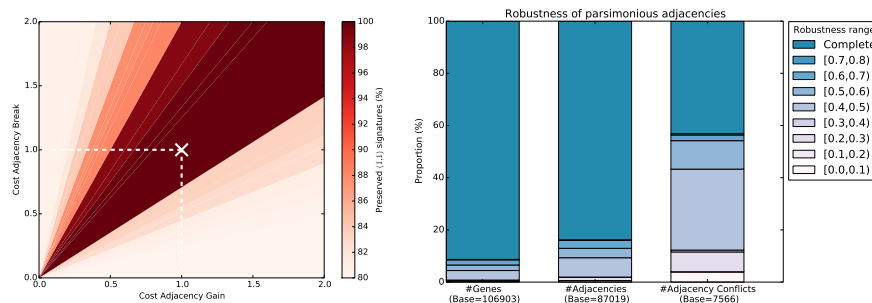


Fig. 3. (A) Average robustness of signatures predicted using the (1,1) cost scheme. At each point (x, y) , the colour indicates the proportion of signatures that are parsimonious, and therefore predicted, for the (1,1) cost scheme, and remain parsimonious for the (x, y) cost scheme. (B) Universally parsimonious adjacencies and syntenic conflicts. (Left) Percentage of ancestral genes present in universally parsimonious adjacencies per level of minimum robustness of the adjacencies, expressed in radians. (Center) Percentage of universally parsimonious adjacencies per level of minimum robustness. (Right) Percentage of conserved conflicting adjacencies per level of minimum robustness.

flict if it contains a gene that does. Among the ancestral adjacencies inferred by DeCo, 16,039 participate in syntenic conflicts, covering 5,817 ancestral genes. This represents a significant level of syntenic conflict and a significant issue in using DeCo to reconstruct ancestral gene orders. It was observed that selecting universally parsimonious ancestral adjacencies, as done in the previous analysis, significantly reduced the number of syntenic conflicts, as almost all discarded ancestral adjacencies participated in syntenic conflicts. Considering syntenic conflicts, we observe (Fig. 3(B, right column)) a positive result, i.e. that filtering by robustness results in a significant decrease of the ratio of conflicting adjacencies. However, even with robust universally parsimonious ancestral adjacencies, one can observe a significant number of adjacencies participating in syntenic conflicts. We discuss these observations in the next section.

5 Discussion and Conclusion

From an application point of view, the ability to exhaustively explore the parameter space leads to the observation that, on the considered instances, the DeCo model is extremely robust. Even taking parsimonious signatures that maximize the number of evolutionary syntenic events (i.e. considering cost schemes that lead to the maximum number of events) results in an average increase of roughly 33% events (1.25 to 1.66), and stays very low, much lower than gene specific events such as gene duplications (average of 3.38 event per reconciled gene tree). This is consistent with the fact that for rare evolutionary events such as genome rearrangements, a parsimony approach is relevant, especially when it can be complemented by efficient algorithms to explore slightly sub-optimal solutions, such as DeClone, and to explore the parameter space. In terms of direct applications of the method developed here and

in [6], gene-tree based reconstruction of ancestral gene orders comes to mind [5]; more precisely, ancestral adjacencies could be determined and scored using a mixture of their Boltzmann probability (that can be computed efficiently using `DeClone`) and robustness to changes of the cost scheme, and conflicts could be cleared out independently and efficiently for each ancestral species using the algorithm of [12] for example.

An interesting observation is that even the set of ancestral adjacencies that are universally-parsimonious and robust to changes in the cost scheme contains a significant number of adjacencies participating in syntenic conflict. We conjecture that the main reason for syntenic conflicts is in the presence of a significant number of erroneous reconciled gene trees. This is supported by the observation that the ancestral species with the highest number of syntenic conflict are also species for which the reconciliation with the mammalian species tree resulted in a significantly larger number of genes than expected (data not shown). This points clearly to errors in either gene tree reconstruction or in the reconciliation with the mammalian species phylogeny, which tends to assign wrong gene duplications in some specific species, resulting an inflation of the number of genes, especially toward the more ancient species [9]. It would be interesting to see if the information about highly supported conflicting adjacencies can be used in reconciled gene tree correction.

From a methodological point of view, we considered here extended signatures for a single ancestral adjacency at a time. It would be natural to extend this concept to the more general case of several ancestral adjacencies considered at once. We conjecture that this case can be addressed without an increase in the asymptotic complexity of computing the polytope; this problem will be considered in the full version of the present work. Next, there exists another way to explore the parameter space of a dynamic programming phylogenetic algorithm. It consists of computing the *Pareto-front* of the input instance [10, 16], rather than optimal signatures for classes of cost schemes. A signature v is said to be *Pareto-optimal* if there is no other signature whose entries are equal or smaller than the corresponding entries in v , and is strictly smaller at at least one coordinate. The Pareto-front is the set of all Pareto-optimal signatures, and can be efficiently computed by dynamic programming [17, 16, 10]. The Pareto-front differs from the approach we describe in the present work in several aspects. An advantage of the Pareto-front is that it is a notion irrespective of the type of cost function being used. This contrasts with the polytope propagation technique, which requires that the cost function be a linear combination of its terms. However, so far, the Pareto-approach has only been used to define a partition of the parameter space when the cost function is restricted to be linear/affine, and it remains to investigate the difference with the polytope approach in this case.

Acknowledgement

The authors wish to express their gratitude towards Anatoliy Tomilov (Ural Federal University) for contributing the open source C++ 11 implementation of the `QuickHull` algorithm used for the convex hull computations within `DeClone`. This research was

supported by an NSERC Discovery Grant to C.C., an SFU Michael Stevenson scholarship to A.R.

References

1. M. S. Bansal, E. J. Alm, and M. Kellis. Reconciliation revisited: Handling multiple optima when reconciling with duplication, transfer, and loss. *J. Comput. Biol.*, 20(10):738–754, 2013.
2. S. Bérard, C. Gallien, B. Boussau, G. J. Szöllösi, V. Daubin, and E. Tannier. Evolution of gene neighborhoods within reconciled phylogenies. *Bioinformatics*, 28(18):382–388, 2012.
3. P. Biller, P. Feijão, and J. Meidanis. Rearrangement-based phylogeny using the single-cut-or-join operation. *IEEE/ACM Trans. Comput. Biol. Bioinf.*, 10(1):122–134, 2013.
4. T. M. Chan. Optimal output-sensitive convex hull algorithms in two and three dimensions. *Discrete Comput. Geom.*, 16:361–368, 1996.
5. C. Chauve, N. El-Mabrouk, L. Gueguen, M. Semeria, and E. Tannier. Duplication, rearrangement and reconciliation: A follow-up 13 years later. In *Models and Algorithms for Genome Evolution*, pages 47–62. Springer, 2013.
6. C. Chauve, Y. Ponty, and J. P. P. Zanetti. Evolution of genes neighborhood within reconciled phylogenies: An ensemble approach. In *BSB*, volume 8826 of *Lecture Notes Comput. Sci.*, pages 49–56. Springer, 2014.
7. Y. Gagnon, M. Blanchette, and N. El-Mabrouk. A flexible ancestral genome reconstruction method based on gapped adjacencies. *BMC Bioinformatics*, 13(S-19):S4, 2012.
8. D. Gusfield, K. Balasubramanian, and D. Naor. Parametric optimization of sequence alignment. *Algorithmica*, 12(4/5):312–326, 1994.
9. M. W. Hahn. Bias in phylogenetic tree reconciliation methods: implications for vertebrate genome evolution. *Genome Biol.*, 8:R141, 2007.
10. R. Libeskind-Hadas, Y. Wu, M. S. Bansal, and M. Kellis. Pareto-optimal phylogenetic tree reconciliation. *Bioinformatics*, 30(12):87–95, 2014.
11. J. Ma, A. Ratan, B. J. Raney, B. B. Suh, L. Zhang, W. Miller, and D. Haussler. DUPCAR: reconstructing contiguous ancestral regions with duplications. *J. Comput. Biol.*, 15(8):1007–1027, 2008.
12. J. Manuch, M. Patterson, R. Wittler, C. Chauve, and E. Tannier. Linearization of ancestral multichromosomal genomes. *BMC Bioinformatics*, 13(S-19):S11, 2012.
13. L. Pachter and B. Sturmfels. Parametric inference for biological sequence analysis. *Proc. Natl. Acad. Sci. U.S.A.*, 101(46):16138–16143, 2004.
14. L. Pachter and B. Sturmfels. Tropical geometry of statistical models. *Proc. Natl. Acad. Sci. U.S.A.*, 101(46):16132–16137, 2004.
15. L. Pachter and B. Sturmfels, editors. *Algebraic Statistics for Computational Biology*. Cambridge University Press, 2005.
16. C. Saule and R. Giegerich. Observations on the feasibility of exact Pareto optimization. In *Proceedings of CMSR 2014*, pages 43–56, 2014.
17. T. Schnattinger, U. Schöning, and H. A. Kestler. Structural RNA alignment by multi-objective optimization. *Bioinformatics*, 29(13):1607–1613, 2013.
18. E. Tannier, C. Zheng, and D. Sankoff. Multichromosomal median and halving problems under different genomic distances. *BMC Bioinformatics*, 10, 2009.