A Canonical Tree Decomposition for Chirotopes

Mathilde Bouvel ✉
Université de Lorraine, CNRS, INRIA, LORIA, F-54000 Nancy, France

Valentin Feray ✉
Université de Lorraine, CNRS, IECL, F-54000 Nancy, France

Xavier Goaoc ✉
Université de Lorraine, CNRS, INRIA, LORIA, F-54000 Nancy, France

Florent Koechlin ✉
Université Sorbonne Paris Nord, LIPN, CNRS UMR 7030, F-93340 Villetaneuse, France

Abstract

We introduce and study a notion of decomposition of planar point sets (or rather of their chirotopes) as trees decorated by smaller chirotopes. This decomposition is based on the concept of mutually avoiding sets, and adapts in some sense the modular decomposition of graphs in the world of chirotopes. The associated tree always exists and is unique up to some appropriate constraints. We also show how to compute the number of triangulations of a chirotope efficiently, starting from its tree and the (weighted) numbers of triangulations of its parts.

2012 ACM Subject Classification Mathematics of computing → Combinatoric problems; Theory of computation → Computational geometry

Keywords and phrases Order type, modular decomposition, counting triangulations, mutually avoiding point sets, generating functions, rewriting systems

Digital Object Identifier 10.4230/LIPIcs.SoCG.2024.23

Related Version *Full Version*: `https://arxiv.org/abs/2403.10311` [12]

Supplementary Material *Software*: `https://github.com/koechlin/chirotope_tree`
archived at `swh:1:dir:b7043f3f7acae1f7c7c21e32a6ca0cc08a7f1cb4`

Funding *Florent Koechlin*: This work was done while FK was a postdoctoral fellow in LORIA, funded by ANR ASPAG (ANR-17-CE40-0017), IUF and INRIA.

Acknowledgements The authors thank Emo Welzl for discussion leading to footnote 2 and the anonymous referees for helpful comments.

1 Introduction

Two planar point sets are *mutually avoiding* if no line through two points from one set separates two points from the other set (Figure 1 gives an example). Any set of n points in general position contains two mutually avoiding subsets of size $\Omega(\sqrt{n})$ each [6], a bound that is asymptotically best possible [26], and smaller mutually avoiding subsets are actually abundant [25, Theorem 1.3]. Mutually avoiding subsets have been applied to the study of crossing families [6] and empty k-gons [26]. In this paper, we investigate their use for the *recursive decomposition* of planar point sets. We also illustrate how such decompositions facilitate the analysis of point sets on the task of *counting triangulations*.

Our approach is combinatorial and is inspired from the *modular decomposition method*, which is an important tool in algorithmic graph theory (see *e.g.* [17]). Intuitively, we encode a point set by a tree whose nodes correspond to subsets, each point appearing in a single

40th International Symposium on Computational Geometry (SoCG 2024).
Editors: Wolfgang Mulzer and Jeff M. Phillips; Article No. 23; pp. 23:1–23:18
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

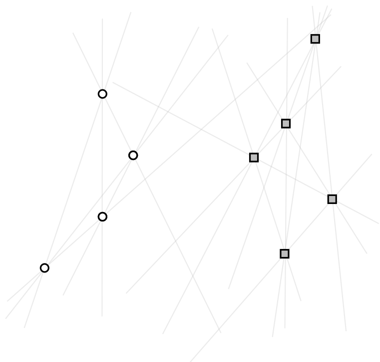

Figure 1 A set of 9 points with two mutually avoiding subsets of 4 and 5 points, respectively.

node, so that subtrees joined by an edge encode mutually avoiding subsets. Formally, the tree rather encodes the *chirotope* of the point set, that is the function mapping each ordered triple of points to its orientation (see below). Here are the main contributions presented here.

- We introduce *chirotope trees* and show how they can be used to build and analyze (chirotopes of) large point sets out of smaller ones.
- We show how the number of triangulations of a chirotope given by a chirotope tree can be computed from the (weighted) numbers of triangulations of its nodes.
- We identify a uniquely defined *canonical tree decomposition* of a chirotope that describes how much it decomposes into mutually avoiding parts.
- We show that the proportion of realizable chirotopes that are indecomposable tends to 1.

The first three results hold both for abstract and realizable chirotopes. We conclude the paper with several new open questions raised by canonical tree decompositions.

Throughout the paper, all point sets are finite, planar[1] and in general position, meaning that **no three points are ever aligned** (for simplicity). Given a finite set X, we write $(X)_3$ for the set of triples (x, y, z) of distinct elements in X. We write $[n]$ for the set $\{1, 2, \ldots, n\}$.

1.1 Context and motivation

Let us first provide some context on the objects, methods and questions that we consider.

Realizable chirotopes. The *chirotope of a set* $\mathcal{P} = \{\mathfrak{p}_\ell\}_{\ell \in X}$ of points in general position labeled by X is the function

$$\chi_{\mathcal{P}} : \begin{cases} \quad (X)_3 & \to \quad \{-1, +1\} \\ (x, y, z) & \mapsto \quad \begin{cases} +1 & \text{if } \mathfrak{p}_x, \mathfrak{p}_y, \mathfrak{p}_z \text{ are in counterclockwise order,} \\ -1 & \text{if } \mathfrak{p}_x, \mathfrak{p}_y, \mathfrak{p}_z \text{ are in clockwise order.} \end{cases} \end{cases}$$

This function encodes the *labeled order type* [15] of the point set. We say that chirotopes of point sets are *realizable* to distinguish them from their combinatorial (abstract) generalization.

[1] It would be non-trivial to adapt our construction to higher dimension.

Representing realizable chirotopes. Deciding if a given function $(X)_3 \to \{-1,1\}$ is a realizable chirotope turns out to be a challenging problem: it is equivalent to the existential theory of the reals [11, Theorem 8.7.2] and NP-hard [24]. Whether this problem is in NP is open, and, interestingly, a positive answer is equivalent to the existence of a (pseudorandom) generator that produces a realizable chirotope of size n in time polynomial in n while ensuring that every element has nonzero probability.[2] The space of realizable chirotopes is therefore difficult to explore, which makes it hard to test efficiently some conjectures in discrete geometry[3] or to check experimentally a geometric algorithm in the exact geometric computing paradigm [23]. The present work grew out of an attempt to devise new ways of constructing and representing (large) realizable chirotopes.

Abstract chirotopes. Most of our results hold in a more general setting. A *chirotope* on a finite set X is a function $\chi : (X)_3 \to \{-1,1\}$ that satisfies the following properties:

(symmetry) for any distinct $x, y, z \in X$,

$$\chi(x,y,z) = \chi(y,z,x) = \chi(z,x,y) = -\chi(z,y,x) = -\chi(y,x,z) = -\chi(x,z,y); \tag{1}$$

(interiority) for any distinct $t, x, y, z \in X$,

$$\chi(t,y,z) = \chi(x,t,z) = \chi(x,y,t) = 1 \quad \Rightarrow \chi(x,y,z) = 1; \tag{2}$$

(transitivity) for any distinct $s, t, x, y, z \in X$,

$$\chi(t,s,x) = \chi(t,s,y) = \chi(t,s,z) = \chi(x,y,t) = \chi(y,z,t) = 1 \quad \Rightarrow \quad \chi(x,z,t) = 1. \tag{3}$$

These functions are in correspondence with the relabeling classes of acyclic uniform oriented matroids of rank 3 [11]. We note that any realizable chirotope is a chirotope. We may use the term *abstract chirotope* to mean a chirotope that is not necessarily realizable.

Notions and properties for point sets that can be expressed by orientations generalize to abstract chirotopes. For example, an element $x \in X$ is *extreme* in a chirotope χ on X if there exists $y \in X \setminus \{x\}$ such that $\chi(x,y,z)$ is the same for all $z \in X \setminus \{x,y\}$. With this definition, Carathéodory's theorem generalizes to chirotopes [11, Theorem 9.2.1 (1)].

▶ **Lemma 1.** *An element $x \in X$ is not extreme in a chirotope χ on X if and only if there exist three distinct elements a, b, c in $X \setminus \{x\}$ such that $\chi(x,a,b) = \chi(x,b,c) = \chi(x,c,a) = 1$.*

Triangulations. Counting the triangulations supported by a given set of n points is a classical problem in computational geometry (see the discussion in [18]). The fastest known algorithm is due to Marx and Miltzow [18] and has complexity $\mathcal{O}(n^{(11+o(1))\sqrt{n}})$. It is quite involved, and a simpler solution, due to Alvarez and Seidel [5], runs in time $\mathcal{O}(n^2 2^n)$.

Given a point set $P = \{\mathfrak{p}_\ell\}_{\ell \in X}$, two segments $\mathfrak{p}_x\mathfrak{p}_y$ and $\mathfrak{p}_z\mathfrak{p}_t$ with distinct endpoints in P cross if and only if $\chi_P(x,y,z) = -\chi_P(x,y,t)$ and $\chi_P(z,t,x) = -\chi_P(z,t,y)$. We can therefore define the crossing of segments for chirotopes. A *segment* in a chirotope χ on X is a pair of elements of X, and the segments xy and zt *cross in* χ if they satisfy the above condition. A *triangulation* of χ is an inclusion-maximal family of segments such that no two cross in χ. The algorithm of Alvarez and Seidel easily generalizes to abstract chirotopes.

[2] A random generator can serve as a verifier for the problem, with the random bitstring as the certificate. Conversely, we can run a random function and a random certificate through the verifier; if accepted, we return that function, otherwise we return a fixed chirotope, *e.g.* n points in convex position.

[3] For instance whether there exists a set of 30 points with no empty hexagon [21].

Modular decomposition. The gist of modular decomposition is to break down discrete structures to allow efficient recursion. These decompositions usually start by partitioning the elements of the structure so that any two elements in a part (also called a *module*) are indistinguishable "from the outside". Choosing this partition as coarse as possible while being nontrivial, and iterating the decomposition within each part yields a decomposition tree. These ideas originated in graph theory [14], where modules gather vertices with the same neighbors outside of the module, and several hereditary classes of graphs have well-behaved modular decompositions, *e.g.* comparability graphs, permutation graphs and cographs [17]. Other examples of structures for which modular decompositions were developed include boolean functions, set systems and permutations [19, 3]. Each structure requires an ad hoc analysis that the proposed notion of modules leads to a well-defined decomposition, where each object has a unique decomposition tree. The proportion of objects with nontrivial decomposition is often vanishingly small (with exceptions, *e.g.* permutations [4]). Nevertheless, such decompositions proved useful *e.g.* in devising fixed-parameter polynomial algorithms for hard algorithmic problems [17, §7] or in solving counting problems [13, 10].

1.2 Our results

We now define our decomposition and state our main results. Sections 2 to 6 give the key ideas behind these results, the details being deferred to the full version [12].

Bowtie decomposition. The fact that two point sets P and Q are mutually avoiding has two interesting consequences at the level of chirotopes. First, the chirotope $\chi_{P\cup\{p^*\}}$ is, up to relabeling, independent of the choice of the point $p^* \in Q$. Second, the chirotope $\chi_{P\cup Q}$ is completely determined by the chirotopes $\chi_{P\cup\{p^*\}}$ and $\chi_{Q\cup\{q^*\}}$ for any choices of $p^* \in Q$ and $q^* \in P$. We can use this to decompose $\chi_{P\cup Q}$ in terms of $\chi_{P\cup\{p^*\}}$ and $\chi_{Q\cup\{q^*\}}$.

Let us express this decomposition. A *sign function* on a set X is a function $(X)_3 \to \{-1,+1\}$ that satisfies the symmetry condition (1). Let X and Y be disjoint sets, and let $x^* \notin X$ and $y^* \notin Y$. Given two sign functions χ on $X \cup \{x^*\}$ and ξ on $Y \cup \{y^*\}$, we define the *bowtie* $\kappa \stackrel{\text{def}}{=} \chi \,{}_{x^*}\!\bowtie_{y^*} \xi$ as the sign function on $X \cup Y$ satisfying:

$$\begin{cases} \kappa(x_1,x_2,x_3) = \chi(x_1,x_2,x_3) & \text{if } x_1,x_2,x_3 \text{ are all in } X; \\ \kappa(x_1,x_2,y) = \chi(x_1,x_2,x^*) & \text{if } x_1,x_2 \text{ are in } X \text{ and } y \text{ is in } Y; \\ \kappa(x,y_2,y_3) = \xi(y^*,y_2,y_3) & \text{if } x \text{ is in } X \text{ and } y_2,y_3 \text{ are in } Y; \\ \kappa(y_1,y_2,y_3) = \xi(y_1,y_2,y_3) & \text{if } y_1,y_2,y_3 \text{ are all in } Y. \end{cases} \tag{4}$$

This defines κ on $(X\cup Y)_3$ via the symmetry condition (1). Note that for mutually avoiding point sets P and Q, $\chi_{P\cup Q} = \chi_{P\cup\{p^*\}} \,{}_{p^*}\!\bowtie_{q^*} \chi_{Q\cup\{q^*\}}$ for any choice of $p^* \in Q$ and $q^* \in P$.

Bowtie products. While defined on sign functions, the bowtie operator also allows to combine smaller chirotopes into larger ones. This works under the following conditions.

▶ **Proposition 2.** *Let X and Y be disjoint sets, with $|X|,|Y| \geq 2$, and let $x^* \notin X$ and $y^* \notin Y$. Let χ and ξ be chirotopes on $X \cup \{x^*\}$ and $Y \cup \{y^*\}$.*

(i) *$\chi \,{}_{x^*}\!\bowtie_{y^*} \xi$ is a chirotope if and only if x^* and y^* are extreme in χ and ξ.*

(ii) *$\chi \,{}_{x^*}\!\bowtie_{y^*} \xi$ is a realizable chirotope if and only if χ and ξ are realizable and x^* and y^* are extreme in χ and ξ.*

(iii) *If (i) holds, then the extreme elements of $\chi \,{}_{x^*}\!\bowtie_{y^*} \xi$ are the elements of $X \cup Y$ extreme in χ or in ξ.*

We also show (Lemma 7) that $\bowtie$ has associativity and commutativity properties. We say that a chirotope κ is *decomposable* if there exist chirotopes χ and ξ, each on a strictly smaller set, such that $\kappa = \chi\,{}_{x^*}\!\!\bowtie_{y^*} \xi$. If no such χ and ξ exist we say that κ is *indecomposable*.

Chirotope trees. Formally, a *chirotope tree* is a tree whose nodes are decorated with chirotopes on disjoint ground sets, and whose edges select an element in the ground set of each of their extremities so that (i) no element is selected more than once, and (ii) each selected element is extreme in its chirotope. Any element selected by an edge of the chirotope tree is called a *proxy*; for illustration purposes, we write the decorating chirotope as the label of the node (if it is realizable, we may draw a realization) and we draw edges of the tree and proxy elements in red, with the edge connecting the proxy elements. For example:

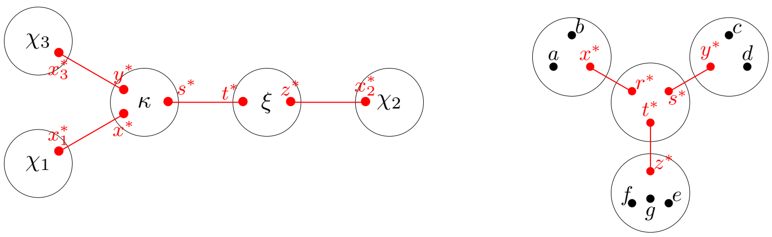

Chirotope of a chirotope tree. To any chirotope tree $T = (V_T, E_T)$ we associate a sign function χ_T as follows. For a node v of T, let χ_v denote the chirotope decorating it, and let X_v denote the set of non-proxy elements of the ground set of χ_v. For $x \in \cup_{w \in V_T} X_w$, the *representative* $R_T(x, v)$ of x in the node v is x if $x \in X_v$, and otherwise it is the label of the proxy selected in χ_v by the first edge on the path in T from v to the node containing x. For distinct $x, y, z \in \cup_{w \in V_T} X_w$, we let $v(x, y, z)$ be the intersection node of the paths in T from x to y,[4] from y to z and from x to z (if two or three elements among x, y, z are in the same X_w, then we set $v(x, y, z) = w$). We define a sign function χ_T on $\cup_{w \in V_T} X_w$ by[5]

$$\chi_T(x, y, z) \stackrel{\text{def}}{=} \chi_v(R(x, v), R(y, v), R(z, v)) \qquad \text{for } v \stackrel{\text{def}}{=} v(x, y, z). \tag{5}$$

For example, if T is the chirotope tree on the right in the previous picture we have $\chi_T(a, c, f) = \kappa(r^*, s^*, t^*) = -1$, where κ is the chirotope decorating the central node. The definition of χ_T ensures the following property: if removing an edge in a tree T produces two subtrees T_1 and T_2, then χ_T is the bowtie product of χ_{T_1} and χ_{T_2} (Lemma 8).

▶ **Proposition 3.** *For any chirotope tree T, χ_T is a chirotope. Moreover, if all chirotopes decorating a node in T are realizable, then χ_T is realizable.*

By Proposition 2(iii), the extreme points of χ_T are the non-proxy extreme points of its decorating chirotopes. Hence, if T has k nodes, then χ_T has at least $k + 2$ extreme points.

Canonical chirotope trees. Let us consider a chirotope tree T and a node v of T. If $\chi_v = \kappa\,{}_{s^*}\!\!\bowtie_{t^*} \xi$ is decomposable, then v can be replaced by two nodes decorated with κ and ξ, connected by an edge, so that the resulting chirotope tree T' satisfies $\chi_{T'} = \chi_T$ (see Section 3 for details). Hence, a chirotope χ corresponds to many chirotope trees, at least

[4] Rather: from the node containing x to the node containing y, and similarly for (y, z) and (x, z).
[5] $R(x, v)$ replaces, *i.e.* serves as a proxy for, x for computing the sign $\chi_T(x, y, z)$ in the node v.

one of which is decorated only by indecomposable chirotopes. However, even requesting the chirotopes decorating the nodes to be indecomposable does not ensure that a unique chirotope tree represents a given chirotope. Here are two trees with the same associated chirotope on $\{a, b, c, d, e, f, g\}$.

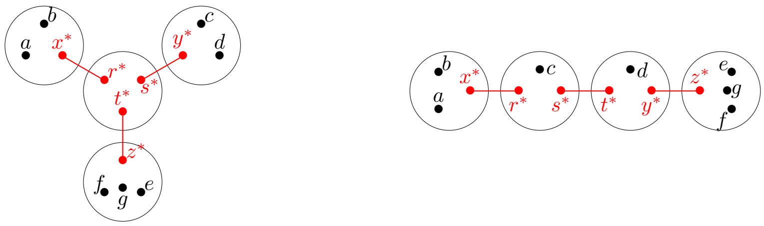

It turns out that the chirotopes that are *convex* in the sense that all their elements are extreme are the only source of redundancy (the chirotope in the above example is not convex, but its restriction to $\{a, b, c, d, e\}$ is). This leads us to define a chirotope tree as *canonical* if every node is decorated by a convex or indecomposable chirotope, and if no edge connects two nodes decorated by convex chirotopes.

▶ **Theorem 4.** *For any chirotope χ, there is a unique canonical chirotope tree T with $\chi_T = \chi$.*

Counting triangulations. Let $\mathcal{T}_\kappa$ denote the set of triangulations of a chirotope κ on X. Given $x^* \in X$, we let $P_{\kappa,x^*}(s) = \sum_{T\in\mathcal{T}_\kappa} s^{\deg_T(x^*)}$ denote the generating polynomial of the triangulations of κ, marking the degree of x^*. We prove that if $\kappa = \chi\, {}_{x^*}\!\bowtie_{y^*} \xi$, then

$$|\mathcal{T}_\kappa| = \sum_{\substack{T'\in\mathcal{T}_\chi \\ T''\in\mathcal{T}_\xi}} \binom{\deg_{T'}(x^*) + \deg_{T''}(y^*) - 2}{\deg_{T'}(x^*) - 1} = \sum_{a,b\geq 2} \tbinom{a+b-2}{a-1}[s^a]P_{\chi,x^*}(s)\ [t^b]P_{\xi,y^*}(t). \qquad (6)$$

More generally, given a chirotope tree T, we can compute the number of triangulations of χ_T given, for each node v, the generating polynomial P_v of the triangulations of its decorating chirotope χ_v that marks not only the degrees of each proxy, but also the presence of each pair of proxies as an edge in the triangulation. If χ_v has n_v points, k of which are proxies, then P_v has $k + \binom{k}{2}$ variables and can be computed from χ_v in time $O\left(2^{n_v+4k} n_v^{k+2}\right)$ by a simple modification of the Alvarez-Seidel algorithm [5].

▶ **Proposition 5.** *Let T be a chirotope tree with m edges, in which each node has degree at most k, and such that χ_T has size n. The number of triangulations of χ_T can be computed from the polynomials $\{P_v\colon v$ a node of $T\}$ in time $O\left(2^{4k} n^{2k} m\right)$ in the Real-RAM model.*

As illustrations, we analyze a family of chirotope chains, obtaining a closed formula for its number of triangulations (Proposition 12) via the *kernel method* from analytic combinatorics, and implement our method for trees of arity 3 (see [12, Section 9] for the details and access to the code). This implementation is merely a proof of concept, and it is beyond the scope of this paper to optimize it or benchmark it, but we can mention that on a laptop, it took a few seconds to count the triangulations of the example of Figure 3 page 15, that assembles a set of 254 points from decorating chirotopes of size 9, and that it handles random ternary chirotope trees with 146 nodes (1024 points) in less than 5 minutes.

Proportion of decomposable chirotopes among realizable ones. Finally, let t_n denote the number of *realizable* chirotopes of size n and d_n denote the number of those that are decomposable.

▶ **Theorem 6.** *For large n, we have* $3\left(n - \mathcal{O}(n^{-2})\right) t_{n-1} \leq d_n \leq \mathcal{O}(n^{-3}) t_n$.

Since t_n is of order n^{4n}, we expect t_{n-1}/t_n to behave as $\Theta(n^{-4})$ (but only an upper bound of this order is known). If this is indeed the case, the lower bound in our theorem behaves as $\Theta(n^{-3})t_n$, i.e. our bounds are tight up to multiplicative constants. Unfortunately, we do not have an analogue result for abstract chirotopes.

1.3 Discussion and related work

Originality with respect to other modular decomposition theories. The theory of modular decomposition [19] does not offer any unifying model that specializes to various structures, but rather proposes general guidelines for developing such a theory. The implementation of these guidelines usually requires analyses that are specific to the objects at hand, and the case of chirotopes is no exception: almost every step requires some specific geometric idea(s).

Among realizable chirotopes, the proportion of indecomposable ones tends to 1. This is analogue to the well-known fact that most graphs are *prime graphs* for the modular decomposition. The proof is, however, much more difficult and interesting than in the graph case because (i) the number of realizable chirotopes grows slower than the number of graphs on n vertices and (ii) this number is only known up to exponential corrections. In particular, our proof uses some geometric constructions and is specific to the realizable setting.

Let us also mention that even if most graphs are prime graphs for the modular decomposition, this decomposition has turned out useful in many ways: design of efficient algorithms [17], enumeration and study of specific classes [9], ... Similarly, for chirotopes, our tree decomposition allows one for example to construct families of large chirotopes for which enumerating triangulations is algorithmically easy (see Section 4).

Recursive constructions of point sets. Recursive constructions of point sets abund in discrete geometry, and with some care several of them directly translate into recursively defined canonical chirotope trees (*e.g.* Horton sets). The closest predecessor to (and inspiration for) this work is the recursive decomposition of *chains* used by Rutschmann and Wettstein [20, Theorem 15] to produce a new lower bound on the number of triangulations of a n-points set. Their representation applies to a smaller subset of order types (the *chains*, of which there are only exponentially many of size n), and decomposes every chain into parts of constant size. It allows to count the triangulations of a chain of size n in $\mathcal{O}(n^2)$ time based on ideas similar to the proof of Equation (6), but tailored to the setting of chains. See the full version [12] for a comparison to other related works [7, 1].

Counting triangulations. An extensive comparison between our method for counting triangulations and the experimental results of Alvarez and Seidel [5] and Rutschmann and Wettstein [20] is beyond the scope of this paper, all the more that the methods apply to different classes of point sets and operate on different types of input. We nevertheless note the following points.

- The general method of Alvarez and Seidel takes as input an arbitrary point set. They tested three methods on high-memory hardware for various types of point sets, and none could handle any example of size 35 or more in less than 10 minutes.
- Our method takes as input a (chirotope presented by a) decomposition tree. We tested it on randomly-generated trees with decorating chirotopes of size 9 summing up to sets of ~ 1000 points; it took less than 5 minutes on a laptop to count the triangulations.
- The method of Rutschmann and Wettstein is based on formulas specific to the class of chains and could handle examples of size 2^{21}.

To us, this is an evidence that our approach is relevant for point sets that are highly decomposable. This is similar to some of the usual benefits of modular decomposition: some hard problems enjoy simple and effective solutions for instances that decompose well.

2 Chirotope trees

Here we implement the modular decomposition guidelines for (realizable) chirotopes.

When is a bowtie product a (realizable) chirotope? Let us first explain why a bowtie $\chi\,{}_{x^*}\!\bowtie_{y^*}\xi$ is a (realizable) chirotope exactly when χ and ξ are (realizable) chirotopes and x^* and y^* are extreme elements. See [12, Section 3] for the detailed proof.

Sketch of proof of Proposition 2 (i) and (ii). Suppose that two functions χ and ξ are such that $\chi\,{}_{x^*}\!\bowtie_{y^*}\xi$ is a chirotope. Since the properties of symmetry, interiority and transitivity that characterize a chirotope are preserved by restriction, χ and ξ must themselves be chirotopes. Furthermore, by Lemma 1, if x^* is not extreme in χ, then we can find three elements a, b, c such that $\chi(x^*, a, b) = \chi(x^*, b, c) = \chi(x^*, c, a) = 1$. A simple case analysis shows that these three elements and any two from Y must violate the transitivity axiom.

Going in the other direction, suppose that χ and ξ are chirotopes and x^* and y^* are extreme. Simple case analysis reveals that $\chi\,{}_{x^*}\!\bowtie_{y^*}\xi$ satisfies the interiority axiom and that if five elements violate the transitivity axiom, then three must come from X and two from Y (or vice-versa) and the three from X must surround x^* (in the sense of Lemma 1). This contradicts the assumption that x^* is extreme and we conclude that $\chi\,{}_{x^*}\!\bowtie_{y^*}\xi$ is indeed a chirotope. Altogether, this proves Statement (i) for abstract chirotopes.

Now to Statement (ii). If $\chi\,{}_{x^*}\!\bowtie_{y^*}\xi$ is a realizable chirotope then (i) implies that x^* and y^* are extreme in χ and ξ, respectively. Moreover, any point set realizing $\chi\,{}_{x^*}\!\bowtie_{y^*}\xi$ contains a subset realizing χ and another one realizing ξ. Hence χ and ξ must be realizable chirotopes. For the reverse direction, the key idea is that for any realizable chirotope χ and any element x extreme for χ, there exists a realization $\mathcal{P} = \{\mathfrak{p}_x, x \in X\}$ of χ such that $\mathfrak{p}_x$ is in an unbounded cell of the arrangement of the set of lines $\big\{(\mathfrak{p}_y\mathfrak{p}_z) \colon y, z \in X \setminus \{x\}, y \neq z\big\}$. This can be proved starting with any realization and applying a suitable projective transform. ◀

Iterating bowtie products. A nice feature of the bowtie when we want to perfom several such operations is that it is commutative and associative under certain conditions.

▶ **Lemma 7.** *Let χ, ξ, κ be chirotopes on disjoint sets $X \cup \{x^*\}$, $Y \cup \{y_1^*, y_2^*\}$ and $Z \cup \{z^*\}$, with $|X|, |Z| \geq 2, |Y| \geq 1$. If the starred elements are extreme in their respective chirotopes, then $\chi\,{}_{x^*}\!\bowtie_{y_1^*}\xi = \xi\,{}_{y_1^*}\!\bowtie_{x^*}\chi$ and $(\chi\,{}_{x^*}\!\bowtie_{y_1^*}\xi)\,{}_{y_2^*}\!\bowtie_{z^*}\kappa = \chi\,{}_{x^*}\!\bowtie_{y_1^*}(\xi\,{}_{y_2^*}\!\bowtie_{z^*}\kappa)$.*

The proof is elementary (see [12, Section 3]). Using the associativity of bowties requires attention to the ground sets of the chirotopes at play. For instance, with the notation of Lemma 7, writing $(\xi\,{}_{y_1^*}\!\bowtie_{x^*}\chi)\,{}_{y_2^*}\!\bowtie_{z^*}\kappa = \xi\,{}_{y_1^*}\!\bowtie_{x^*}(\chi\,{}_{y_2^*}\!\bowtie_{z^*}\kappa)$ is absurd, as in the right-hand term y_2^* does not belong to the ground set of χ (even though the left-hand term is well-defined).

These commutativity and associativity properties make it possible to omit certain ordering informations in expressions using bowtie products. For instance, Lemma 7 implies that

$$(\chi_1\,{}_{x_1^*}\!\bowtie_{x_2^*}(\chi_2\,{}_{y_2^*}\!\bowtie_{x_3^*}\chi_3))\,{}_{z_2^*}\!\bowtie_{x_4^*}\chi_4 = \chi_1\,{}_{x_1^*}\!\bowtie_{x_2^*}((\chi_2\,{}_{z_2^*}\!\bowtie_{x_4^*}\chi_4)\,{}_{y_2^*}\!\bowtie_{x_3^*}\chi_3)$$

so one is tempted to represent these (and other equivalent) expressions by the following tree.

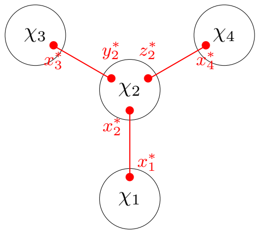

The notion of chirotope tree systematizes this representation of expressions by trees.

Chirotope trees encode chirotopes (proof of Proposition 3). The sign function χ_T of a chirotope tree T is defined so that the following property holds.

▶ **Lemma 8.** *Let T be a chirotope tree and $e = \{v_1, v_2\}$ be an edge of T, and let s_1 and s_2 be the elements selected by e in v_1 and v_2 respectively. If T_1 and T_2 denote the two chirotope trees obtained by removing e from T, with $v_i \in T_i$, then $\chi_T = \chi_{T_1} {}_{s_1}\!\bowtie_{s_2} \chi_{T_2}$.*

This lemma is proved via elementary manipulations of the definitions and case analysis; see [12, Section 4].

The proof of Proposition 3 (stating that χ_T is a chirotope) is now a straightforward induction using Lemma 8 and the fact that the bowtie of two (realizable) chirotopes whose starred elements are extreme is a (realizable) chirotope (Proposition 2). Iterating this lemma, we can actually compute the chirotope χ_T of a chirotope tree starting from the chirotopes decorating each node, and iterating bowtie operations.

3 Canonical chirotope trees

Canonical chirotope trees arise naturally from two operations on chirotope trees.

Contraction. Let T be a chirotope tree and $e = \{v, v'\}$ be an edge in T that selects s^* and t^* in χ_v and $\chi_{v'}$, respectively. The *contraction of e in T* is the tree T' obtained from T by merging the nodes v and v' into a new node v_{new}, decorated with $\chi_{new} \stackrel{\text{def}}{=} \chi_v {}_{s^*}\!\bowtie_{t^*} \chi_{v'}$. We denote this transformation by $T \xrightarrow{e} T'$. By Proposition 2, χ_{new} is a chirotope, and is realizable if both χ_v and $\chi_{v'}$ are realizable. Also, the extreme elements of χ_v and $\chi_{v'}$ different from s^* and t^* are extreme elements of χ_{new} so that the proxy elements corresponding to other edges of T are still extreme in their respective chirotopes in T'; it follows that T' is indeed a chirotope tree. Here is an example of a contraction of a tree (with $\chi = \kappa {}_{s^*}\!\bowtie_{t^*} \xi$).

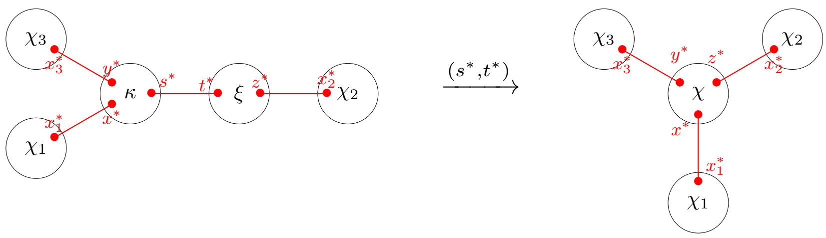

Split. Let T be a chirotope tree with a node v_0 whose decoration χ has a nontrivial bowtie decomposition $\chi = \kappa\, {}_{s^*}\!\bowtie_{t^*} \xi$. The *split of* T *according to* $\chi = \kappa\, {}_{s^*}\!\bowtie_{t^*} \xi$ is the tree T' obtained from T by replacing v_0 by two nodes v_1 and v_2 that are decorated by κ and ξ, respectively, and connected by an edge selecting s^* in v_1 and t^* in v_2. We denote this transformation by $T \xrightarrow[\kappa\, {}_{s^*}\bowtie_{t^*} \xi]{\chi} T'$. Note that for any edge e connecting v_0 to another vertex w in T, the element s selected by e in v_0 belongs to the ground set of either κ or ξ. Therefore, e is naturally seen as an edge of T', connecting either v_1 or v_2 to w, depending on the set in which s lives. Here is an example of a split of a tree according to the bowtie decomposition $\chi = \kappa\, {}_{s^*}\!\bowtie_{t^*} \xi$.

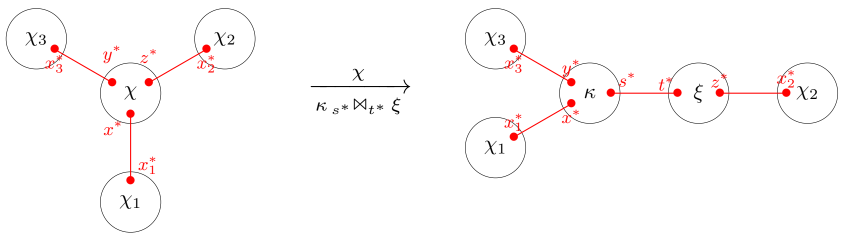

Properties. The contraction and split decomposition operations are inverse of one another in the sense that (with the notation of the previous two paragraphs)

$$T \xrightarrow[\kappa\, {}_{s^*}\bowtie_{t^*} \xi]{\chi} T' \;\Rightarrow\; T' \xrightarrow{(s^*,t^*)} T, \qquad \text{and} \quad T \xrightarrow{(s^*,t^*)} T' \;\Rightarrow\; T' \xrightarrow[\chi_v\, {}_{s^*}\bowtie_{t^*} \chi_{v'}]{\chi_{new}} T. \tag{7}$$

Moreover, contractions and splits do not change the associated chirotope [12, Section 4].

▶ **Proposition 9.** *If* $T \xrightarrow{e} T'$ *or if* $T \xrightarrow[\kappa\, {}_{y^*}\bowtie_{z^*} \xi]{\chi_{v_0}} T'$*, then* $\chi_T = \chi_{T'}$.

Canonical chirotope trees are unique (proof of Theorem 4). As explained in Section 1.2, we define a chirotope tree as *canonical* if every node is decorated by a convex or indecomposable chirotope, and if no edge connects two nodes decorated by convex chirotopes. Consider the following rewriting rules on the set of chirotope trees:

- $T \xrightarrow{\diamond} T'$ if T' is obtained from T by contracting an edge between two *convex* chirotopes;
- $T \xrightarrow{\bowtie} T'$ if T' is obtained from T by splitting a *nonconvex* node of T.

We write $T \Rightarrow T'$ if $T \xrightarrow{\diamond} T'$ or $T \xrightarrow{\bowtie} T'$. As usual, for any rewriting rule $\to$, we write $\to^*$ for the rewriting rule consisting in any number (possibly zero) of successive applications of $\to$.

Let T_χ^{tr} denote the chirotope tree with a single node decorated with χ. By Proposition 9, starting with T_χ^{tr} and applying $\Rightarrow$ any number of times only produces chirotope trees whose associated chirotope is χ. Moreover, by definition, a chirotope tree is canonical if and only if the rewriting rule $\Rightarrow$ does not apply to it. Our proof that every chirotope χ admits a unique canonical chirotope tree (Theorem 4) decomposes as follows (see [12, Section 4]).

- We prove that $\Rightarrow$ *terminates*, *i.e.* every rewriting sequence $T_0 \Rightarrow T_1 \Rightarrow \ldots \Rightarrow T_n \Rightarrow \ldots$ is finite; this implies the existence of at least one canonical chirotope tree associated with χ.
- We prove the *accessibility* from T_χ^{tr} of all canonical chirotope trees T with $\chi_T = \chi$; i.e. we always have $T_\chi^{tr} \Rightarrow^* T$.
- We finally prove that $\Rightarrow$ is *confluent*: starting with T_χ^{tr} and iterating $\Rightarrow$ in any way always leads to the same final state T. Together with the previous point, this implies the uniqueness of a canonical chirotope tree associated with χ.

4 Counting triangulations

We now examine how to count the triangulations of a chirotope described as a chirotope tree.

4.1 The principle

Let $\kappa = \chi \,{}_{x^*}\!\bowtie_{y^*} \xi$ where χ and ξ are chirotopes on $X \cup \{x^*\}$ and $Y \cup \{y^*\}$, with X and Y disjoint, $x^* \notin X$ and $y^* \notin Y$. Let $\pi_{Y\to x}$ denote the map that sends every element of Y to x. We extend this map to triangulations by putting, for any triangulation T of κ

$$\pi_{Y\to x}(T) \stackrel{\text{def}}{=} \{\{\pi_{Y\to x}(a), \pi_{Y\to x}(b)\} \colon \{a,b\} \in T\} \setminus \{\{x,x\}\}.$$

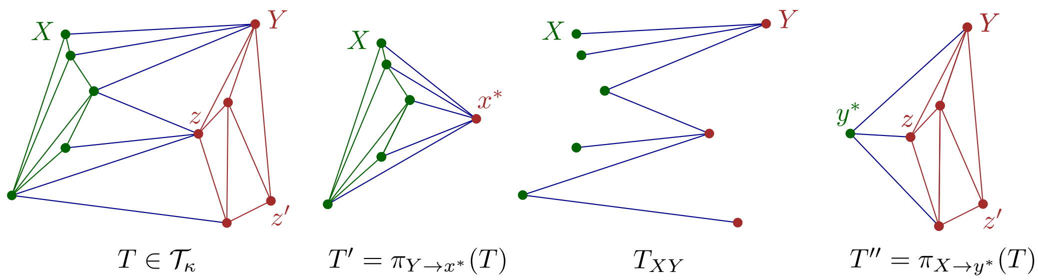

Figure 2 A triangulation T of $\kappa = \chi \,{}_{x^*}\!\bowtie_{y^*} \xi$, the projection of T on $X \cup \{x^*\}$, the set T_{XY}, and the projection of T on $Y \cup \{y^*\}$.

Given a set of edges T of $\chi \,{}_{x^*}\!\bowtie_{y^*} \xi$, not necessarily a triangulation, and two sets A and B, we let T_{AB} denote the set of edges of T with one element in A and the other in B. As illustrated in Figure 2, any triangulation T of κ decomposes into triangulations of χ and ξ and a set T_{XY} of non-crossing edges. This can be turned into a bijection (see [12, Section 7] for a proof).

▶ **Proposition 10.** *With the above notation, for any triangulation T of κ, the sets $T' \stackrel{\text{def}}{=} \pi_{Y\to x^*}(T)$ and $T'' \stackrel{\text{def}}{=} \pi_{X\to y^*}(T)$ are triangulations of χ and ξ, respectively, and T_{XY} is a maximal set of noncrossing edges between the neighbors of x^* in T' and the neighbors of y^* in T''.*

Conversely, for any triangulations T' of χ and T'' of ξ and any maximal set H of non-crossing edges between the neighbors of x^* in T' and the neighbors of y^* in T'', there exists a unique triangulation $T \in \mathcal{T}_\kappa$ such that $T' = \pi_{Y\to x^*}(T)$, $T'' = \pi_{X\to y^*}(T)$ and $H = T_{XY}$.

Recall that for a chirotope α and an element a^*, the polynomial $P_{\alpha,a^*}(s) = \sum_{T\in\mathcal{T}_\alpha} s^{\deg_T(a^*)}$ counts the triangulations of α while marking the degree of a^*. Here are examples with $P_{\chi,x^*}(s) = s^3(s+1)$ and $P_{\xi,y^*}(t) = t^2(1+t+t^2)$.

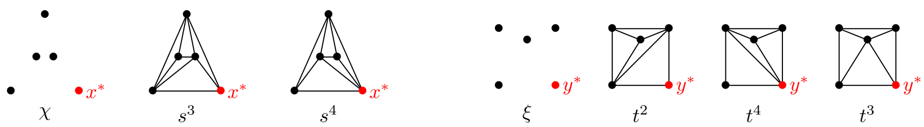

Proposition 10 implies Formula (6) which expresses $|\mathcal{T}_\kappa|$ for $\kappa \stackrel{\text{def}}{=} \chi \,{}_{x^*}\!\bowtie_{y^*} \xi$ from $P_{\chi,x^*}(s)$ and $P_{\xi,y^*}(t)$. To iterate this computation through several bowtie products, we need to determine not only $|\mathcal{T}_\kappa|$, but $P_{\kappa,z^*}(t)$ for some extreme point z^* of κ, to be used as a proxy point in a subsequent bowtie product. For $z^* \in Y$ we replace $P_{\xi,y^*}(t)$ by two bivariate polynomials:

$$Q^{\in}_{\xi,y^*,z^*}(t,u) = \sum_{\substack{T\in\mathcal{T}_\xi\\ \{y^*,z^*\}\in T}} t^{\deg_T(y^*)}u^{\deg_T(z^*)} \text{ and } Q^{\notin}_{\xi,y^*,z^*}(t,u) = \sum_{\substack{T\in\mathcal{T}_\xi\\ \{y^*,z^*\}\notin T}} t^{\deg_T(y^*)}u^{\deg_T(z^*)}.$$

Note that $P_{\xi,y^*}(t) = Q^{\in}_{\xi,y^*,z^*}(t,1) + Q^{\notin}_{\xi,y^*,z^*}(t,1)$. In the example

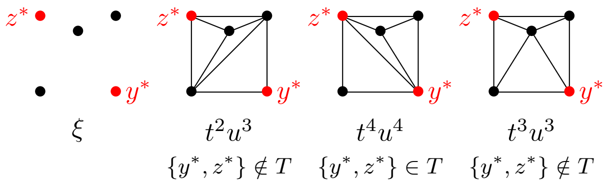

we have $Q^{\in}_{\xi,y^*,z^*}(t,u) = t^4u^4$ and $Q^{\notin}_{\xi,y^*,z^*}(t,u) = t^2u^3(1+t)$. Here is our refinement of Formula (6), see [12, Section 8] for a proof.

▶ **Proposition 11.** *Let $\kappa = \chi\,{}_{x^*}\!\bowtie_{y^*} \xi$ where χ and ξ are chirotopes on $X\cup\{x^*\}$ and $Y\cup\{y^*\}$, with X and Y disjoint, $x^*\notin X$ and $y^*\notin Y$. For any $z^*\in Y$ extreme in ξ we have*

$$P_{\kappa,z^*}(u) = \sum_{a,b\geq 2} \binom{a+b-2}{a-1}[s^a]P_{\chi,x^*}(s)[t^b]Q^{\notin}_{\xi,y^*,z^*}(t,u)$$
$$+\sum_{a,b\geq 2} R_{a,b}(u)[s^a]P_{\chi,x^*}(s)\ [t^b]Q^{\in}_{\xi,y^*,z^*}(t,u),\ \textit{with}\ R_{a,b}(u) = \sum_{i=0}^{a-1}\binom{a+b-i-3}{b-2}u^i.$$

4.2 Example: asymptotic analysis of a family of chirotope chains

Consider the two following chirotopes on $\{x^*,y^*,z,c\}$:

$\chi^{(0)} =$ 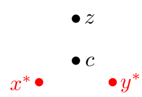 and $\chi^{(1)} =$ 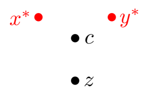

Each of $\chi^{(0)}$ and $\chi^{(1)}$ has a unique triangulation, which contains the edge $\{x^*,y^*\}$ and where each vertex has degree 3. Thus,

$$\begin{array}{rcccl} P_{\chi^{(0)},y^*}(s) &=& P_{\chi^{(1)},y^*}(s) &=& s^3,\\ Q^{\in}_{\chi^{(0)},x^*,y^*}(t,u) &=& Q^{\in}_{\chi^{(1)},x^*,y^*}(t,u) &=& t^3u^3,\\ Q^{\notin}_{\chi^{(0)},x^*,y^*}(t,u) &=& Q^{\notin}_{\chi^{(1)},x^*,y^*}(t,u) &=& 0. \end{array}$$

For $i\in\mathbb{N}$ let $\chi^{(0)}_i$ denote a copy of $\chi^{(0)}$ relabeled by $x^*\mapsto x^*_i$, $y^*\mapsto y^*_i$, $z\mapsto z_i$, $c\mapsto c_i$. We define $\chi^{(1)}_i$ to be a copy of $\chi^{(1)}$ with the same relabeling. For any word $\sigma=\sigma_1\sigma_2\ldots\sigma_k\in\{0,1\}^k$ we put $\chi^{(\sigma)} \stackrel{\text{def}}{=} \chi^{(\sigma_1)}_1\,{}_{y^*_1}\!\bowtie_{x^*_2}\chi^{(\sigma_2)}_2\,{}_{y^*_2}\!\bowtie_{x^*_3}\cdots{}_{y^*_{k-1}}\!\bowtie_{x^*_k}\chi^{(\sigma_k)}_k$. This expression translates into a chirotope tree, actually a chain, for example for $\chi^{(01101)}$:

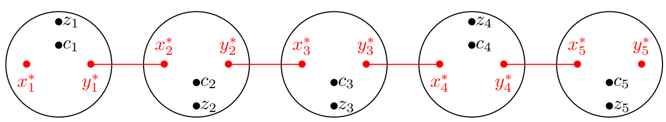

The chirotopes $\chi^{(0)}$ and $\chi^{(1)}$ are indecomposable (and non-convex), so these trees are canonical and Theorem 4 ensures that $\chi^{(\sigma)} = \chi^{(\sigma')}$ if and only if $\sigma = \sigma'$. Moreover, $\chi^{(0)}$ and $\chi^{(1)}$ are realizable so $\chi^{(\sigma)}$ is realizable by Proposition 2 (ii) and is the chirotope of a set of $2k+2$ points, k of which are interior. It turns out that the number of triangulations of $\chi^{(\sigma)}$ depends only on the length of σ, and there is a closed formula for this number!

▶ **Proposition 12.** *For every $\sigma \in \{0,1\}^k$ we have*

$$\left|\mathcal{T}_{\chi^{(\sigma)}}\right| = \frac{3(2k+1)\binom{4k+2}{2k+1}}{(2k+2)(4k+1)} - \frac{4^k\binom{2k+2}{k+1}}{2k+1} \sim_{k\to\infty} \frac{3-2\sqrt{2}}{\sqrt{2\pi}}\,\frac{16^k}{k^{3/2}}.$$

Recall that for $\sigma \in \{0,1\}^k$, the chirotope $\chi^{(\sigma)}$ has $2k+2$ points. In particular, up to a multiplicative constant, $\left|\mathcal{T}_{\chi^{(\sigma)}}\right|$ has the same asymptotics as the number of triangulations of $2k+2$ points in convex position.

Proof. We make a detour via infinite words and we use bivariate generating functions. Let $\tau \in \{0,1\}^{\mathbb{N}}$ be an infinite binary word, let τ_k be its kth letter, and let $\tau^{[k]}$ denote its prefix of length k. We introduce the shorthands

$$P_k(u) \overset{\text{def}}{=} P_{\chi^{(\tau^{[k]})},y_k^*}(u) \quad \text{and} \quad Q_k(t,u) \overset{\text{def}}{=} Q^{\in}_{\chi_k^{(\tau_k)},x_k^*,y_k^*}(t,u) = t^3u^3 \text{ as before.}$$

Since $\chi^{(\tau^{[k+1]})} = \chi^{(\tau^{[k]})}\ {}_{y_k^*}\!\bowtie_{x_{k+1}^*} \chi_{k+1}^{(\tau_{k+1})}$, we can apply Proposition 11 with $z^* = y_{k+1}^*$ to get

$$P_{k+1}(u) = \sum_{a,b\geq 2} R_{a,b}(u)\ \ [s^a]P_k(s)\ \ [t^b]Q_{k+1}(t,u) = u^3\sum_{a=0}^{2k+1} R_{a,3}(u)\ [s^a]P_k(s).$$

Using $R_{a,3}(u) = \sum_{i=0}^{a-1}\binom{a-i}{1}u^i = \frac{u}{(1-u)^2}(u^a-1) + \frac{a}{1-u}$, we can develop

$$P_{k+1}(u) = \frac{u^4}{(1-u)^2}\underbrace{\sum_{a=0}^{2k+1}[s^a]P_k(s)\,u^a}_{=P_k(u)} - \frac{u^4}{(1-u)^2}\underbrace{\sum_{a=0}^{2k+1}[s^a]P_k(s)}_{=P_k(1)} + \frac{u^3}{(1-u)}\underbrace{\sum_{a=0}^{2k+1}a\,[s^a]P_k(s)}_{=P_k'(1)},$$

which rewrites as
$$P_{k+1}(s) - \frac{s^4}{(1-s)^2}P_k(s) = -\frac{s^4}{(1-s)^2}P_k(1) + \frac{s^3}{1-s}P_k'(1)\,. \tag{8}$$

Let us introduce the formal series $F(s,u) \overset{\text{def}}{=} \sum_{k\geq 1} P_k(s)u^k$. Note that $F(1,u) = \sum_{k\geq 1} P_k(1)u^k$ is the generating series of the number of triangulations of $\chi^{(\tau^{[k]})}$. Since $\chi^{(\tau^{[k]})}$ is realizable, we have $P_k(1) \leq 30^{2k+2}$ [22]. Moreover, each P_k has degree $2k+1$ and nonnegative coefficients, so that, for $s > 0$, we have $|P_k(s)| \leq (30s)^{2k+2}$. This implies that F is analytic in the two variables s and u on the domain $D = \{(s,u) : |s|^2 \cdot |u| < \frac{1}{30}\}$, which contains in particular the point $(1,0)$. Now, multiplying Equation (8) by u^{k+1} and summing up for $k \geq 1$, we obtain the functional equation on D

$$\left(1 - \frac{us^4}{(1-s)^2}\right)F(s,u) = us^3\left(1 - \frac{s}{(1-s)^2}F(1,u) + \frac{1}{1-s}\partial_s F(1,u)\right), \tag{9}$$

where $\partial_s F(s,u)$ denotes the derivative of F according to the first variable s.

The right-hand term of Equation (9) is linear in $F(1,u)$ and $\partial_s F(1,u)$. Let us apply the standard *kernel method* (see, e.g., [8]): since $F(s,u)$ is analytic near $(1,0)$, we can evaluate Equation (9) for some $s(u)$ that cancels the coefficient $1-\frac{us^4}{(1-s)^2}$, called the *kernel*, and such that $s(u)$ is close to 1 when u is close to 0. Two functions satisfy these conditions:

$$s_1(u) = \frac{-1+\sqrt{1+4\sqrt{u}}}{2\sqrt{u}} \quad \text{and} \quad s_2(u) = \frac{1-\sqrt{1-4\sqrt{u}}}{2\sqrt{u}}.$$

Substituting into Equation (9) yields the following linear system:

$$\begin{cases} \dfrac{s_1(u)}{(1-s_1(u))^2}F(1,u) - \dfrac{1}{1-s_1(u)}\partial_s F(1,u) = 1 \\ \dfrac{s_2(u)}{(1-s_2(u))^2}F(1,u) - \dfrac{1}{1-s_2(u)}\partial_s F(1,u) = 1 \end{cases}. \tag{10}$$

Solving this linear system gives us the closed form $F(1,u) = s_1(u) + s_2(u) - s_1(u)s_2(u) - 1$. Recovering the number of triangulations of $\chi^{(\tau^{[k]})}$ then amounts to computing a Taylor expansion, and the asymptotic expression follows either by using Stirling formula, or directly from the formula of $F(1,u)$ by singularity analysis. We elaborate on these computations in [12, Section 8]. ◀

4.3 The general method (proof of Proposition 5)

The method generalizes to chirotope trees of arbitrary arity (see [12, Section 9] for details).

Sketch of proof of Proposition 5. For each node v, we define a polynomial P_v that counts the triangulations of the chirotope χ_v decorating v, while marking two types of information: (i) the degree of each proxy of χ_v, and (ii) which edges between proxies of χ_v are contained in the triangulation. If χ_v has k proxies, then P_v has k variables $\{x_1,\dots,x_k\}$ for (i) and $\binom{k}{2}$ variables $\{y_{i,j}\}_{1\le i<j\le k}$ for (ii). Moreover if we write P_v under the form $P_v = \sum_{h\in\mathcal{H}} R_h(x_1,\dots,x_k)\cdot h$ where $\mathcal{H}$ is the set of monomials in the variables $y_{i,j}$ appearing in P_v, then each monomial $h\in\mathcal{H}$ is of degree at most 1 in each $y_{i,j}$, each polynomial $R_h(x_1,\dots,x_k)$ is of degree at most n in each x_i, and $|\mathcal{H}| = \mathcal{O}(16^k)$ (here we use that the proxies are in convex position).

Proposition 11 can be generalized to describe how the polynomial P_v changes when a leaf of a chirotope tree is merged with its parent node. If the parent node has k proxies then this computation can be done by doing $\mathcal{O}(n^2 16^k)$ multiplications of $(k-1)$-variate polynomials of degree at most n, so in $\mathcal{O}(n^{2k}16^k)$ time. Starting from a chirotope tree T and merging the leaves one after another, we reduce the tree to a single node, whose polynomial counts the triangulations of χ_T. In the course of this computation, no polynomial has partial degree more than n with respect to any of their variables, so the bound $\mathcal{O}(n^{2k}16^k)$ is valid for every merging, so that the total time complexity is bounded by $\mathcal{O}(mn^{2k}16^k)$. We believe that this bound is not tight but improving it is beyond the scope of this paper. ◀

As a proof of concept, we implemented this method for trees of arity 2 and 3, both the modification of the Alvarez-Seidel algorithm and the recursive counting (see [12, Section 9]). Our implementation runs within Sagemath to take advantage of basic functions for the manipulation of multivariate polynomials, but is mostly basic python. We made little effort in optimizing it, yet it only takes a few seconds to compute that the number of triangulations of the chirotope of size 254 presented by the tree of Figure 3 is

$$\begin{aligned}&592\,966\,751\,293\,974\,711\,252\,579\,414\,478\,724\,131\,868\,483\,318\,559\,312\,640\,993\,804\,562\,350\,446\\ &\quad 478\,462\,502\,194\,102\,338\,465\,347\,793\,563\,468\,964\,711\,887\,069\,260\,192\,677\,783\,5079\,385\,675\\ &\quad 578\,362\,313\,461\,572\,573\,372\,584\,158\,103\,703\,847\,713\,232\,664 \approx 5.92966751.10^{180}\end{aligned}$$

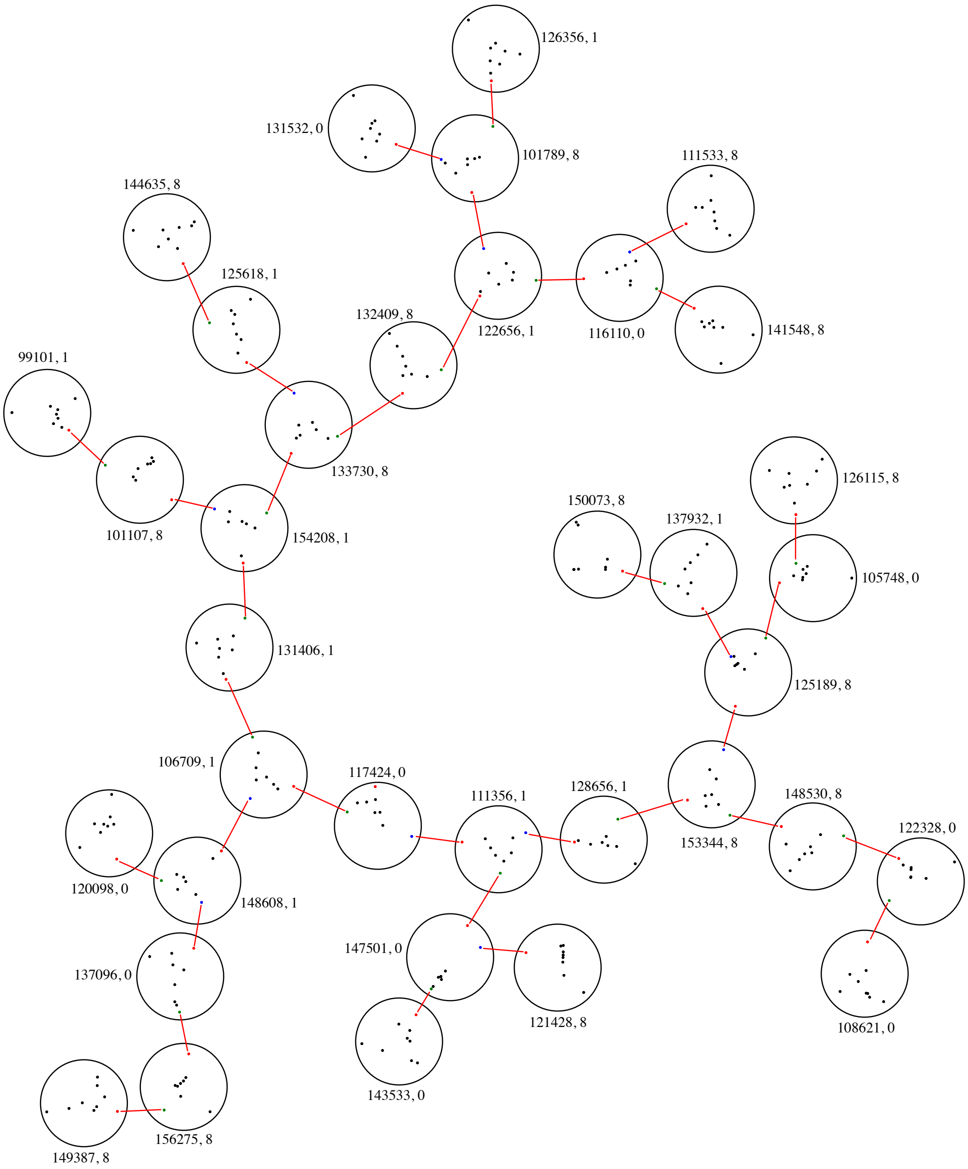

Figure 3 A chirotope tree of 36 nodes, each decorated with a chirotope of size 9, adding up to 254 elements. The numbers next to each node indicate the identifier of the decorating chirotope in the order type database [2] and the label of the red proxy in this database.

5 The number of decomposable chirotopes (Theorem 6)

We now outline the proof of Theorem 6, whose details are presented in [12, Section 5]. Let $\mathcal{X}_n$ be the set of realizable chirotopes on $[n]$, $\mathcal{X}_n^* \subset \mathcal{X}_n$ the subset of those in which n is extreme, $\mathcal{D}_n \subset \mathcal{X}_n$ the subset of those that are decomposable, and $\mathcal{I}_n^* = \mathcal{X}_n^* \setminus \mathcal{D}_n$ the subset of those that are indecomposable and in which n is extreme. We put $t_n = |\mathcal{X}_n|$, $t_n^* = |\mathcal{X}_n^*|$, $d_n = |\mathcal{D}_n|$ and $i_n^* = |\mathcal{I}_n^*|$. We claim that $i_n^* \geq \frac{3}{n}(t_n - d_n)$. Indeed, letting $\operatorname{ext}(\chi)$ denote the number of extreme elements of a chirotope χ, we have by symmetry

$$\sum_{\chi \in \mathcal{X}_n \setminus \mathcal{D}_n} \operatorname{ext}(\chi) = \left|\{(\chi, i) \in (\mathcal{X}_n \setminus \mathcal{D}_n) \times [n] \colon i \text{ is extreme in } \chi\}\right| = n i_n^*.$$

It follows that $\dfrac{i_n^*}{t_n - d_n} = \dfrac{1}{n(t_n-d_n)} \displaystyle\sum_{\chi \in \mathcal{X}_n \setminus \mathcal{D}_n} \mathrm{ext}(\chi) \geq \dfrac{1}{n(t_n-d_n)} 3(t_n - d_n) = \dfrac{3}{n}$.

Now, given $\chi \in \mathcal{D}_n$, let us fix a nontrivial decomposition $\chi = \chi_1'\,{}_{x^*}\!\bowtie_{y^*} \chi_2'$, with χ_1' a chirotope on $M \cup \{x^*\}$ for some subset $M \subset [n]$, and χ_2' a chirotope on $([n] \setminus M) \cup \{y^*\}$. We let $n_1 \stackrel{\text{def}}{=} |M|$ and $n_2 \stackrel{\text{def}}{=} n - n_1$. We then let χ_1 denote the chirotope on $[n_1+1]$ obtained by renaming, in χ_1', the (extreme) element x^* as $n_1 + 1$ and M as $[n_1]$ increasingly. Similarly, we let χ_2 denote the chirotope on $[n_2+1]$ obtained by renaming in χ_2' the set $[n] \setminus M$ as $[n_2]$ increasingly and the (extreme) element y^* as $n_2 + 1$. By construction, $\chi_1 \in \mathcal{X}^*_{n_1+1}$ and $\chi_2 \in \mathcal{X}^*_{n_2+1}$. The map $(\chi, M) \mapsto (\chi_1, \chi_2, M)$ is an injection (its reverse amounts to relabeling χ_1 and χ_2 according to M and taking a bowtie product). As a consequence we have

$$d_n \leq \sum_{\substack{n_1+n_2=n \\ n_1,n_2 \geq 2}} \binom{n}{n_1} t^*_{n_1+1}\, t^*_{n_2+1}. \tag{11}$$

On the other hand, the chirotopes $\chi_1\,{}_{x^*}\!\bowtie_{y^*} \chi_2$ are different for any $\chi_1 \in \mathcal{X}_3$ and $\chi_2 \in \mathcal{X}_{n-1} \setminus \mathcal{D}_{n-1}$ (this can easily be seen *e.g.* via Theorem 4). Since there are 2 chirotopes on $\{1,2,3\}$, it follows that $d_n \geq 2\binom{n}{2} i^*_{n-1}$ and we have

$$3n(t_{n-1} - d_{n-1}) \;\leq\; d_n \;\leq \sum_{\substack{n_1+n_2=n \\ n_1,n_2 \geq 2}} \binom{n}{n_1} t^*_{n_1+1}\, t^*_{n_2+1}. \tag{12}$$

The most difficult part of the proof is to show that the right hand sum is $\mathcal{O}(n^{-3} t_n)$. We fix some $\varepsilon > 0$ independent of n and consider separately the contributions where $\min(n_1, n_2) < \varepsilon n$ and those where $\min(n_1, n_2) \geq \varepsilon n$. The former contribution is easily controlled using two classical facts [16]: $t_n \leq e^{2n}(n-1)^{4n}$ and there exists $K \in (0, 1/8]$ such that, for all $n \geq 3$, we have $t_{n+1} \geq K\, n^4\, t_n$. Controlling the latter contribution requires new geometric ideas and is beyond the scope of this extended abstract (see [12, Section 5]).

6 Some open problems

Let us conclude this presentation by sampling some of directions of enquiry this work opens.

1. Here we used canonical chirotope trees as a tool to put together large chirotopes out of smaller ones, and we therefore considered the chirotope tree as given. One could, however, start with some point set and set out to compute the decomposition tree. This raises the following computational questions. How efficiently can one find a partition of a given (realizable) chirotope into two mutually avoiding parts or decide that none exists? Or, going even further, compute the canonical chirotope tree of a given (realizable) chirotope?

2. Can crossing-free structures besides triangulations be counted efficiently from the chirotope tree? What about other statistics of a chirotope such as the number of k-sets or the number of crossing pairs?

3. What is the proportion of indecomposable (abstract) chirotopes? (Our proof of Theorem 6 uses that $t_n/t_{n-1} \geq \Omega(n^4)$ and we know no analogue of this in the abstract setting.)

References

1 Péter Ágoston, Gábor Damásdi, Balázs Keszegh, and Dömötör Pálvölgyi. Orientation of good covers, 2022. `arXiv:2206.01723`.

2 Oswin Aichholzer, Franz Aurenhammer, and Hannes Krasser. Enumerating order types for small point sets with applications. *Order*, 19(3):265–281, 2002. `doi:10.1023/A:1021231927255`.

3 Michael H. Albert and Mike D. Atkinson. Simple permutations and pattern restricted permutations. *Discrete Mathematics*, 300(1-3):1–15, 2005.

4 Michael H. Albert, Mike D. Atkinson, and Martin Klazar. The enumeration of simple permutations. *Journal of Integer Sequences*, 6, 2003.

5 Victor Alvarez and Raimund Seidel. A simple aggregative algorithm for counting triangulations of planar point sets and related problems. In *Proceedings of the twenty-ninth annual symposium on Computational geometry*, pages 1–8, 2013.

6 Boris Aronov, Paul Erdős, Wayne Goddard, Daniel J. Kleitman, Michael Klugerman, Janos Pach, and Leonard J. Schulman. Crossing families. *Combinatorica*, 14(2):127–134, 1994. `doi:10.1007/BF01215345`.

7 Martin Balko, Jan Kynčl, Stefan Langerman, and Alexander Pilz. Induced ramsey-type results and binary predicates for point sets. *The Electronic Journal in Combinatorics*, 24:1–22, 2017.

8 Cyril Banderier, Philippe Flajolet, Danièle Gardy, Mireille Bousquet-Mélou, Alain Denise, and Dominique Gouyou-Beauchamps. Generating functions for generating trees. *Discrete Mathematics*, 246(1–3):29–55, 2002.

9 Frédérique Bassino, Mathilde Bouvel, Valentin Féray, Lucas Gerin, Mickaël Maazoun, and Adeline Pierrot. Random cographs: Brownian graphon limit and asymptotic degree distribution. *Random Struct. Algorithms*, 60(2):166–200, 2022. `doi:10.1002/rsa.21033`.

10 Frédérique Bassino, Mathilde Bouvel, and Dominique Rossin. Enumeration of pin-permutations. *Electronic Journal of Combinatorics*, 18, 2011.

11 Anders Björner, Michel Las Vergnas, Bernd Sturmfels, Neil White, and Gunter M Ziegler. *Oriented matroids*. Number 46 in Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1999.

12 Mathilde Bouvel, Valentin Féray, Xavier Goaoc, and Florent Koechlin. A canonical tree decomposition for order types, and some applications, 2024. `arXiv:2403.10311`.

13 Cédric Chauve, Éric Fusy, and Jérémie Lumbroso. An exact enumeration of distance-hereditary graphs. In *2017 Proceedings of the Fourteenth Workshop on Analytic Algorithmics and Combinatorics (ANALCO'17)*, pages 31–45, 2017.

14 Tibor Gallai. Transitiv orientierbare Graphen. *Acta Mathematica Academiae Scientiarum Hungarica*, 18(1–2):25–66, March 1967.

15 Jacob E. Goodman and Richard Pollack. Multidimensional sorting. *SIAM J. Comput.*, 12(3):484–507, 1983. `doi:10.1137/0212032`.

16 Jacob E. Goodman and Richard Pollack. Allowable sequences and order types in discrete and computational geometry. In *New trends in discrete and computational geometry*, pages 103–134. Berlin: Springer-Verlag, 1993.

17 Michel Habib and Christophe Paul. A survey of the algorithmic aspects of modular decomposition. *Computer Science Review*, 4(1):42–59, February 2010.

18 Dániel Marx and Tillmann Miltzow. Peeling and Nibbling the Cactus: Subexponential-Time Algorithms for Counting Triangulations and Related Problems. In Sándor Fekete and Anna Lubiw, editors, *32nd International Symposium on Computational Geometry (SoCG 2016)*, volume 51 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 52:1–52:16, Dagstuhl, Germany, 2016. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. `doi:10.4230/LIPIcs.SoCG.2016.52`.

19 Rolf H. Möhring and Franz J. Radermacher. Substitution decomposition for discrete structures and connections with combinatorial optimization. *Annals of Discrete Math*, 19:257–356, 1984.

20 Daniel Rutschmann and Manuel Wettstein. Chains, Koch chains, and point sets with many triangulations. *Journal of the ACM*, 70(3):1–26, 2023.

21 Manfred Scheucher. A SAT attack on Erdős-Szekeres numbers in R^d and the empty hexagon theorem. *Computing in Geometry and Topology*, 2(1):2:1–2:13, 2023.

22 Micha Sharir and Adam Sheffer. Counting triangulations of planar point sets. *The Electronic Journal on Combinatorics*, 18(1):1–70, 2011.

23 Vikram Sharma and Chee K Yap. Robust geometric computation. In *Handbook of Discrete and Computational Geometry*, pages 1189–1223. Chapman and Hall/CRC, 2017.

24 Peter Shor. Stretchability of pseudolines is NP-hard. *Applied Geometry and Discrete Mathematics-The Victor Klee Festschrift*, 1991.

25 Andrew Suk and Ji Zeng. A positive fraction Erdős–Szekeres theorem and its applications. *Discrete & Computational Geometry*, pages 1–18, 2023.

26 Pavel Valtr. On mutually avoiding sets. In *The Mathematics of Paul Erdös II*, pages 324–328. Springer, 1997.