

Scaling Limits of Markov-Branching Trees and Applications



Lecture Notes of the XII Simposio de Probabilidad y Procesos Estocásticos 16–20 Novembre 2015, Mérida, Yucatán

Bénédicte Haas

Abstract The goal of these lecture notes is to survey some of the recent progress on the description of large-scale structure of random trees. We use the framework of Markov-Branching sequences of trees and discuss several applications.

Keywords Random trees · Scaling limits · Self-similar fragmentations · self-similar Markov processes

Mathematics Subject Classification 05C05, 60F17, 60J05, 60J25, 60J80

1 Introduction

The goal of these lecture notes is to survey some of the recent progress on the description of large-scale structure of random trees. Describing the structure of large (random) trees, and more generally large graphs, is an important goal of modern probabilities and combinatorics. Beyond the purely probabilistic or combinatorial aspects, motivations come from the study of models from biology, theoretical computer science or mathematical physics.

The question we will typically be interested in is the following. For $(T_n, n \geq 1)$ a sequence of random unordered (i.e. non-planar) trees, where, for each n , T_n is a tree of size n (the size of a tree may be its number of vertices or its number of leaves, for example): does there exist a deterministic sequence $(a_n, n \geq 1)$ and a *continuous*

B. Haas (✉)

Université Paris 13, Sorbonne Paris Cité, LAGA, CNRS (UMR 7539), Villetaneuse, France
e-mail: haas@math.univ-paris13.fr

random tree \mathcal{T} such that

$$\frac{T_n}{a_n} \xrightarrow{n \rightarrow \infty} \mathcal{T}?$$

To make sense of this question, we will view T_n as a metric space by “replacing” its edges with segments of length 1, and then use the notion of Gromov-Hausdorff distance to compare compact metric spaces. When such a convergence holds, the continuous limit highlights some properties of the discrete objects that approximate it, and vice-versa.

As a first example, consider T_n a tree picked uniformly at random in the set of trees with n vertices labelled by $\{1, \dots, n\}$. The tree T_n has to be understood as a *typical element* of this set of trees. In this case the answer to the previous question dates back to a series of works by Aldous in the beginning of the 1990s [8–10]: Aldous showed that

$$\frac{T_n}{2\sqrt{n}} \xrightarrow[n \rightarrow \infty]{(d)} \mathcal{T}_{\text{Br}} \quad (1)$$

where the limiting tree is called the Brownian Continuum Random Tree (CRT), and can be constructed from a standard Brownian excursion. This result has various interesting consequences, e.g. it gives the asymptotics in distribution of the diameter, the height (if we consider rooted versions of the trees) and several other statistics related to the tree T_n . Consequently it also gives the asymptotic proportion of trees with n labelled vertices that have a diameter larger than $x\sqrt{n}$ or/and a height larger than $y\sqrt{n}$, etc. Some of these questions on statistics of uniform trees were already treated in previous works, the strength of Aldous’s result is that it describes the asymptotics of the *whole* tree T_n .

Aldous has actually established a version of the convergence (1) in a much broader context, that of *conditioned Galton–Watson trees* with finite variance. In this situation, to fit to our context, T_n is an unordered version of the genealogical tree of a Galton–Watson process (with a given, fixed offspring distribution with mean one and finite variance) conditioned on having a total number of vertices equal to n , $n \geq 1$. Multiplied by $1/\sqrt{n}$, this tree converges in distribution to the Brownian CRT multiplied by a constant that only depends on the variance of the offspring distribution. This should be compared with (and is related to) the convergence of rescaled sums of i.i.d. random variables towards the normal distribution and its functional analog, the convergence of rescaled random walks towards the Brownian motion. It turns out that the above sequence of uniform labelled trees can be seen as a sequence of conditioned Galton–Watson trees (when the offspring distribution is a Poisson distribution) and more generally that several sequences of *combinatorial* trees reduce to conditioned Galton–Watson trees. In the early 2000s, Duquesne [44] extended Aldous’s result to conditioned Galton–Watson trees with offspring distributions in the domain of attraction of a stable law. We also refer to [46, 70] for related results. In most of these cases the scaling sequences (a_n) are asymptotically much smaller, i.e. $a_n \ll \sqrt{n}$, and other continuous trees arise in the limit, the so-called family of *stable Lévy trees*. All these results on conditioned Galton–Watson

trees are now well established, and have a lot of applications in the study of large random graphs (see e.g. Miermont's book [78] for the connections with random maps and Addario-Berry et al. [4] for connections with Erdős–Rényi random graphs in the critical window).

The classical proofs to establish the scaling limits of Galton–Watson trees consist in considering specific ordered versions of the trees and rely on a careful study of their so-called *contour functions*. It is indeed a common approach to encode trees into functions (similarly to the encoding of the Brownian tree by the Brownian excursion), which are more familiar objects. It turns out that for Galton–Watson trees, the contour functions are closely related to random walks, whose scaling limits are well known. Let us also mention that another common approach to study large random combinatorial structures is to use technics of analytic combinatorics, see [54] for a complete overview of the topic. None of these two methods will be used here.

In these lecture notes, we will focus on another point of view, that of sequences of random trees that satisfy a certain *Markov-Branching property*, which appears naturally in a large set of models and includes conditioned Galton–Watson trees. This property is a sort of discrete fragmentation property which roughly says that in each tree of the sequence, the subtrees above a given height are independent with a law that depends only on their total size. Under appropriate assumptions, we will see that Markov-Branching sequences of trees, suitably rescaled, converge to a family of continuous fractal trees, called the *self-similar fragmentation trees*. These continuous trees are related to the self-similar fragmentation processes studied by Bertoin in the 2000s [14], which are models used to describe the evolution of objects that randomly split as time passes. The main results on Markov-Branching trees presented here were developed in the paper [59], which has its roots in the earlier paper [63]. Several applications have been developed in these two papers, and in more recent works [15, 60, 89]: to Galton–Watson trees with arbitrary degree constraints, to several combinatorial trees families, including the Pólya trees (i.e. trees uniformly distributed in the set of rooted, unlabelled, unordered trees with n vertices, $n \geq 1$), to several examples of dynamical models of tree growth and to sequence of *cut-trees*, which describe the genealogy of some deletion procedure of edges in trees. The objective of these notes is to survey and gather these results, as well as further related results.

In Sect. 2 below, we will start with a series of definitions related to discrete trees and then present several classical examples of sequences of random trees. We will also introduce there the Markov-Branching property. In Sect. 3 we set up the topological framework in which we will work, by introducing the notions of real trees and Gromov–Hausdorff topology. We also recall there the classical results of Aldous [9] and Duquesne [44] on large conditioned Galton–Watson trees. Section 4 is the core of these lecture notes. We present there the results on scaling limits of Markov-Branching trees, and give the main ideas of the proofs. The key ingredient is the study of an integer-valued Markov chain describing the sizes of the subtrees containing a typical leaf of the tree. Section 5 is devoted to the applications mentioned above. Last, Sect. 6 concerns further perspectives and related models (multi-type trees, local limits, applications to other random graphs).

All the sequences of trees we will encounter here have a power growth. There is however a large set of random trees that naturally arise in applications that do not have such a behavior. In particular, many models of trees arising in the analysis of algorithms have a logarithmic growth. See e.g. Drmota's book [42] for an overview of the most classical models. These examples do not fit into our framework.

2 Discrete Trees, Examples and Motivations

2.1 Discrete Trees

Our objective is mainly to work with unordered trees. We give below a precise definition of these objects and mention nevertheless the notions of ordered or/and labelled trees to which we will sometimes refer.

A **discrete tree** (or graph-theoretic tree) is a finite or countable graph (V, E) that is connected and has no cycle. Here V denotes the set of vertices of the graph and E its set of edges. Note that two vertices are then connected by exactly one path and that $\#V = \#E + 1$ when the tree is finite.

In the following, we will often denote a (discrete) tree by the letter \mathfrak{t} , and for $\mathfrak{t} = (V, E)$ we will use the slight abuse of notation $v \in \mathfrak{t}$ to mean $v \in V$.

A tree \mathfrak{t} can be seen as a metric space, when endowed with the **graph distance** d_{gr} : given two vertices $u, v \in \mathfrak{t}$, $d_{\text{gr}}(u, v)$ is defined as the number of edges of the unique path from u to v .

A **rooted tree** (\mathfrak{t}, ρ) is an ordered pair where \mathfrak{t} is a tree and $\rho \in \mathfrak{t}$. The vertex ρ is then called the root of \mathfrak{t} . This gives a genealogical structure to the tree. The root corresponds to the generation 0, its neighbors can be interpreted as its children and form the generation 1, the children of its children form the generation 2, etc. We will usually call the **height** of a vertex its generation, and denote it by $\text{ht}(v)$ (the height of a vertex is therefore its distance to the root). The height of the tree is then

$$\text{ht}(\mathfrak{t}) = \sup_{v \in \mathfrak{t}} \text{ht}(v)$$

and its **diameter**

$$\text{diam}(\mathfrak{t}) = \sup_{u, v \in \mathfrak{t}} d_{\text{gr}}(u, v).$$

The **degree of a vertex** $v \in \mathfrak{t}$ is the number of connected components obtained when removing v (in other words, it is the number of neighbors of v). A vertex v different from the root and of degree 1 is called a **leaf**. In a rooted tree, the **out-degree of a vertex** v is the number of children of v . Otherwise said, $\text{out-degree}(v) = \text{degree}(v) - \mathbb{1}_{\{v \neq \text{root}\}}$. A (full) **binary tree** is a rooted tree where all vertices but the leaves have out-degree 2. A **branch-point** is a vertex of degree at least 3.

In these lecture notes, we will mainly work with rooted trees. Moreover we will consider, unless specifically mentioned, that two isomorphic trees are equal, or, when the trees are rooted, that **two root-preserving isomorphic trees are equal**. Such trees can be considered as *unordered unlabelled* trees, in opposition to the following definitions.

Ordered or/and Labelled Trees In the context of rooted trees, it may happen that one needs to order the children of the root, and then, recursively, the children of each vertex in the tree. This gives an **ordered** (or planar) tree. Formally, we generally see such a tree as a subset of the infinite Ulam–Harris tree

$$\mathcal{U} = \bigcup_{n=0}^{\infty} \mathbb{N}^n$$

where $\mathbb{N} := \{1, 2, \dots\}$ and $\mathbb{N}^0 = \{\emptyset\}$. The element \emptyset is the root of the Ulam–Harris tree, and any other $u = u_1 u_2 \dots u_n \in \mathcal{U} \setminus \{\emptyset\}$ is connected to the root via the unique shortest path

$$\emptyset \rightarrow u_1 \rightarrow u_1 u_2 \rightarrow \dots \rightarrow u_1 \dots u_n.$$

The height (or generation) of such a sequence u is therefore its length, n . We then say that $\mathfrak{t} \subset \mathcal{U}$ is a (finite or infinite) rooted *ordered* tree if:

- $\emptyset \in \mathfrak{t}$
- if $u = u_1 \dots u_n \in \mathfrak{t} \setminus \{\emptyset\}$, then $u = u_1 \dots u_{n-1} \in \mathfrak{t}$ (the parent of an individual in \mathfrak{t} that is not the root is also in \mathfrak{t})
- if $u = u_1 \dots u_n \in \mathfrak{t}$, there exists an integer $c_u(\mathfrak{t}) \geq 0$ such that the element $u_1 \dots u_n j \in \mathfrak{t}$ if and only if $1 \leq j \leq c_u(\mathfrak{t})$.

The number $c_u(\mathfrak{t})$ corresponds to the number of children of u in \mathfrak{t} , i.e., its out-degree.

We will also sometimes consider **labelled** trees. In these cases, the vertices are labelled in a bijective way, typically by $\{1, \dots, n\}$ if there are n vertices (whereas in an unlabelled tree, the vertices but the root are indistinguishable). Partial labelling is also possible, e.g. by labelling only the leaves of the tree.

In the following we will always specify when a tree is ordered or/and labelled. When not specified, it is implicitly unlabelled, unordered.

Counting Trees It is sometimes possible, but not always, to have explicit formulae for the number of trees of a specific structure. For example, it is known that the number of trees with n labelled vertices is

$$n^{n-2} \quad (\text{Cayley formula}),$$

and consequently, the number of rooted trees with n labelled vertices is

$$n^{n-1}.$$

The number of rooted ordered binary trees with $n + 1$ leaves is

$$\frac{1}{n+1} \binom{2n}{n}$$

(this number is called the n th Catalan number) and the number of rooted ordered trees with n vertices is

$$\frac{1}{n} \binom{2n-2}{n-1}.$$

On the other hand, there is no explicit formula for the number of rooted (unlabelled, unordered) trees. Otter [79] shows that this number is asymptotically proportional to

$$c\kappa^n n^{-3/2}$$

where $c \sim 0.4399$ and $\kappa \sim 2.9557$. This should be compared to the asymptotic expansion of the n th Catalan number, which is proportional (by Stirling's formula) to $\pi^{-1/2} 4^n n^{-3/2}$.

We refer to the book of Drmota [42] for more details and technics, essentially based on generating functions.

2.2 First Examples

We now present a first series of classical families of random trees. Our goal will be to describe their scaling limits when the sizes of the trees grow, as discussed in the Introduction. This will be done in Sect. 5. Most of these families (but not all) share a common property, the Markov-Branching property that will be introduced in the next section.

Combinatorial Trees Let \mathbb{T}_n denote a finite set of trees with n vertices, all sharing some structural properties. E.g. \mathbb{T}_n may be the set of all rooted trees with n vertices, or the set of all rooted ordered trees with n vertices, or the set of all binary trees with n vertices, etc. We are interested in the asymptotic behavior of a “typical element” of \mathbb{T}_n as $n \rightarrow \infty$. That is, we pick a tree *uniformly* at random in \mathbb{T}_n , denote it by T_n and study its scaling limit. The global behavior of T_n as $n \rightarrow \infty$ will represent some of the features shared by most of the trees. For example, if the probability that the height of T_n is larger than $n^{\frac{1}{2}+\varepsilon}$ tends to 0 as $n \rightarrow \infty$, this means that the proportion of trees in the set that have a height larger than $n^{\frac{1}{2}+\varepsilon}$ is asymptotically negligible, etc. We will more specifically be interested in the following cases:

- T_n is a uniform rooted tree with n vertices
- T_n is a uniform rooted ordered tree with n vertices

- T_n is a uniform tree with n labelled vertices
- T_n is a uniform rooted ordered binary tree with n vertices (n odd)
- T_n is a uniform rooted binary tree with n vertices (n odd),

etc. Many variations are of course possible, in particular one may consider trees picked uniformly amongst sets of trees with a given structure and n leaves, or more general degree constraints. Some of these uniform trees will appear again in the next example.

Galton–Watson Trees Galton–Watson trees are random trees describing the genealogical structure of Galton–Watson processes. These are simple mathematical models for the evolution of a population that continue to play an important role in probability theory and in applications. Let η be a probability on \mathbb{Z}_+ (η is called the *offspring distribution*) and let $m := \sum_{i \geq 1} i \eta(i) \in [0, \infty]$ denote its mean. Informally, in a Galton–Watson tree with offspring distribution η , each vertex has a random number of children distributed according to η , independently. We will always assume that $\eta(1) < 1$ in order to avoid the trivial case where each individual has a unique child. Formally, an η -Galton–Watson tree T^η is usually seen as an ordered rooted tree and defined as follows (recall the Ulam–Harris notation \mathcal{U}):

- $c_\emptyset(T^\eta)$ is distributed according to η
- conditionally on $c_\emptyset(T^\eta) = p$, the p ordered subtrees $\tau_i = \{u \in \mathcal{U} : iu \in T^\eta\}$ descending from $i = 1, \dots, p$ are independent and distributed as T^η .

From this construction, one sees that the distribution of T^η is given by:

$$\mathbb{P}(T^\eta = \mathbf{t}) = \prod_{v \in \mathbf{t}} \eta_{c_v(\mathbf{t})} \quad (2)$$

for all rooted ordered tree \mathbf{t} . This definition of Galton–Watson trees as ordered trees is the simplest, avoiding any symmetry problems. However in the following we will mainly see these trees up to isomorphism, which roughly means that we can “forget the order”.

Clearly, if we call Z_k the number of individuals at height k , then $(Z_k, k \geq 1)$ is a Galton–Watson process starting from $Z_0 = 1$. It is well known that the extinction time of this process,

$$\inf\{k \geq 0 : Z_k = 0\}$$

is finite with probability 1 when $m \leq 1$ and with a probability $\in [0, 1)$ when $m > 1$. The offspring distribution η and the tree T^η are said to be *subcritical* when $m < 1$, *critical* when $m = 1$ and *supercritical* when $m > 1$. From now on, we assume that

$$m = 1$$

and for integers n such that $\mathbb{P}(\#T^\eta = n) > 0$, we let $T_n^{\eta, \vee}$ denote a *non-ordered* version of the Galton–Watson tree T^η conditioned to have n vertices. Sometimes, we

will need to keep the order and we will let $T_n^{\eta, \text{v}, \text{ord}}$ denote this ordered conditioned version. We point out that in most cases, *but not all*, a subcritical or a supercritical Galton–Watson tree conditioned to have n vertices is distributed as a critical Galton–Watson tree conditioned to have n vertices with a different offspring distribution. So the assumption $m = 1$ is not too restrictive. We refer to [66] for details on that point.

It turns out that conditioned Galton–Watson trees are closely related to combinatorial trees. Indeed, one can easily check with (2) that:

- if $\eta = \text{Geo}(1/2)$, $T_n^{\eta, \text{v}, \text{ord}}$ is uniform amongst the set of rooted ordered trees with n vertices
- if $\eta = \text{Poisson}(1)$, $T_n^{\eta, \text{v}}$ is uniform amongst the set of rooted trees with n labelled vertices
- if $\eta = \frac{1}{2}(\delta_0 + \delta_2)$, $T_n^{\eta, \text{v}, \text{ord}}$ is uniform amongst the set of rooted ordered binary trees with n vertices.

We refer e.g. to Aldous [9] for additional examples.

Hence, studying the large-scale structure of conditioned Galton–Watson trees will also lead to results in the context of combinatorial trees. As mentioned in the Introduction, the scaling limits of large conditioned Galton–Watson trees are now well known. Their study has been initiated by Aldous [8–10] and then expanded by Duquesne [44]. This will be reviewed in Sect. 3. However, there are some sequences of combinatorial trees that *cannot* be reinterpreted as Galton–Watson trees, starting with the example of the uniform rooted tree with n vertices or the uniform rooted binary tree with n vertices. Studying the scaling limits of these trees remained open for a while, because of the absence of symmetry properties. These scaling limits are presented in Sect. 5.2.

In another direction, one may also wonder what happens when considering versions of Galton–Watson trees conditioned to have n leaves, instead of n vertices, or more general degree constraints. This is discussed in Sect. 5.1.2.

Dynamical Models of Tree Growth We now turn to several sequences of finite rooted random trees that are built recursively by adding at each step new edges on the pre-existing tree. We start with a well known algorithm that Rémy [88] introduced to generate uniform binary trees with n leaves.

Rémy’s Algorithm The sequence $(T_n(\mathbb{R}), n \geq 1)$ is constructed recursively as follows:

- Step 1: $T_1(\mathbb{R})$ is the tree with one edge and two vertices: one root, one leaf
- Step n : given $T_{n-1}(\mathbb{R})$, choose uniformly at random one of its edges and graft on “its middle” one new edge-leaf. By this we mean that the selected edge is split into two so as to obtain two edges separated by a new vertex, and then a new edge-leaf is glued on the new vertex. This gives $T_n(\mathbb{R})$.

It turns out (see e.g. [88]) that the tree $T_n(\mathbb{R})$, to which has been subtracted the edge between the root and the first branch point, is distributed as a binary critical Galton–Watson tree conditioned to have $2n - 1$ vertices, or equivalently n leaves

(after forgetting the order in the GW-tree). As so, we deduce its asymptotic behavior from that of Galton–Watson trees. However this model can be extended in several directions, most of which are not related to Galton–Watson trees. We detail three of them.

Ford’s α -Model [55] Let $\alpha \in [0, 1]$. We construct a sequence $(T_n(\alpha), n \geq 1)$ by modifying Rémy’s algorithm as follows:

- Step 1: $T_1(\alpha)$ is the tree with one edge and two vertices: one root, one leaf
- Step n : given $T_{n-1}(\alpha)$, give a weight $1 - \alpha$ to each edge connected to a leaf, and α to all other edges (the internal edges). The total weight is $n - 1 - \alpha$. Now, if $n \neq 2$ or $\alpha \neq 1$, choose an edge at random with a probability proportional to its weight and graft on “its middle” one new edge-leaf. This gives $T_n(\alpha)$. When $n = 2$ and $\alpha = 1$ the total weight is 0 and we decide to graft anyway on the middle of the edge of T_1 one new edge-leaf.

Note that when $\alpha = 1/2$ the weights are the same on all edges and we recover Rémy’s algorithm. When $\alpha = 0$, the new edge is always grafted uniformly on an edge-leaf, which gives a tree $T_n(0)$ known as the *Yule tree* with n leaves. When $\alpha = 1$, we obtain a deterministic tree called the *comb tree*. This family of trees indexed by $\alpha \in [0, 1]$ was introduced by Ford [55] in order to interpolate between the Yule, the uniform and the comb models. His goal was to propose new models for phylogenetic trees.

k-Ary Growing Trees [60] This is another extension of Rémy’s algorithm, where now several edges are added at each step. Consider an integer $k \geq 2$. The sequence $(T_n(k), n \geq 1)$ is constructed recursively as follows:

- Step 1: $T_1(k)$ is the tree with one edge and two vertices: one root, one leaf
- Step n : given $T_{n-1}(k)$, choose uniformly at random one of its edges and graft on “its middle” $k - 1$ new edges-leaf. This gives $T_n(k)$.

When $k = 2$, we recover Rémy’s algorithm. For larger k , there is no connection with Galton–Watson trees.

Marginals of Stable Trees: Marchal’s Algorithm In [73], Marchal considered the following algorithm, that attributes weights also to the vertices. Fix a parameter $\beta \in (1, 2]$ and construct the sequence $(T_n(\beta), n \geq 1)$ as follows:

- Step 1: $T_1(\beta)$ is the tree with one edge and two vertices: one root, one leaf
- Step n : given $T_{n-1}(\beta)$, attribute the weight
 - $\beta - 1$ on each edge
 - $d - 1 - \beta$ on each vertex of degree $d \geq 3$.

The total weight is $n\beta - 1$. Then select at random an edge or vertex with a probability proportional to its weight and graft on it a new edge-leaf. This gives $T_n(\beta)$.

The reason why Marchal introduced this algorithm is that $T_n(\beta)$ is actually distributed as the shape of a tree with edge-lengths that is obtained by sampling n leaves at random in the stable Lévy tree with index β . The class of stable Lévy trees plays an important role in the theory of random trees. It is introduced in Sect. 3.2 below.

Note that when $\beta = 2$, vertices of degree 3 are never selected (their weight is 0). So the trees $T_n(\beta)$, $n \geq 1$ are all binary, and we recover Rémy's algorithm.

Of course, several other extensions of trees built by adding edges recursively may be considered, some of which are mentioned in Sects. 5.3.3 and 6.1.

Remark In these dynamical models of tree growth, we build *on a same probability space* the sequence of trees, contrary to the examples of Galton–Watson trees or combinatorial trees that give sequences of *distributions* of trees. In this situation, one may expect to have more than a convergence in distribution for the rescaled sequences of trees. We will see in Sect. 5.3 that it is indeed the case.

2.3 The Markov-Branching Property

Markov-Branching trees were introduced by Aldous [11] as a class of random binary trees for phylogenetic models and later extended to non-binary cases in Broutin et al. [30], and Haas et al. [63]. It turns out that many natural models of sequence of trees satisfy the **Markov-Branching property (MB-property)** for short, starting with the example of conditioned Galton–Watson trees and most of the examples of the previous section.

Consider

$$(T_n, n \geq 1)$$

a sequence of trees where T_n is a rooted (unordered, unlabelled) tree with n leaves. The MB-property is a property of the sequence of *distributions* of $T_n, n \geq 1$. Informally, the MB-property says that for each tree T_n , *given that*

the root of T_n splits it in p subtrees with respectively $\lambda_1 \geq \dots \geq \lambda_p$ leaves,

then T_n is distributed as the tree obtained by gluing on a common root p independent trees with respective distributions those of $T_{\lambda_1}, \dots, T_{\lambda_p}$. The way the leaves are distributed in the sub-trees above the root, in each T_n , for $n \geq 1$, will then allow to fully describe the distributions of the $T_n, n \geq 1$.

We now explain rigorously how to build such sequences of trees. We start with a sequence of probabilities $(q_n, n \geq 1)$, where for each n , q_n is a probability on the set of partitions of the integer n . If $n \geq 2$, this set is defined by

$$\mathcal{P}_n := \left\{ \lambda = (\lambda_1, \dots, \lambda_p), \lambda_i \in \mathbb{N}, \lambda_1 \geq \dots \geq \lambda_p \geq 1 : \sum_{i=1}^p \lambda_i = n \right\},$$

whereas if $n = 1$, $\mathcal{P}_1 := \{(1), \emptyset\}$ (we need to have a cemetery point). For a partition $\lambda \in \mathcal{P}_n$, we denote by $p(\lambda)$ its length, i.e. the number of terms in the sequence λ . The probability q_n will determine how the n leaves of T_n are distributed into the subtrees above its root. We call such a probability a *splitting distribution*. In order that effective splittings occur, we will always assume that

$$q_n((n)) < 1, \quad \forall n \geq 1.$$

We need to define a notion of *gluing* of trees. Consider t_1, \dots, t_p , p discrete rooted (unordered) trees. Informally, we want to glue them on a same common root in order to form a tree $\langle t_1, \dots, t_p \rangle$ whose root splits into the p subtrees t_1, \dots, t_p . Formally, this can e.g. be done as follows. Consider first ordered versions of the trees $t_1^{\text{ord}}, \dots, t_p^{\text{ord}}$ seen as subsets of the Ulam–Harris tree \mathcal{U} and then define a new ordered tree by

$$\langle t_1^{\text{ord}}, \dots, t_p^{\text{ord}} \rangle := \{\emptyset\} \cup_{i=1}^p i t_i^{\text{ord}}.$$

The tree $\langle t_1, \dots, t_p \rangle$ is then defined as the unordered version of $\langle t_1^{\text{ord}}, \dots, t_p^{\text{ord}} \rangle$.

Definition 2.1 For each $n \geq 1$, let q_n be a probability on \mathcal{P}_n such that $q_n((n)) < 1$. From the sequence $\mathbf{q} = (q_n, n \geq 1)$ we construct recursively a sequence of distributions $(\mathcal{L}_n^{\mathbf{q}})$ such that for all $n \geq 1$, $\mathcal{L}_n^{\mathbf{q}}$ is carried by the set of rooted trees with n leaves, as follows:

- $\mathcal{L}_1^{\mathbf{q}}$ is the distribution of a line-tree with $G + 1$ vertices and G edges where G is a geometric distribution:

$$\mathbb{P}(G = k) = q_1(\emptyset)(1 - q_1(\emptyset))^k, \quad k \geq 0,$$

- for $n \geq 2$, $\mathcal{L}_n^{\mathbf{q}}$ is the distribution of

$$\langle T_1, \dots, T_{p(\Lambda)} \rangle$$

where Λ is a partition of n distributed according to q_n , and given Λ , the trees $T_1, \dots, T_{p(\Lambda)}$ are independent with respective distributions $\mathcal{L}_{\Lambda_1}^{\mathbf{q}}, \dots, \mathcal{L}_{\Lambda_{p(\Lambda)}}^{\mathbf{q}}$.

A sequence $(T_n, n \geq 1)$ of random rooted trees such that $T_n \sim \mathcal{L}_n^{\mathbf{q}}$ for each $n \in \mathbb{N}$ is called a **MB-sequence of trees indexed by the leaves, with splitting distributions $(q_n, n \geq 1)$** .

This construction may be re-interpreted as follows: we start from a collection of n *indistinguishable* balls, and with probability $q_n(\lambda_1, \dots, \lambda_p)$, split the collection into p sub-collections with $\lambda_1, \dots, \lambda_p$ balls. Note that there is a chance $q_n((n)) < 1$ that the collection remains unchanged during this step of the procedure. Then, re-iterate the splitting operation independently for each sub-collection using this time the probability distributions $q_{\lambda_1}, \dots, q_{\lambda_p}$. If a sub-collection consists of a single




Fig. 1 A sample tree T_{11} . The first splitting arises with probability $q_{11}(4, 4, 3)$

ball, it can remain single with probability $q_1((1))$ or get wiped out with probability $q_1(\emptyset)$. We continue the procedure until all the balls are wiped out. The tree T_n is then the genealogical tree associated with this process: it is rooted at the initial collection of n balls and its n leaves correspond to the n isolated balls just before they are wiped out, See Fig. 1 for an illustration.

We can define similarly MB-sequences of (distributions of) trees indexed by their number of *vertices*. Consider here a sequence $(p_n, n \geq 1)$ such that p_n is a probability on \mathcal{P}_n with no restriction but

$$p_1((1)) = 1.$$

Mimicking the previous balls construction and starting from a collection of n indistinguishable balls, we first remove a ball, split the $n - 1$ remaining balls in sub-collections with $\lambda_1, \dots, \lambda_p$ balls with probability $p_{n-1}((\lambda_1, \dots, \lambda_p))$, and iterate independently on sub-collections until no ball remains. Formally, this gives:

Definition 2.2 For each $n \geq 1$, let p_n be a probability on \mathcal{P}_n , such that $p_1((1)) = 1$. From the sequence $(p_n, n \geq 1)$ we construct recursively a sequence of distributions $(\mathcal{V}_n^{\mathbf{p}})$ such that for all $n \geq 1$, $\mathcal{V}_n^{\mathbf{p}}$ is carried by the set of trees with n vertices, as follows:

- $\mathcal{V}_1^{\mathbf{p}}$ is the deterministic distribution of the tree reduced to one vertex,
- for $n \geq 2$, $\mathcal{V}_n^{\mathbf{p}}$ is the distribution of

$$\langle T_1, \dots, T_{p(\Lambda)} \rangle$$

where Λ is a partition of $n - 1$ distributed according to p_{n-1} , and given Λ , the trees $T_1, \dots, T_{p(\Lambda)}$ are independent with respective distributions $\mathcal{V}_{\Lambda_1}^{\mathbf{p}}, \dots, \mathcal{V}_{\Lambda_{p(\Lambda)}}^{\mathbf{p}}$.

A sequence $(T_n, n \geq 1)$ of random rooted trees such that $T_n \sim \mathcal{V}_n^{\mathbf{p}}$ for each $n \in \mathbb{N}$ is called a **MB-sequence of trees indexed by the vertices, with splitting distributions** $(\mathbf{p}_n, n \geq 1)$.

More generally, the MB-property can be extended to sequences of trees $(T_n, n \geq 1)$ with arbitrary degree constraints, i.e. such that for all n , T_n has n vertices in A , where A is a given subset of \mathbb{Z}_+ . We will not develop this here and refer the interested reader to [89] for more details.

Some Examples

1. A deterministic example. Consider the splitting distributions on \mathcal{P}_n

$$q_n(\lceil n/2 \rceil, \lfloor n/2 \rfloor) = 1, \quad n \geq 2,$$

as well as $q_1(\emptyset) = 1$. Let $(T_n, n \geq 1)$ the corresponding MB-sequence indexed by leaves. Then T_n is a deterministic discrete binary tree, whose root splits in two subtrees with both $n/2$ leaves when n is even, and respectively $(n+1)/2$, $(n-1)/2$ leaves when n is odd. Clearly, when $n = 2^k$, the height of T_n is exactly k , and more generally for large n , $\text{ht}(T_n) \sim \ln(n)/\ln(2)$.

2. A basic example. For $n \geq 2$, let q_n be the probability on \mathcal{P}_n defined by

$$q_n((n)) = 1 - \frac{1}{n^\alpha} \quad \text{and} \quad q_n(\lceil n/2 \rceil, \lfloor n/2 \rfloor) = \frac{1}{n^\alpha} \quad \text{for some } \alpha > 0,$$

and let $q_1(\emptyset) = 1$. Let $(T_n, n \geq 1)$ be an MB-sequence indexed by leaves with splitting distributions (q_n) . Then T_n is a discrete tree with vertices with degrees $\in \{1, 2, 3\}$ where the distance between the root and the first branch point (i.e. the first vertex of degree 3) is a Geometric distribution on \mathbb{Z}_+ with success parameter $n^{-\alpha}$. The two subtrees above this branch point are independent subtrees, independent of the Geometric r.v. just mentioned, and whose respective distances between the root and first branch point are Geometric distributions with respectively $(\lceil n/2 \rceil)^{-\alpha}$ and $(\lfloor n/2 \rfloor)^{-\alpha}$ parameters. Noticing the weak convergence

$$\frac{\text{Geo}(n^{-\alpha})}{n^\alpha} \xrightarrow[n \rightarrow \infty]{(d)} \text{Exp}(1)$$

one may expect that $n^{-\alpha} T_n$ has a limit in distribution. We will later see that it is indeed the case.

3. Conditioned Galton–Watson trees. Let $T_n^{\eta,1}$ be a Galton–Watson tree with offspring distribution η , conditioned on having n leaves, for integers n for which this is possible. The branching property is then preserved by conditioning and the sequence $(T_n^{\eta,1}, n : \mathbb{P}(\#\text{leaves } T^\eta) > 0)$ is Markov-Branching, with splitting

distributions

$$q_n^{\text{GW},\eta}(\lambda) = \eta(p) \times \frac{p!}{\prod_{i=1}^p m_i(\lambda)!} \times \frac{\prod_{i=1}^p \mathbb{P}(\#\text{leaves } T^\eta = \lambda_i)}{\mathbb{P}(\#\text{leaves } T^\eta = n)}$$

for all $\lambda \in \mathcal{P}_n, n \geq 2$, where $\#\text{leaves } T^\eta$ is the number of leaves of the unconditioned Galton–Watson tree T^η , and $m_i(\lambda) = \#\{1 \leq j \leq p : \lambda_j = i\}$. The probability $q_1^{\text{GW},\eta}$ is given by $q_1^{\text{GW},\eta}((1)) = \eta(1)$.

Similarly, if $T_n^{\eta,\vee}$ denotes a Galton–Watson tree with offspring distribution η , conditioned on having n vertices, the sequence $(T_n^{\eta,\vee}, \mathbb{P}(\#\text{vertices } T^\eta) > 0)$ is MB, with splitting distributions

$$p_{n-1}^{\text{GW},\eta}(\lambda) = \eta(p) \times \frac{p!}{\prod_{i=1}^p m_i(\lambda)!} \times \frac{\prod_{i=1}^p \mathbb{P}(\#\text{vertices } T^\eta = \lambda_i)}{\mathbb{P}(\#\text{vertices } T^\eta = n)} \quad (3)$$

for all $\lambda \in \mathcal{P}_{n-1}, n \geq 3$ where $\#\text{vertices } T^\eta$ is the number of leaves of the unconditioned GW-tree T^η . Details can be found in [59, Section 5].

- 4. Dynamical models of tree growth.** Rémy’s, Ford’s, Marchal’s and the k -ary algorithms all lead to MB-sequences of trees indexed by leaves. To be precise, we have to remove in each of these trees the edge adjacent to the root to obtain MB-sequences of trees (the roots have all a unique child). The MB-property can be proved by induction on n . By construction, the distribution of the leaves in the subtrees above the root is closely connected to urns models. We have the following expressions for the splitting distributions:

Ford’s α -Model For $k \geq \frac{n}{2}, n \geq 2$,

$$\begin{aligned} q_n^{\text{Ford},\alpha}(k, n-k) &= \left(1 + \mathbb{1}_{k \neq \frac{n}{2}}\right) \frac{\Gamma(k-\alpha)\Gamma(n-k-\alpha)}{\Gamma(n-\alpha)\Gamma(1-\alpha)} \left(\frac{\alpha}{2} \binom{n}{k} + (1-2\alpha) \binom{n-2}{k-1}\right), \end{aligned}$$

and $q_1(\emptyset) = 1$. See [55] for details. In particular, taking $\alpha = 1/2$ one sees that

$$\begin{aligned} q_n^{\text{Rémy}}(k, n-k) &= \frac{1}{4} \left(1 + \mathbb{1}_{k \neq \frac{n}{2}}\right) \frac{\Gamma(k-1/2)\Gamma(n-k-1/2)}{\Gamma(n-1/2)\Gamma(1-1/2)} \binom{n}{k}, \quad k \geq \frac{n}{2}, n \geq 2. \end{aligned}$$

k-Ary Growing Trees Note that in these models, there are $1+(k-1)(n-1)$ leaves in the tree $T_n(k)$, so that the indices do not exactly correspond to the definitions of the Markov-Branching properties seen in the previous section. However, by relabelling, defining for $m = 1+(k-1)(n-1)$ the tree $\overline{T}_m(k)$ to be the tree $T_n(k)$ to which the edge adjacent to the root has been removed, we obtain an MB-sequence $(\overline{T}_m(k), m \in (k-1)\mathbb{N} + 2 - k)$. The splitting distributions are

defined for $m = 1 + (k - 1)(n - 1)$, $n \geq 2$ and $\lambda = (\lambda_1, \dots, \lambda_k) \in \mathcal{P}_m$ such that $\lambda_i = 1 + (k - 1)\ell_i$, for some $\ell_i \in \mathbb{Z}_+$ for all i (note that $\sum_{i=1}^k \ell_i = n - 2$) by

$$q_m^k(\lambda) = \sum_{\mathbf{n}=(n_1, \dots, n_k) \in \mathbb{N}^k: \mathbf{n}^\downarrow = \lambda} \bar{q}_m(\mathbf{n})$$

where \mathbf{n}^\downarrow is the decreasing rearrangement of the elements of \mathbf{n} and

$$\begin{aligned} \bar{q}_m(\mathbf{n}) &= \frac{1}{k(\Gamma(\frac{1}{k}))^{k-1}} \left(\prod_{i=1}^k \frac{\Gamma(\frac{1}{k} + n_i)}{n_i!} \right) \\ &\times \frac{(n-2)!}{\Gamma(\frac{1}{k} + n - 1)} \left(\sum_{j=1}^{n_1+1} \frac{n_1!}{(n_1 - j + 1)!} \frac{(n-j-1)!}{(n-2)!} \right). \end{aligned}$$

See [60, Section 3].

Marchal's Algorithm For $\lambda = (\lambda_1, \dots, \lambda_p) \in \mathcal{P}_n$, $n \geq 2$,

$$\begin{aligned} q_n^{\text{Marchal}, \beta}(\lambda) &= \frac{n!}{\lambda_1! \dots \lambda_p! m_1(\lambda)! \dots m_n(\lambda)!} \frac{\beta^{2-p} \Gamma(2 - \beta^{-1}) \Gamma(p - \beta)}{\Gamma(n - \beta^{-1}) \Gamma(2 - \beta)} \prod_{i=1}^p \frac{\Gamma(\lambda_j - \beta^{-1})}{\Gamma(1 - \beta^{-1})} \end{aligned}$$

where $m_i(\lambda) = \#\{1 \leq j \leq p : \lambda_j = i\}$. This is a consequence of [46, Theorem 3.2.1] and [75, Lemma 5].

- 5. Cut-trees.** *Cut-tree of a uniform Cayley tree.* Consider C_n a uniform Cayley tree of size n , i.e. a tree picked uniformly at random amongst the set of rooted tree with n labelled vertices. This tree has the following *recursive property* (see Pitman [85, Theorem 5]): removing an edge uniformly at random in C_n gives two trees, which given their numbers of vertices, $k, n - k$ say, are independent uniform Cayley trees of respective sizes $k, n - k$. Now, consider the following deletion procedure: remove in C_n one edge uniformly at random, then remove another edge in the remaining set of $n - 2$ edges uniformly at random and so on until all edges have been removed. It was shown by Janson [65] and Panholzer [83] that the number of steps needed to isolate the root divided by \sqrt{n} converges in distribution to a Rayleigh distribution (i.e. with density $x \exp(-x^2/2)$ on \mathbb{R}_+). Bertoin [15] was more generally interested in the number of steps needed to isolate ℓ distinguished vertices, and in that aim he introduced the *cut-tree* T_n^{cut} of C_n . The tree T_n^{cut} is the genealogical tree of the above deletion procedure, i.e. it describes the genealogy of the connected components, see Fig. 2 for an illustration and [15] for a precise construction of T_n^{cut} . Let us just mention here that T_n^{cut} is a rooted binary tree with n leaves, and that Pitman's recursive property




Fig. 2 On the left, a version of the tree C_7 , with edges labelled in order of deletion. On the right the associated cut-tree T_7^{cut} , whose vertices are the different connected components arising in the deletion procedure

implies that $(T_n^{\text{cut}}, n \geq 1)$ is MB. The corresponding splitting probabilities are:

$$q_n^{\text{Cut,Cayley}}(k, n - k) = \frac{(n - k)^{n-k-1} k^{k-1} (n - 2)!}{(n - k)! k! n^{n-3}}, \quad n/2 < k \leq n - 1,$$

the calculations are detailed in [15, 84].

Cut-tree of a uniform recursive tree. A recursive tree with n vertices is a tree with vertices labelled by $1, \dots, n$, rooted at 1, such that the sequence of labels of vertices along any branch from the root to a leaf is increasing. It turns out that the cut-tree of a uniform recursive tree is also MB and with splitting probabilities

$$q_n^{\text{Cut,Recursive}}(k, n - k) = \frac{n}{(n - 1)} \left(\frac{1}{k(k + 1)} + \frac{1}{(n - k)(n - k + 1)} \right) \quad n/2 < k \leq n - 1,$$

see [16].

Remark The first example is a simple example of models where *macroscopic branchings are frequent*, unlike the second example where *macroscopic branchings are rare* (they occur with probability $n^{-\alpha} \rightarrow 0$). By macroscopic branchings, we mean that the way that the n leaves (or vertices) are distributed above the root gives at least two subtrees with a size proportional to n . Although it is not completely obvious yet, nearly all other examples above have rare macroscopic branchings (in a sense that will be specified later) and this is typically the context in which we will study the scaling limits of MB-trees. Typically the tree T_n will then grow as a power of n . When macroscopic branchings are frequent, there is no scaling limit in general for the Gromov–Hausdorff topology, a topology introduced in the next section. However it is known that the height of the tree T_n is then often of order $c \ln(n)$. This case has been studied in [30].

3 The Example of Galton–Watson Trees and Topological Framework

We start with an informal version of the prototype result of Aldous on the description of the scaling limits of conditioned Galton–Watson trees. Let η be a critical offspring distribution with *finite variance* $\sigma^2 \in (0, \infty)$, and let $T_n^{\eta, \nu}$ denote a Galton–Watson tree with offspring distribution η , conditioned to have n vertices (in the following it is implicit that we only consider integers n such that this conditioning is possible). Aldous [10] showed that

$$\frac{\sigma}{2} \times \frac{T_n^{\eta, \nu}}{\sqrt{n}} \xrightarrow[n \rightarrow \infty]{(d)} \mathcal{T}_{\text{Br}} \quad (4)$$

where the continuous tree \mathcal{T}_{Br} arising in the limit is the Brownian Continuum Random Tree, sometimes simply called the Brownian tree. Note that the limit only depends on η via its variance σ^2 .

This result by Aldous was a breakthrough in the study of large random trees, since it was the first to describe the behavior of the tree as a whole. We will discuss this in more details in Sect. 3.2. Let us first introduce the topological framework in order to make sense of this convergence.

3.1 Real Trees and the Gromov–Hausdorff Topology

Since the pioneering works of Evans et al. [52] in 2003 and Duquesne and Le Gall [47] in 2005, the theory of *real trees* (or \mathbb{R} -trees) has been intensively used in probability. These trees are metric spaces having a “tree property” (roughly, this means that for each pair of points x, y in the metric space, there is a unique path going from x to y —see below for a precise definition). This point of view allows behavior such as infinite total length of the tree, vertices with infinite degree, and density of the set of leaves.

In these lecture notes, all the real trees we will consider are compact metric spaces. For this reason, we restrict ourselves to the theory of compact real trees. We now briefly recall background on real trees and the Gromov–Hausdorff and Gromov–Hausdorff–Prokhorov distances, and refer to [51, 71] for more details on this topic.

Real Trees A *real tree* is a metric space (\mathcal{T}, d) such that, for any points x and y in \mathcal{T} ,

- there is an isometry $\varphi_{x,y} : [0, d(x, y)] \rightarrow \mathcal{T}$ such that $\varphi_{x,y}(0) = x$ and $\varphi_{x,y}(d(x, y)) = y$
- for every continuous, injective function $c : [0, 1] \rightarrow \mathcal{T}$ with $c(0) = x, c(1) = y$, one has $c([0, 1]) = \varphi_{x,y}([0, d(x, y)])$.

Note that a discrete tree may be seen as a real tree by “replacing” its edges by line segments. Unless specified, it will be implicit throughout these notes that these line segments are all of length 1.

We denote by $[[x, y]]$ the line segment $\varphi_{x,y}([0, d(x, y)])$ between x and y . A *rooted* real tree is an ordered pair $((\mathcal{T}, d), \rho)$ such that (\mathcal{T}, d) is a real tree and $\rho \in \mathcal{T}$. This distinguished point ρ is called the root. The height of a point $x \in \mathcal{T}$ is defined by

$$\text{ht}(x) = d(\rho, x)$$

and the height of the tree itself is the supremum of the heights of its points, while the diameter is the supremum of the distance between two points:

$$\text{ht}(\mathcal{T}) = \sup_{x \in \mathcal{T}} d(\rho, x) \quad \text{diam}(\mathcal{T}) = \sup_{x, y \in \mathcal{T}} d(x, y).$$

The *degree* of a point x is the number of connected components of $\mathcal{T} \setminus \{x\}$. We call *leaves* of \mathcal{T} all the points of $\mathcal{T} \setminus \{\rho\}$ which have degree 1. Given two points x and y , we define $x \wedge y$ as the unique point of \mathcal{T} such that $[[\rho, x]] \cap [[\rho, y]] = [[\rho, x \wedge y]]$. It is called the *branch point* of x and y if its degree is larger than or equal to 3. For $a > 0$, we define the rescaled tree $a\mathcal{T}$ as (\mathcal{T}, ad) (the metric d thus being implicit and dropped from the notation).

As mentioned above, we will only consider compact real trees. We now want to measure how close two such metric spaces are. We start by recalling the definition of Hausdorff distance between compact subsets of a metric space.

Hausdorff Distance If A and B are two nonempty compact subsets of a metric space (E, d) , the Hausdorff distance between A and B is defined by

$$d_{E, \text{H}}(A, B) = \inf \{ \varepsilon > 0; A \subset B^\varepsilon \text{ and } B \subset A^\varepsilon \},$$

where A^ε and B^ε are the closed ε -enlargements of A and B , i.e. $A^\varepsilon = \{x \in E : d(x, A) \leq \varepsilon\}$ and similarly for B^ε .

The Gromov–Hausdorff extends this concept to compact real trees (or more generally compact metric spaces) that are not necessarily compact subsets of a single metric space, by considering embeddings in an arbitrary common metric space.

Gromov–Hausdorff Distance Given two compact rooted trees (\mathcal{T}, d, ρ) and $(\mathcal{T}', d', \rho')$, let

$$d_{\text{GH}}(\mathcal{T}, \mathcal{T}') = \inf \{ \max (d_{\mathcal{Z}, \text{H}}(\phi(\mathcal{T}), \phi'(\mathcal{T}')), d_{\mathcal{Z}}(\phi(\rho), \phi'(\rho'))) \},$$

where the infimum is taken over all pairs of isometric embeddings ϕ and ϕ' of \mathcal{T} and \mathcal{T}' in the same metric space $(\mathcal{Z}, d_{\mathcal{Z}})$, for all choices of metric spaces $(\mathcal{Z}, d_{\mathcal{Z}})$.

We will also be concerned with *measured* trees, that are real trees equipped with a probability measure on their Borel sigma-field. To this effect, recall first the definition of the Prokhorov distance between two probability measures μ and μ' on a metric space (E, d) :

$$d_{E,P}(\mu, \mu') = \inf \{ \varepsilon > 0 ; \forall A \in \mathcal{B}(E), \mu(A) \leq \mu'(A^\varepsilon) + \varepsilon \text{ and } \mu'(A) \leq \mu(A^\varepsilon) + \varepsilon \}.$$

This distance metrizes the weak convergence on the set of probability measures on (E, d) .

Gromov–Hausdorff–Prokhorov Distance Given two measured compact rooted trees $(\mathcal{T}, d, \rho, \mu)$ and $(\mathcal{T}', d', \rho', \mu')$, we let

$$d_{\text{GHP}}(\mathcal{T}, \mathcal{T}') = \inf \{ \max (d_{\mathcal{Z},H}(\phi(\mathcal{T}), \phi'(\mathcal{T}')), d_{\mathcal{Z}}(\phi(\rho), \phi'(\rho')), d_{\mathcal{Z},P}(\phi_*\mu, \phi'_*\mu')) \},$$

where the infimum is taken on the same space as before and $\phi_*\mu, \phi'_*\mu'$ are the push-forwards of μ, μ' by ϕ, ϕ' .

The Gromov–Hausdorff distance d_{GH} indeed defines a distance on the set of compact rooted real trees taken up to root-preserving isomorphisms. Similarly, The Gromov–Hausdorff–Prokhorov distance d_{GHP} is a distance on the set of compact measured rooted real trees taken up to root-preserving and measure-preserving isomorphisms. Moreover these two metric spaces are Polish, see [52] and [3]. We will always identify two (measured) rooted \mathbb{R} -trees when they are isometric and still use the notation (\mathcal{T}, d) (or \mathcal{T} when the choice of the metric is clear) to design their isometry class.

Statistics It is easy to check that the function that associates with a compact rooted tree its diameter is continuous (with respect to the GH-topology on the set of compact rooted real trees and the usual topology on \mathbb{R}). Similarly, the function that associates with a compact rooted tree its height is continuous. The function that associates with a compact rooted measured tree the distribution of the height of a leaf chosen according to the probability on the tree is continuous as well (with respect to the GHP-topology on the set of compact rooted measured real trees and the weak topology on the set of probability measures on \mathbb{R}). Consequently, the existence of scaling limits with respect to the GHP-topology will directly imply scaling limits for the height, the diameter and the height of a typical vertex of the trees.

3.2 Scaling Limits of Conditioned Galton–Watson Trees

We can now turn to rigorous statements on the scaling limits of conditioned Galton–Watson trees. We reformulate the above result (4) by Aldous in the finite variance case and also present the result by Duquesne [44] when the offspring distribution η is heavy tailed, in the domain of attraction of a stable distribution. In the following,

η always denotes a critical offspring distribution, $T_n^{\eta, \nu}$ is an η -GW tree conditioned to have n vertices, and $\mu_n^{\eta, \nu}$ is the uniform probability on its vertices. The following convergences hold for the Gromov–Hausdorff–Prokhorov topology.

Theorem 3.1

- (i) (Aldous [10]) *Assume that η has a finite variance σ^2 . Then, there exists a random compact real tree, called the Brownian tree and denoted \mathcal{T}_{Br} , endowed with a probability measure μ_{Br} supported by its set of leaves, such that as $n \rightarrow \infty$*

$$\left(\frac{\sigma T_n^{\eta, \nu}}{2\sqrt{n}}, \mu_n^{\eta, \nu} \right) \xrightarrow[\text{GHP}]{(d)} (\mathcal{T}_{\text{Br}}, \mu_{\text{Br}}).$$

- (ii) (Duquesne [44]) *If $\eta_k \sim \kappa k^{-1-\alpha}$ as $k \rightarrow \infty$ for $\alpha \in (1, 2)$, then there exists a random compact real tree \mathcal{T}_α , called the stable Lévy tree with index α , endowed with a probability measure μ_α supported by its set of leaves, such that as $n \rightarrow \infty$*

$$\left(\frac{T_n^{\eta, \nu}}{n^{1-1/\alpha}}, \mu_n^{\eta, \nu} \right) \xrightarrow[\text{GHP}]{(d)} \left(\left(\frac{\alpha(\alpha-1)}{\kappa\Gamma(2-\alpha)} \right)^{1/\alpha} \alpha^{1/\alpha-1} \cdot \mathcal{T}_\alpha, \mu_\alpha \right).$$

The result by Duquesne actually extends to cases where the offspring distribution η is in the domain of attraction of a stable distribution with index $\alpha \in (1, 2]$. See [44] for details.

The Brownian tree was first introduced by Aldous in the early 1990s in the series of papers [8–10]. This tree can be constructed in several ways, the most common being the following. Let $(\mathbf{e}(t), t \in [0, 1])$ be a normalized Brownian excursion, which, formally, can be defined from a standard Brownian motion B by letting

$$\mathbf{e}(t) = \frac{|B_{g+t(d-g)}|}{\sqrt{d-g}}, \quad 0 \leq t \leq 1,$$

where $g := \sup\{s \leq 1 : B_s = 0\}$ and $d = \inf\{s \geq 1 : B_s = 0\}$ (note that $d - g > 0$ a.s. since $B_1 \neq 0$ a.s.). Then consider for $x, y \in [0, 1], x \leq y$, the non-negative quantity

$$d_{\mathbf{e}}(x, y) = \mathbf{e}(x) + \mathbf{e}(y) - 2 \inf_{z \in [x, y]} \{\mathbf{e}(z)\},$$

and then the equivalent relation $x \sim_{\mathbf{e}} y \Leftrightarrow d_{\mathbf{e}}(x, y) = 0$. It turns out that the quotient space $[0, 1]/\sim_{\mathbf{e}}$ endowed with the metric induced by $d_{\mathbf{e}}$ (which indeed gives a true metric) is a compact real tree. The Brownian excursion \mathbf{e} is called the *contour function* of this tree. Equipped with the measure $\mu_{\mathbf{e}}$, which is the push-forward of the Lebesgue measure on $[0, 1]$, this gives a version $([0, 1]/\sim_{\mathbf{e}}, d_{\mathbf{e}}, \mu_{\mathbf{e}})$ of the measured tree $(\mathcal{T}_{\text{Br}}, \mu_{\text{Br}})$. To get a better intuition of what this means, as well

as more details and other constructions of the Brownian tree, we refer to the three papers by Aldous [8–10] and to the survey by Le Gall [71].

In the early 2000s, the family of stable Lévy trees $(\mathcal{T}_\alpha, \alpha \in (1, 2))$ —where by convention \mathcal{T}_2 is $\sqrt{2} \cdot \mathcal{T}_{\text{Br}}$ —was introduced by Duquesne and Le Gall [46, 47] in the more general framework of *Lévy trees*, building on earlier work of Le Gall and Le Jan [72]. These trees can be constructed in a way similar as above from continuous functions built from the stable Lévy processes. This construction is complex and we will not detail it here. Others constructions are possible, see e.g. [50, 56]. The stable trees are important objects of the theory of random trees. They are intimately related to continuous state branching processes, fragmentation and coalescence processes. They appear as scaling limits of various models of trees and graphs, starting with the Galton–Watson examples above and some other examples discussed in Sect. 5. In particular, it is noted that it was only proved recently that Galton–Watson trees conditioned by their number of leaves or more general arbitrary degree restrictions also converge in the scaling limit to stable trees, see Sect. 5.1.2 and the references therein.

In the last few years, the geometric and fractal aspects of stable trees have been studied in great detail: Hausdorff and packing dimensions and measures [45, 47, 48, 57]; spectral dimension [34]; spinal decompositions and invariance under uniform re-rooting [49, 64]; fragmentation into subtrees [75, 76]; and embeddings of stable trees into each other [35]. We simply point out it here that the Brownian tree is *binary*, in the sense that all its points have their degree in $\{1, 2, 3\}$ almost surely, whereas the stable trees $\mathcal{T}_\alpha, \alpha \in (1, 2)$ have only points with degree in $\{1, 2, \infty\}$ almost surely (every branch point has an infinite number of “children”).

Applications to Combinatorial Trees Using the connections between some families of combinatorial trees and Galton–Watson trees mentioned in Sect. 2.1, we obtain the following scaling limits (in all cases, μ_n denotes the uniform probability on the vertices of the tree T_n):

- If T_n is uniform amongst the set of rooted ordered trees with n vertices,

$$\left(\frac{T_n}{\sqrt{n}}, \mu_n \right) \xrightarrow[\text{GHP}]{(d)} (\mathcal{T}_{\text{Br}}, \mu_{\text{Br}}).$$

- If T_n is uniform amongst the set of rooted trees with n labelled vertices,

$$\left(\frac{T_n}{\sqrt{n}}, \mu_n \right) \xrightarrow[\text{GHP}]{(d)} (2\mathcal{T}_{\text{Br}}, \mu_{\text{Br}}).$$

- If T_n is uniform amongst the set of rooted binary ordered trees with n vertices,

$$\left(\frac{T_n}{\sqrt{n}}, \mu_n \right) \xrightarrow[\text{GHP}]{(d)} (2\mathcal{T}_{\text{Br}}, \mu_{\text{Br}}).$$

As a consequence, this provides the behavior of several statistics of the trees, that first interested combinatorists.

We will not present the original proofs by Aldous [10] and Duquesne [44], but will rather focus on the fact that they may be recovered by using the MB-property. This is the goal of the next two sections, where we will present in a general setting some results on the scaling limits for MB-sequences of trees. As already mentioned, the main idea of the proofs of Aldous [10] and Duquesne [44] is rather based on the study of the so-called *contour functions* of the trees. We refer to Aldous and Duquesne papers, as well as Le Gall’s survey [71] for details. See also Duquesne and Le Gall [46] and Kortchemski [69, 70] for further related results.

4 Scaling Limits of Markov-Branching Trees

Our goal is to set up an asymptotic criterion on the splitting probabilities (q_n) of an MB-sequence of trees so that this sequence, suitably normalized, converges to a non-trivial continuous limit. We follow here the approach of the paper [59] that found its roots in the previous work [63] were similar results where proved under stronger assumptions. A remark on these previous results is made at the end of this section.

The splitting probability q_n corresponds to a “discrete” fragmentation of the integer n into smaller integers. To set up the desired criterion, we first need to introduce a continuous counterpart for these partitions of integers, namely

$$\mathcal{S}^\downarrow = \left\{ \mathbf{s} = (s_1, s_2, \dots) : s_1 \geq s_2 \geq \dots \geq 0 \text{ and } \sum_{i \geq 1} s_i = 1 \right\}$$

which is endowed with the distance $d_{\mathcal{S}^\downarrow}(\mathbf{s}, \mathbf{s}') = \sup_{i \geq 1} |s_i - s'_i|$. Our main hypothesis on (q_n) then reads:

Hypothesis (H) There exist $\gamma > 0$ and ν a non-trivial σ -finite measure on \mathcal{S}^\downarrow satisfying $\int_{\mathcal{S}^\downarrow} (1 - s_1) \nu(d\mathbf{s}) < \infty$ and $\nu(1, 0, \dots) = 0$, such that

$$n^\gamma \sum_{\lambda \in \mathcal{P}_n} q_n(\lambda) \left(1 - \frac{\lambda_1}{n}\right) f\left(\frac{\lambda_1}{n}, \dots, \frac{\lambda_p}{n}, 0, \dots\right) \xrightarrow{n \rightarrow \infty} \int_{\mathcal{S}^\downarrow} (1 - s_1) f(\mathbf{s}) \nu(d\mathbf{s}).$$

for all continuous $f : \mathcal{S}^\downarrow \rightarrow \mathbb{R}$.

We will see in Sect. 5 that most of the examples of splitting probabilities introduced in Sect. 2.3 satisfy this hypothesis. As a first, easy, example, consider the “basic example” introduced there (Example 2): $q_n((n)) = 1 - n^{-\alpha}$ and $q_n(\lceil n/2 \rceil, \lfloor n/2 \rfloor) = n^{-\alpha}$, $\alpha > 0$. Then, clearly, (H) is satisfied with

$$\gamma = \alpha \quad \text{and} \quad \nu(d\mathbf{s}) = \delta_{\left(\frac{1}{2}, \frac{1}{2}, 0, \dots\right)}.$$

The interpretation of the hypothesis (H) is that macroscopic branchings are rare, in the sense that the macroscopic splitting events $n \mapsto ns$, $\mathbf{s} \in \mathcal{S}^\downarrow$ with $s_1 < 1 - \varepsilon$ occur with a probability asymptotically proportional to $n^{-\gamma} \mathbb{1}_{\{s_1 < 1 - \varepsilon\}} \nu(\mathbf{ds})$, for a.e. fixed $\varepsilon \in (0, 1)$.

The main result on the scaling limits of MB-trees indexed by the leaves is the following.

Theorem 4.1 ([59]) *Let $(T_n, n \geq 1)$ be a MB-sequence indexed by the leaves and assume that its splitting probabilities satisfy (H). Then there exists a compact, measured real tree $(\mathcal{T}_{\gamma, \nu}, \mu_{\gamma, \nu})$ such that*

$$\left(\frac{T_n}{n^\gamma}, \mu_n \right) \xrightarrow[\text{GHP}]{(d)} (\mathcal{T}_{\gamma, \nu}, \mu_{\gamma, \nu}),$$

where μ_n is the uniform probability on the leaves of T_n .

The goal of this section is to detail the main steps of the proof of this result and to discuss some properties of the limiting measured tree, which belongs to the so-called family of *self-similar fragmentation trees* (the distribution of such a tree is entirely characterized by the parameters γ and ν). In that aim we will first study how the height of a leaf chosen uniformly at random in T_n grows (Sects. 4.1 and 4.2). Then we will review some results on self-similar fragmentation trees (Sect. 4.3). Last we will explain how one can use the scaling limit of the height of a leaf chosen at random to obtain, by induction, the scaling limit of the subtree spanned by k leaves chosen independently, for all k , and then finish the proof of Theorem 4.1 with a tightness criterion (Sect. 4.4).

There is a similar result for MB-sequences indexed by the vertices.

Theorem 4.2 ([59]) *Let $(T_n, n \geq 1)$ be a MB-sequence indexed by the vertices and assume that its splitting probabilities satisfy (H) for some $0 < \gamma < 1$. Then there exists a compact, measured real tree $(\mathcal{T}_{\gamma, \nu}, \mu_{\gamma, \nu})$ such that*

$$\left(\frac{T_n}{n^\gamma}, \mu_n \right) \xrightarrow[\text{GHP}]{(d)} (\mathcal{T}_{\gamma, \nu}, \mu_{\gamma, \nu}),$$

where μ_n is the uniform probability on the vertices of T_n .

Theorem 4.2 is actually a direct corollary of Theorem 4.1, for the following reason. Consider an MB-sequence indexed by the vertices with splitting probabilities (p_n) and for all n , branch on each internal vertex of the tree T_n an edge with a leaf. This gives a tree \bar{T}_n with n leaves. It is then obvious that $(\bar{T}_n, n \geq 1)$ is an MB-sequence indexed by the *leaves*, with splitting probabilities (q_n) defined by

$$q_n(\lambda_1, \dots, \lambda_p, 1) = p_{n-1}(\lambda_1, \dots, \lambda_p), \quad \text{for all } (\lambda_1, \dots, \lambda_p) \in \mathcal{P}_{n-1}$$

(and $q_n(\lambda) = 0$ for all other $\lambda \in \mathcal{P}_n$). It is moreover easy to see that (q_n) satisfies (H) with parameters (γ, ν) , $0 < \gamma < 1$, if and only if (p_n) does. Hence Theorem 4.1

implies Theorem 4.2, since \overline{T}_n , endowed with the uniform probability, is at distance less than one from (T_n, μ_n) for the GHP-distance.

We will present in Sect. 5 several applications of these two theorems. Let us just consider here the “basic example” of Sect. 2.3 (Example 2). We have already noticed that its splitting probabilities satisfy Hypothesis (H), with parameters α and $\delta_{(1/2, 1/2, 0, \dots)}$. Hence in this case, the corresponding sequence of MB-trees T_n divided by n^α and endowed with the uniform probability measure on its leaves converges for the GHP-topology towards a $(\alpha, \delta_{(1/2, 1/2, 0, \dots)})$ -self-similar fragmentation tree.

Remark These two statements are also valid when replacing in (H) and in the theorems the power sequence n^γ by any regularly varying sequence with index $\gamma > 0$. We recall that a sequence (a_n) is said to vary regularly with index $\gamma > 0$ if for all $c > 0$,

$$\frac{a_{\lfloor cn \rfloor}}{a_n} \xrightarrow{n \rightarrow \infty} c^\gamma.$$

We refer to [24] for backgrounds on that topic. For simplicity, in the following we will only work with power sequences, but the reader should have in mind that everything holds similarly for regularly varying sequences.

Convergence in Probability In [63], scaling limits are established for some MB-sequences that moreover satisfy a property of *sampling consistency*, namely that for all n , T_n is distributed as the tree with n leaves obtained by removing a leaf picked uniformly at random in T_{n+1} , as well as the adjacent edge. This consistency property is demanding and the approach developed in [59] allows to do without it. However we note that if the MB-sequence is *strongly sampling consistent*, one can actually establish under suitable conditions a *convergence in probability* of the rescaled trees, which is of course an improvement. By strongly sampling consistent, we mean that versions of the trees can be built on a same probability space so that if T_n° denotes the tree with n leaves obtained by removing an edge-leaf picked uniformly in T_{n+1} , then (T_n, T_{n+1}) is distributed as (T_n°, T_{n+1}) . We refer to [63] for details.

4.1 A Markov Chain in the Markov-Branching Sequence of Trees

Consider $(T_n, n \geq 1)$ an MB-sequence of trees indexed by the leaves, with splitting distribution $(q_n, n \geq 1)$. Before studying the scaling limit of the trees in their whole, we start by studying the scaling limit of a *typical leaf*. For example, in each T_n , we mark one of the n leaves uniformly at random and we want to determine how the height of the marked leaf behaves as $n \rightarrow \infty$. In that aim, let \star_n denote this marked leaf and let $\star_n(k)$ denote its ancestor at generation k , $0 \leq k \leq \text{ht}(\star_n)$ (so that $\star_n(0)$ is the root of T_n and $\star_n(\text{ht}(\star_n)) = \star_n$). Let also $T_n^*(k)$ be the subtree composed of




Fig. 3 A Markov chain in the Markov-Branching trees. Here $n = 9$, the marked leaf is circled and the subtrees of descendants of $\star_9(k)$ for $0 \leq k \leq 4$ are dotted. Moreover $X_9(0) = 9, X_9(1) = 5, X_9(2) = 3, X_9(3) = 2, X_9(4) = 1$ and $X_9(i) = 0, \forall i \geq 5$

the descendants of $\star_n(k)$ in T_n , formally,

$$T_n^*(k) := \{v \in T_n : \star_n(k) \in [[\rho, v]]\}, \quad k \leq \text{ht}(\star_n)$$

and $T_n^*(k) := \emptyset$ if $k > \text{ht}(\star_n)$. We then set

$$X_n(k) := \#\{\text{leaves of } T_n^*(k)\}, \quad \forall k \in \mathbb{Z}_+ \tag{5}$$

with the convention that $X_n(k) = 0$ for $k > \text{ht}(\star_n)$ (Fig. 3).

Proposition 4.3 *The process $(X_n(k), k \geq 0)$ is a \mathbb{Z}_+ -valued non-increasing Markov chain starting from $X_n(0) = n$, with transition probabilities*

$$p(i, j) = \sum_{\lambda \in \mathcal{P}_i} q_i(\lambda) m_j(\lambda) \frac{j}{i} \quad \text{for all } 1 \leq j \leq i, \text{ with } i \geq 1 \tag{6}$$

and $p(1, 0) = q_1(\emptyset) = 1 - p(1, 1)$.

Proof The Markov property is a direct consequence of the Markov branching property. Indeed, given $X_n(1) = i_1, \dots, X_n(k-1) = i_{k-1}$, the tree $T_n^*(k-1)$ is distributed as $T_{i_{k-1}}$ if $i_{k-1} \geq 1$ and is the emptyset otherwise. In particular, when $i_{k-1} = 0$, the conditional distribution of $X_n(k)$ is the Dirac mass at 0. When $i_{k-1} \geq 1$, we use the fact that \star_n is in $T_n^*(k-1)$, hence, still conditioning on the same event, we have that \star_n is uniformly distributed amongst the i_{k-1} leaves of $T_{i_{k-1}}$. Otherwise said, given $X_n(1) = i_1, \dots, X_n(k-1) = i_{k-1}$ with $i_{k-1} \geq 1$, $(T_n^*(k-1), \star_n)$ is distributed as $(T_{i_{k-1}}, \star_{i_{k-1}})$ and consequently $X_n(k)$ is distributed as $X_{i_{k-1}}(1)$. Hence the Markov property of the chain $(X_n(k), k \geq 0)$. It remains to compute the transition probabilities:

$$p(n, k) = \mathbb{P}(X_n(1) = k) = \sum_{\lambda \in \mathcal{P}_n} q_n(\lambda) \mathbb{P}(X_n(1) = k | \Lambda_n = \lambda)$$

where Λ_n denotes the partition of n corresponding to the distribution of the leaves in the subtrees of T_n above the root. Since \star_n is chosen uniformly amongst the set

of leaves, we clearly have that

$$\mathbb{P}(X_n(1) = k | \Lambda_n = \lambda) = \frac{k}{n} \times \#\{j : \lambda_j = k\}, \quad \forall k \geq 1.$$

□

Hence studying the scaling limit of the height of the marked leaf in the tree T_n reduces to studying the scaling limit of the absorption time A_n of the Markov chain $(X_n(k), k \geq 0)$ at 0:

$$A_n := \inf \{k \geq 0 : X_n(k) = 0\}$$

(to be precise, this absorption time is equal to the height of the marked leaf + 1). The study of the scaling limit of $((X_n(k), k \geq 1), A_n)$ as $n \rightarrow \infty$ is the goal of the next section. Before getting in there, let us notice that the Hypothesis (H) on the splitting probabilities $(q_n, n \geq 1)$ of $(T_n, n \geq 1)$, together with (6), implies the following behavior of the transition probabilities $(p(n, k), k \leq n)$:

$$n^\gamma \sum_{k=0}^n p(n, k) \left(1 - \frac{k}{n}\right) g\left(\frac{k}{n}\right) \xrightarrow{n \rightarrow \infty} \int_{[0,1]} g(x) \mu(dx) \quad (7)$$

for all continuous functions $g : [0, 1] \rightarrow \mathbb{R}$, where the measure μ in the limit is a finite, non-zero measure on $[0, 1]$ defined by

$$\int_{[0,1]} g(x) \mu(dx) = \int_{\mathcal{S}^\downarrow} \sum_{i \geq 1} s_i (1 - s_i) g(s_i) \nu(ds). \quad (8)$$

To see this, apply (H) to the continuous function defined by

$$f(\mathbf{s}) = \frac{\sum_{i \geq 1} s_i (1 - s_i) g(s_i)}{1 - s_1} \quad \text{for } \mathbf{s} \neq (1, 0, \dots)$$

and $f(1, 0, \dots) = g(1) + g(0)$.

4.2 Scaling Limits of Non-increasing Markov Chains

As discussed in the previous section, studying the height of a typical leaf in MB-trees amounts to studying the absorption time at 0 of a \mathbb{Z}_+ -valued non-increasing Markov chain. In this section, we study in a general framework the scaling limits of \mathbb{Z}_+ -valued non-increasing Markov chains, under appropriate assumptions on the transition probabilities. At the end of the section we will see how this applies to the

height of a typical leaf in an MB-sequence. In the following,

$$(X_n(k), k \geq 0)$$

denotes a non-increasing \mathbb{Z}_+ -valued Markov chain starting from n ($X_n(0) = n$), with transition probabilities ($p(i, j), 0 \leq j \leq i$) such that

Hypothesis (H') $\exists \gamma > 0$ and μ a non-trivial finite measure on $[0, 1]$ such that

$$n^\gamma \sum_{k=0}^n p(n, k) \left(1 - \frac{k}{n}\right) f\left(\frac{k}{n}\right) \xrightarrow{n \rightarrow \infty} \int_{[0,1]} f(x) \mu(dx)$$

for all continuous functions $f : [0, 1] \rightarrow \mathbb{R}$.

This hypothesis implies that starting from n , *macroscopic jumps* (i.e. with size proportional to n) are rare, since for a.e. $0 < \varepsilon \leq 1$, the probability to do a jump larger than εn is of order $c_\varepsilon n^{-\gamma}$ where $c_\varepsilon = \int_{[0,1-\varepsilon]} (1-x)^{-1} \mu(dx)$ (note that this may tend to ∞ when ε tends to 0).

Now, let

$$A_n := \inf \{k \geq 0 : X_n(i) = X_n(k), \quad \forall i \geq k\}$$

be the first time at which the chain enters an absorption state (note that $A_n < \infty$ a.s. since the chain is non-increasing and \mathbb{Z}_+ -valued). In the next theorem, $\mathbb{D}([0, \infty), [0, \infty))$ denotes the set of non-negative càdlàg processes, endowed with the Skorokhod topology.

Theorem 4.4 ([58]) *Assume (H').*

(i) *Then, in $\mathbb{D}([0, \infty), [0, \infty))$,*

$$\left(\frac{X_n(\lfloor n^\gamma t \rfloor)}{n}, t \geq 0 \right) \xrightarrow[n \rightarrow \infty]{(d)} (\exp(-\xi_{\tau(t)}), t \geq 0),$$

where ξ is a subordinator, i.e. a non-decreasing Lévy process, and τ is the time-change (acceleration of time)

$$\tau(t) := \inf \left\{ u \geq 0 : \int_0^u \exp(-\gamma \xi_r) dr \geq t \right\}, t \geq 0.$$

The distribution of ξ is characterized by its Laplace transform $\mathbb{E}[\exp(-\lambda \xi_t)] = \exp(-t\phi(\lambda))$, with

$$\phi(\lambda) = \mu(\{0\}) + \mu(\{1\})\lambda + \int_{(0,1)} (1-x^\lambda) \frac{\mu(dx)}{1-x}, \quad \lambda \geq 0.$$

(ii) Moreover, jointly with the above convergence,

$$\frac{A_n}{n^\gamma} \xrightarrow[n \rightarrow \infty]{(d)} \int_0^\infty \exp(-\gamma \xi_r) dr = \inf \{t \geq 0 : \exp(-\xi_{\tau(t)}) = 0\}.$$

Comments For background on Lévy processes, we refer to [12]. Let us simply recall here that the law of a subordinator is characterized by three parameters: a measure on $(0, \infty)$ that codes its jumps (which here is the push-forward of $\mu(dx)(1-x)\mathbb{1}_{\{x \in (0,1)\}}$ by the application $x \mapsto -\ln(x)$), a linear drift (here $\mu(\{1\})$) and a killing rate at which the process jumps to $+\infty$ (here $\mu(\{0\})$).

Main Ideas of the Proof of Theorem 4.4

(i) Let $Y_n(t) := n^{-1}X_n(\lfloor n^\gamma t \rfloor)$, for $t \geq 0, n \in \mathbb{N}$. First, using Aldous' tightness criterion [23, Theorem 16.10] and (H') , one can check that the sequence $(Y_n, n \geq 1)$ is tight. It is then sufficient to prove that every possible limit in distribution of subsequences of (Y_n) are distributed as $\exp(-\xi_\tau)$. Let Y' be such a limit and $(n_k, k \geq 1)$ a sequence such that Y_{n_k} converges to Y' in distribution. In the limit, we actually prefer to deal with ξ than with ξ_τ , and for this reason we start by changing time in Y_n by setting

$$\tau_{Y_n}(t) := \inf \left\{ u \geq 0 : \int_0^u Y_n^{-\gamma}(r) dr > t \right\} \quad \text{and} \quad Z_n(t) := Y_n(\tau_{Y_n}(t)), \quad t \geq 0.$$

One can then easily check that (Z_{n_k}) converges in distribution to Z' where $Z' = Y' \circ \tau_{Y'}$, with $\tau_{Y'}(t) := \inf \left\{ u \geq 0 : \int_0^u (Y'(r))^{-\gamma} dr > t \right\}$. It is also easy to reverse the time-change and get that

$$Y'(t) = Z'(\tau_{Y'}^{-1}(t)) = Z' \left(\inf \left\{ u \geq 0 : \int_0^u Z'^{\gamma}(r) dr > t \right\} \right), \quad t \geq 0.$$

With this last equality, we see that it just remains to prove that Z' is distributed as $\exp(-\xi)$. This can be done in three steps:

(a) Observe the following (easy!) fact: if P is the transition function of a Markov chain M with countable state space $\subset \mathbb{R}$, then for any positive function f such that $f^{-1}(\{0\})$ is absorbing,

$$f(M(k)) \prod_{i=0}^{k-1} \frac{f(M(i))}{Pf(M(i))}, \quad k \geq 0$$

is a martingale. As a consequence: for all $\lambda \geq 0$ and $n \geq 1$, if we let $G_n(\lambda) := \mathbb{E}[(X_n(1)/n)^\lambda]$, then,

$$M_n^{(\lambda)}(t) := Z_n^\lambda(t) \left(\prod_{i=0}^{\lfloor n^\gamma \tau_{Y_n}(t) \rfloor - 1} G_{X_n(i)}(\lambda) \right)^{-1}, \quad t \geq 0$$

is a martingale.

- (b) Under (H'), $1 - G_n(\lambda) \underset{n \rightarrow \infty}{\sim} n^{-\gamma} \phi(\lambda)$. Together with the convergence in distribution of (Z_{n_k}) to Z' and the definition of $M_n^{(\lambda)}$, this leads to the convergence (this is the most technical part)

$$M_{n_k}^{(\lambda)} \xrightarrow[k \rightarrow \infty]{(d)} (Z')^\lambda \exp(\phi(\lambda) \cdot),$$

and the martingale property passes to the limit.

- (c) Hence $(Z')^\lambda \exp(\phi(\lambda) \cdot)$ is a martingale for all $\lambda \geq 0$. Using Laplace transforms, it is then easy to see that this implies in turn that $-\ln Z'$ is a non-decreasing process with independent and stationary increments (hence a subordinator), with Laplace exponent ϕ .

Hence $Z' \stackrel{(d)}{=} \exp(-\xi)$.

- (ii) We do not detail this part and refer to [58, Section 4.3]. Let us simply point out that it is *not* a direct consequence of the convergence of (Y_n) to $\exp(-\xi_\tau)$ since convergence of functions in $\mathbb{D}([0, \infty), [0, \infty))$ does not lead, in general, to the convergence of their absorption times (when they exist). \square

This result leads to the following corollary.

Corollary 4.5 *Let $(T_n, n \geq 1)$ be a MB-sequence indexed by the leaves, with splitting probabilities satisfying (H) with parameters (γ, ν) . For each n , let \star_n be a leaf chosen uniformly amongst the n leaves of T_n . Then,*

$$\frac{\text{ht}(\star_n)}{n^\gamma} \xrightarrow[n \rightarrow \infty]{(d)} \int_0^\infty \exp(-\gamma \xi_r) dr$$

where ξ is a subordinator with Laplace exponent $\phi(\lambda) = \int_{\mathcal{S}^\downarrow} \sum_{i \geq 1} (1 - s_i^\lambda) s_i \nu(ds)$, $\lambda \geq 0$.

Proof As seen at the end of the previous section, under (H) the transition probabilities of the Markov chain (5) satisfy assumption (H') with parameters γ and μ , with μ defined by (8). The conclusion follows with Theorem 4.4 (ii). \square

Further Reading Apart from applications to Markov-Branching trees, Theorem 4.4 can be used to describe the scaling limits of various stochastic processes, e.g. random walks with a barrier or the number of collisions in Λ -coalescent

processes, see [58]. Recently, Bertoin and Kortchemski [18] set up results similar to Theorem 4.4 for *non-monotone* Markov chains and develop several applications, to random walks conditioned to stay positive, to the number of particles in some coagulation-fragmentations processes, to random planar maps (see [20] for this last point). Also in [61] similar convergences for bivariate Markov chains towards time-changed Markov additive processes are studied. This will have applications to dynamical models of tree growth in a broader context than the one presented in Sect. 5.3, and more generally to multi-type MB-trees.

4.3 Self-Similar Fragmentation Trees

Self-similar fragmentation trees are random compact measured real trees that describe the genealogical structure of self-similar fragmentation processes with a negative index. It turns out that this set of trees is closely related to the set of trees arising as scaling limits of MB-trees. We start by introducing the self-similar fragmentation processes, following Bertoin [14], and then turn to the description of their genealogical trees, which were first introduced in [57] and then in [91] in a broader context.

4.3.1 Self-Similar Fragmentation Processes

Fragmentation processes are continuous-time processes that describe the evolution of an object that splits repeatedly and randomly as time passes. In the models we are interested in, the fragments are characterized by their mass alone, other characteristics, such as their shape, do not come into account. Many researchers have been working on such models. From a historical perspective, it seems that Kolmogorov [68] was the first in 1941. Since the early 2000s, there has been a full treatment of fragmentation processes satisfying a self-similarity property. We refer to Bertoin's book [14] for an overview of work in this area and a deepening of the results presented here.

We will work on the space of masses

$$\mathcal{S}_-^\downarrow = \left\{ \mathbf{s} = (s_1, s_2, \dots) : s_1 \geq s_2 \geq \dots \geq 0 \text{ and } \sum_{i \geq 1} s_i \leq 1 \right\},$$

which contains the set \mathcal{S}^\downarrow , and which is equipped with the same metric $d_{\mathcal{S}^\downarrow}$.

Definition 4.6 Let $\alpha \in \mathbb{R}$. An α -self-similar fragmentation process is an \mathcal{S}_-^\downarrow -valued Markov process $(F(t), t \geq 0)$ which is continuous in probability and such that, for all $t_0 \geq 0$, given that $F(t_0) = (s_1, s_2, \dots)$, the process $(F(t_0 + t), t \geq 0)$ is distributed as the process G obtained by considering a sequence $(F^{(i)}, i \geq 1)$ of

i.i.d. copies of F and then defining $G(t)$ to be the decreasing rearrangement of the sequences $s_i F^{(i)}(s_i^\alpha t)$, $i \geq 1$, for all $t \geq 0$.

In the following, we will always consider processes starting from a unique mass equal to 1, i.e. $F(0) = (1, 0, \dots)$. At time t , the sequence $F(t)$ should be understood as the decreasing sequence of the masses of fragments present at that time.

It turns out that such processes indeed exist and that their distributions are characterized by three parameters: the *index of self-similarity* $\alpha \in \mathbb{R}$, an *erosion coefficient* $c \geq 0$ that codes a continuous melt of the fragment (when $c = 0$ there is no erosion) and a *dislocation measure* ν , which is a measure ν on \mathcal{S}_-^\downarrow such that $\int_{\mathcal{S}_-^\downarrow} (1 - s_1) \nu(ds) < \infty$. The role of the parameters α and ν can be specified as follows when $c = 0$ and ν is finite: then, each fragment with mass m waits a random time with exponential distribution with parameter $\nu(\mathcal{S}_-^\downarrow)$ and then splits in fragments with masses $m\mathbf{S}$, where \mathbf{S} is distributed according to $\nu/\nu(\mathcal{S}_-^\downarrow)$, independently of the splitting time. When ν is infinite, the fragments split immediately, see [14, Chapter 3] for further details.

The index α has an enormous influence on the behavior of the process: when $\alpha = 0$, all fragments split at the same rate, whereas when $\alpha > 0$ fragments with small masses split slower and when $\alpha < 0$ fragments with small masses split faster. In this last case the fragments split so quickly that the whole initial object is reduced to “dust” in finite time, almost surely, i.e. $\inf\{t \geq 0 : F(t) = (0, \dots)\} < \infty$ a.s.

The Tagged Fragment Process We turn to a connection with the results seen in the previous section. Pick a point uniformly at random in the initial object (this object can be seen as an interval of length 1, for example), independently of the evolution of the process and let $F_*(t)$ be the mass of the fragment containing this marked point at time t . The process $F_*(t)$ is non-increasing and more precisely,

Theorem 4.7 (Bertoin [14], Theorem 3.2 and Corollary 3.1) *The process F_* can be written as*

$$F_*(t) = \exp(-\xi_{\tau(t)}), \quad \forall t \geq 0,$$

where ξ is a subordinator with Laplace exponent

$$\phi(\lambda) = c + \int_{\mathcal{S}_-^\downarrow} \left(1 - \sum_{i \geq 1} s_i\right) \nu(ds) + c\lambda + \int_{\mathcal{S}_-^\downarrow} \sum_{i \geq 1} (1 - s_i^\lambda) s_i \nu(ds), \quad \lambda \geq 0$$

and τ is a time-change depending on the parameter α , $\tau(t) = \inf\{u \geq 0 : \int_0^u \exp(\alpha \xi_r) dr > t\}$.

4.3.2 Self-Similar Fragmentation Trees

It was shown in [57] that to every self-similar fragmentation process with a negative index $\alpha = -\gamma < 0$, no erosion ($c = 0$) and a dislocation measure ν satisfying $\nu(\sum_{i \geq 1} s_i < 1) = 0$ (we say that ν is *conservative*), there is an associated compact rooted measured tree that describes its genealogy. We denote such a tree by $(\mathcal{T}_{\gamma, \nu}, \mu_{\gamma, \nu})$ and precise that the measure $\mu_{\gamma, \nu}$ is fully supported by the set of leaves of $\mathcal{T}_{\gamma, \nu}$ and non-atomic. In [91], Stephenson more generally constructed and studied compact rooted measured trees that describe the genealogy of any self-similar fragmentation process with a negative index. However in this survey, we restrict ourselves to the family of trees $(\mathcal{T}_{\gamma, \nu}, \mu_{\gamma, \nu})$, with $\gamma > 0$ and ν conservative.

The connection between a tree $(\mathcal{T}_{\gamma, \nu}, \mu_{\gamma, \nu})$ and the fragmentation process it is related to can be summarized as follows: for all $t \geq 0$, consider the connected components of $\{v \in \mathcal{T}_{\gamma, \nu} : \text{ht}(v) > t\}$, the set of points in $\mathcal{T}_{\gamma, \nu}$ that have a height strictly larger than t , and let $F(t)$ denote the decreasing rearrangement of the $\mu_{\gamma, \nu}$ -masses of these components. Then F is a fragmentation process, with index of self-similarity $-\gamma$, dislocation measure ν and no erosion. Besides, we note that $\mathcal{T}_{\gamma, \nu}$ possesses a fractal property, in the sense that if we fix a $t \geq 0$ (deterministic) and consider a point x at height t , then any subtree of $\mathcal{T}_{\gamma, \nu}$ descending from this point x (i.e. any connected component of $\{v \in \mathcal{T}_{\gamma, \nu} : x \in [[\rho, v]]\}$), having, say, a $\mu_{\gamma, \nu}$ -mass m , is distributed as $m^\gamma \mathcal{T}_{\gamma, \nu}$.

First Examples The Brownian tree and the α -stable trees that arise as scaling limits of Galton–Watson trees all belong to the family of self-similar fragmentation trees. More precisely,

- Bertoin [13] notices that the Brownian tree $(\mathcal{T}_{\text{Br}}, \mu_{\text{Br}})$ is a self-similar fragmentation tree and calculates its characteristics: $\gamma = 1/2$ and $\nu_{\text{Br}}(s_1 + s_2 < 1) = 0$ and

$$\nu_{\text{Br}}(s_1 \in dx) = \frac{\sqrt{2}}{\sqrt{\pi} x^{3/2} (1-x)^{3/2}}, \quad 1/2 < x < 1.$$

The fact that $\nu_{\text{Br}}(s_1 + s_2 < 1) = 0$ corresponds to the fact the tree is binary: every branch point has two descendants trees, and in the corresponding fragmentation, every splitting events gives two fragments.

- Miermont [75] proves that each stable tree \mathcal{T}_α is self-similar and calculates its characteristics when $\alpha \in (1, 2)$: $\gamma = 1 - 1/\alpha$ and

$$\int_{\mathcal{S}^\downarrow} f(\mathbf{s}) \nu_\alpha(d\mathbf{s}) = C_\alpha \mathbb{E} \left[T_1 f \left(\frac{\Xi_i}{T_1}, i \geq 1 \right) \right],$$

where

$$C_\alpha = \frac{\alpha(\alpha - 1)\Gamma(1 - 1/\alpha)}{\Gamma(2 - \alpha)}$$

and $(\Xi_i, i \geq 1)$ is the sequence of lengths, ranked in decreasing order, of intervals between successive atoms of a Poisson measure on \mathbb{R}_+ with intensity $(\alpha\Gamma(1 - 1/\alpha))^{-1}dr/r^{1+1/\alpha}$, and $T_1 = \sum_i \Xi_i$.

Hausdorff Dimension We first quickly recall the definition of Hausdorff dimension, which is a quantity that measures the “size” of metric spaces. We refer to the book of Falconer [53] for more details on that topic and for an introduction to fractal geometry in general. For all $r > 0$, the r -dimensional Hausdorff measure of a metric space (Z, d_Z) is defined by

$$\mathcal{M}^r(Z) := \lim_{\varepsilon \rightarrow 0} \inf_{\{(C_i)_{i \in \mathbb{N}} : \text{diam} C_i \leq \varepsilon\}} \left\{ \sum_{i \geq 1} \text{diam}(C_i)^r : Z \subset \cup_{i \geq 1} C_i \right\}.$$

where the infimum is taken over all coverings of Z by countable families of subsets $C_i \subset Z, i \in \mathbb{N}$, all with a diameter smaller than ε . The function $r > 0 \mapsto \mathcal{M}^r(Z) \in [0, \infty]$ is finite, non-zero at most one point. The Hausdorff dimension of Z is then given by

$$\dim_{\text{H}}(Z) = \inf \{r > 0 : \mathcal{M}^r(Z) = 0\} = \sup \{r > 0 : \mathcal{M}^r(Z) = \infty\}.$$

The Hausdorff dimension of a fragmentation tree depends mainly on its index of self-similarity. Let $\mathcal{L}(\mathcal{T}_{\gamma,v})$ denote the set of leaves of $\mathcal{T}_{\gamma,v}$. Then we know that,

Theorem 4.8 ([57]) *If $\int_{\mathcal{S}^1} (s_1^{-1} - 1)v(ds) < \infty$, then almost surely*

$$\dim_{\text{H}}(\mathcal{L}(\mathcal{T}_{\gamma,v})) = \frac{1}{\gamma} \quad \text{and} \quad \dim_{\text{H}}(\mathcal{T}_{\gamma,v}) = \max\left(\frac{1}{\gamma}, 1\right).$$

In particular, the Hausdorff dimension of the Brownian tree is 2, and more generally the Hausdorff dimension of the α -stable tree, $\alpha \in (1, 2)$, is $\alpha/(\alpha - 1)$. This recovers a result of Duquesne and Le Gall [47] proved in the framework of Lévy trees (we note that the intersection between the set of Lévy trees and that of self-similar fragmentation trees is exactly the set of stable Lévy trees).

Height of a Typical Leaf The measured tree $(\mathcal{T}_{\gamma,v}, \mu_{\gamma,v})$ has been constructed in such a way that if we pick a leaf L at random in $\mathcal{T}_{\gamma,v}$ according to $\mu_{\gamma,v}$ and we consider for each $t \geq 0$ the $\mu_{\gamma,v}$ -mass of the connected component of $\{v \in \mathcal{T}_{\gamma,v} : \text{ht}(v) > t\}$ that contains this marked leaf (with the convention that this mass is 0 if $\text{ht}(L) \leq t$), then we obtain a process which is distributed as the tagged fragment of the corresponding fragmentation process, as defined in the previous section. In particular, the height of L is distributed as the absorption time of the process $\exp(-\xi_t)$ introduced in Theorem 4.7, i.e.

$$\text{ht}(L) \stackrel{(d)}{=} \int_0^\infty \exp(-\gamma \xi_r) dr \tag{9}$$

where ξ is a subordinator with Laplace exponent $\phi(\lambda) = \int_{\mathcal{S}^\downarrow} \sum_{i \geq 1} (1 - s_i^\lambda) s_i \nu(ds)$, $\lambda \geq 0$. This is exactly the distribution of the limit appearing in Corollary 4.5.

4.4 Scaling Limits of Markov-Branching Trees

We can now explain the main steps of the proof of Theorem 4.1. In that aim, let $(T_n, n \geq 1)$ denote an MB-sequence indexed by the leaves with splitting probabilities satisfying (H), with parameters (γ, ν) in the limit. We actually only give here a hint of the proof of the convergence of the rescaled trees and refer to [59, Section 4.4] to see how to incorporate the measures. The proof of the convergence of the rescaled trees consists in three main steps:

First Step: Convergence of the Height of a Typical Leaf Keeping the notations previously introduced, $\text{ht}(\star_n)$ for the height of a typical leaf in T_n and $\text{ht}(L)$ for the height of a typical leaf in a fragmentation tree $(\mathcal{T}_{\gamma, \nu}, \mu_{\gamma, \nu})$, we get by (9) and Corollary 4.5 that

$$\frac{\text{ht}(\star_n)}{n^\gamma} \xrightarrow[n \rightarrow \infty]{(d)} \text{ht}(L).$$

Second Step: Convergence of Finite-Dimensional Marginals For all integers $k \geq 2$, let $T_n(k)$ be the subtree of T_n spanned by the root and k (different) leaves picked independently, uniformly at random. Similarly, let $\mathcal{T}_{\gamma, \nu}(k)$ be the subtree of $\mathcal{T}_{\gamma, \nu}$ spanned by the root and k leaves picked independently at random according to the measure $\mu_{\gamma, \nu}$. Then (under (H)),

$$\frac{T_n(k)}{n^\gamma} \xrightarrow[n \rightarrow \infty]{(d)} \mathcal{T}_{\gamma, \nu}(k). \quad (10)$$

This can be proved by induction on k . For $k = 1$, this is Step 1 above. For $k \geq 2$, we use the induction hypothesis and the MB-property. Here is the main idea. Consider the decomposition of T_n into subtrees above its first branch point in $T_n(k)$ and take only into account the subtrees having marked leaves. We obtain $m \geq 2$ subtrees with, say, n_1, \dots, n_m leaves respectively ($\sum_{i=1}^m n_i \leq n$), and each of these trees have $k_1 \geq 1, \dots, k_m \geq 1$ marked leaves ($\sum_{i=1}^m k_i = k$). Given $m, n_1, \dots, n_m, k_1, \dots, k_m$, the MB-property ensures that the m subtrees are independent with respective distributions that of $T_{n_1}(k_1), \dots, T_{n_m}(k_m)$. An application of the induction hypothesis to these subtrees leads to the expected result. We refer to [59, Section 4.2] for details.

Third Step: A Tightness Criterion To get the convergence for the GH-topology, the previous result must be completed with a tightness criterion. The idea is to use the following well known result.

Theorem 4.9 ([23], Theorem 3.2) *If $X_n, X, X_n(k), X(k)$ are r.v. in a metric space (E, d) such that $X_n(k) \xrightarrow[n \rightarrow \infty]{(d)} X(k), \forall k$ and $X(k) \xrightarrow[k \rightarrow \infty]{(d)} X$ and for all $\varepsilon > 0$,*

$$\lim_{k \rightarrow \infty} \limsup_{n \rightarrow \infty} \mathbb{P}(d(X_n, X_n(k)) > \varepsilon) = 0 \quad (11)$$

then $X_n \xrightarrow[n \rightarrow \infty]{(d)} X$.

In our context, the finite-dimensional convergence (10) has already been checked. Moreover, since $\mu_{\gamma, v}$ is fully supported on the set of leaves of $\mathcal{T}_{\gamma, v}$, we see by picking an infinite sequence of i.i.d. leaves according to $\mu_{\gamma, v}$, that there exist versions of the $\mathcal{T}_{\gamma, v}(k), k \geq 1$ that converge almost surely to $\mathcal{T}_{\gamma, v}$ as $k \rightarrow \infty$. It remains to establish the tightness criterion (11) for $T_n, T_n(k)$, with respect to the distance d_{GH} . The main tool is the following bounds:

Proposition 4.10 *Under (H), for all $p > 0$, there exists a finite constant C_p such that*

$$\mathbb{P}\left(\frac{\text{ht}(T_n)}{n^\gamma} \geq x\right) \leq \frac{C_p}{x^p}, \quad \forall x > 0, \forall n \geq 1.$$

The proof holds by induction on n , using (H) and the MB-property. We refer to [59, Section 4.3] for details and to see how, using again the MB-property, this helps to control the distance between T_n and $T_n(k)$, to get that for $\varepsilon > 0$:

$$\lim_{k \rightarrow \infty} \limsup_{n \rightarrow \infty} \mathbb{P}\left(d_{\text{GH}}\left(\frac{T_n(k)}{n^\gamma}, \frac{T_n}{n^\gamma}\right) \geq \varepsilon\right) = 0$$

as required.

5 Applications

We now turn to the description of the scaling limits of various models of random trees that are closely linked to the MB-property.

5.1 Galton–Watson Trees

5.1.1 Galton–Watson Trees with n Vertices

A first application of Theorem 4.2 is that it permits to recover the classical results of Aldous and Duquesne (grouped together in Theorem 3.1) on the scaling limits

of Galton–Watson trees conditioned to have n vertices. To see this, one just has to check the two following lemmas, for η a critical offspring distribution, $\eta(1) \neq 1$, and $(p_n^{\text{GW},\eta})$ the associated splitting distributions defined in (3).

Lemma 5.1 *If η has a finite variance σ^2 , then $(p_n^{\text{GW},\eta})$ satisfies (H) with*

$$\gamma = 1/2 \quad \text{and} \quad \nu = \frac{\sigma}{2} \nu_{\text{Br}}.$$

Lemma 5.2 *If $\eta(k) \sim \kappa k^{-\alpha-1}$ for some $\alpha \in (1, 2)$, then $(p_n^{\text{GW},\eta})$ satisfies (H) with*

$$\gamma = 1 - 1/\alpha \quad \text{and} \quad \nu = \left(\kappa \Gamma(2 - \alpha) \alpha^{-1} (\alpha - 1)^{-1} \right)^{1/\alpha} \nu_\alpha.$$

The measures ν_{Br} and ν_α are the dislocation measures of the Brownian and α -stable tree, respectively, and are defined in Sect. 4.3.2. Together with the scaling limit results on MB-trees, this gives Theorem 3.1 (i) and (ii) respectively. The proofs of these lemmas are not completely obvious. We give here a rough idea of the main steps of the proof of Lemma 5.1, and refer to [59, Section 5] for more details and for the proof of Lemma 5.2.

Sketch of the Main Steps of the Proof of Lemma 5.1 Recall that T^η denotes a Galton–Watson tree with offspring distribution η . To simplify, we assume that the support of η generates \mathbb{Z} , so that $\mathbb{P}(\#\text{vertices } T^\eta = n) > 0$ for all n large enough. The Otter–Dwass formula (or cyclic lemma) [86, Chapter 6] then implies that

$$\mathbb{P}(\#\text{vertices } T^\eta = n) = n^{-1} \mathbb{P}(S_n = -1)$$

where S_n is a random walk with i.i.d. increments of law $(\eta(i+1), i \geq -1)$. Together with the local limit Theorem, which ensures that $\mathbb{P}(S_n = -1) \underset{n \rightarrow \infty}{\sim} (2\pi\sigma^2n)^{-1/2}$, this leads to

$$\mathbb{P}(\#\text{vertices } T^\eta = n) \underset{n \rightarrow \infty}{\sim} (2\pi\sigma^2)^{-1/2} n^{-3/2}.$$

(We note that this argument is also fundamental in the study of large Galton–Watson trees via their contour functions). Then the idea is to use this approximation in the definition of $p_n^{\text{GW},\eta}$ to show that the two sums

$$\begin{aligned} & \sqrt{n} \sum_{\lambda \in \mathcal{P}_n} p_n^{\text{GW},\eta}(\lambda) \left(1 - \frac{\lambda_1}{n}\right) f\left(\frac{\lambda}{n}\right) \quad \text{and} \\ & \frac{\sigma}{\sqrt{2\pi}} \frac{1}{n} \sum_{\lambda_1 = \lceil n/2 \rceil}^n f\left(\frac{\lambda_1}{n}, \frac{n - \lambda_1}{n}, \dots\right) \left(\frac{\lambda_1}{n}\right)^{-3/2} \left(\frac{n - \lambda_1}{n}\right)^{-3/2} \end{aligned}$$

are asymptotically equivalent (this is the technical part), for all continuous functions $f : \mathcal{S}^\downarrow \rightarrow \mathbb{R}$. The conclusion follows, since the second sum is a Riemann sum that converges to $\sigma(\sqrt{2\pi})^{-1} \int_{1/2}^1 f(x, 1-x, \dots)x^{-3/2}(1-x)^{-3/2}dx$.

5.1.2 Galton–Watson Trees with Arbitrary Degree Constraints

One may then naturally wonder if Theorem 4.1 could also be used to get the scaling limits of Galton–Watson trees conditioned to have n leaves. The answer is yes, and moreover this can be done in a larger context, using a simple generalization of Theorems 4.1 and 4.2 to MB-trees with arbitrary degree constraints. This generalization was done by Rizzolo [89] using an idea similar to the one presented below Theorem 4.2 to get this theorem from Theorem 4.1. It is quite heavy to state neatly, so we let the reader see the paper [89] and focus here on the applications developed in this paper to Galton–Watson trees.

The classical theorems of Aldous and Duquesne on conditioned Galton–Watson trees can be extended to Galton–Watson trees conditioned to have a number of vertices with out-degree in a given set. To be more precise, fix $A \subset \mathbb{Z}_+$ and consider an offspring distribution η with mean 1 and variance $0 < \sigma^2 < \infty$. For integers n for which such a conditioning is possible, let $T_n^{\eta,A}$ denote a version of a η -Galton–Watson tree conditioned to have exactly n vertices with out-degree in A . For example, if $A = \mathbb{Z}_+$, this is the model of the previous section, whereas if $A = \{0\}$, $T_n^{\eta,A}$ is a η -Galton–Watson tree conditioned to have n leaves.

Theorem 5.3 (Kortchemski [69] and Rizzolo [89]) *As $n \rightarrow \infty$,*

$$\left(\frac{T_n^{\eta,A}}{\sqrt{n}}, \mu_n^{\eta,A} \right) \xrightarrow[\text{GHP}]{(d)} \left(\frac{2}{\sigma \sqrt{\eta(A)}} \mathcal{T}_{\text{Br}}, \mu_{\text{Br}} \right).$$

The proof of Rizzolo [89] relies on his theorem on scaling limits of MB-trees with arbitrary degree constraints. The most technical part is to evaluate the splitting probabilities, which is done by generalizing the Otter-Dwass formula. The proof of Kortchemski [69] is more in the spirit of the proofs of Aldous and Duquesne and consists in studying the contour functions of the conditioned trees. We note that [69] also includes cases where η has an infinite variance, and is in the domain of attraction of a stable distribution (the limit is then a multiple of a stable tree). It should be possible to recover this more general case via the approach of Rizzolo.

An Example of Application to Combinatorial Trees Indexed by the Number of Leaves Let T_n be a tree uniformly distributed amongst the set of rooted ordered trees with n leaves with no vertex with out-degree 1. One checks, using (2), that T_n is distributed as an η -Galton–Watson tree conditioned to have n leaves, with η defined by $\eta(i) = (1 - 2^{-1/2})^{i-1}, i \geq 2, \eta(1) = 0$ and $\eta(0) = 2 - 2^{1/2}$. The

variance of η is $4(\sqrt{2} - 1)$, so that finally,

$$\left(\frac{T_n}{\sqrt{n}}, \mu_n \right) \xrightarrow[\text{GHP}]{(d)} \left(\frac{1}{2^{1/4}(\sqrt{2} - 1)} \mathcal{T}_{\text{Br}}, \mu_{\text{Br}} \right)$$

where μ_n is the uniform probability on the leaves of T_n .

5.2 Pólya Trees

The above results on conditioned Galton–Watson trees give the scaling limits of several sequences of combinatorial trees, as already mentioned. There is however a significant case which does not fall within the Galton–Watson framework, that of *uniform Pólya trees*. By Pólya trees we simply mean rooted finite trees (non-ordered, non-labelled). They are named after Pólya [87] who developed an analytical treatment of this family of trees, based on generating functions. In this section, we let $T_n(\mathbb{P})$ be uniformly distributed amongst the set of Pólya trees with n vertices.

These trees are more complicated to study than uniform rooted trees with labelled vertices, or uniform rooted, ordered trees, because of their lack of symmetry. In this direction, Drmota and Gittenberger [43] showed that the shape of $T_n(\mathbb{P})$ is *not* a conditioned Galton–Watson tree. However Aldous [9] conjectured in 1991 that the scaling limit of $(T_n(\mathbb{P}))$ should nevertheless be the Brownian tree, up to a multiplicative constant. Quite recently, several papers studied the scaling limits of Pólya trees, with different points of view. Using techniques of analytic combinatorics, Broutin and Flajolet [26] studied the scaling limit of the height of a uniform *binary* Pólya tree with n vertices, whereas Drmota and Gittenberger [43] studied the profile of $T_n(\mathbb{P})$ (the profile is the sequence of the sizes of each generation of the tree) and showed that it converges after an appropriate rescaling to the local time of a Brownian excursion. Marckert and Miermont [74] obtained a full scaling limit picture of uniform *binary* Pólya trees: by appropriate trimming procedures, they showed that rescaled by \sqrt{n} , they converge in distribution towards a multiple of the Brownian tree.

More recently, with different methods, the following result was proved.

Theorem 5.4 (Haas–Miermont [59] and Panagiotou–Stufler [81]) *As $n \rightarrow \infty$,*

$$\left(\frac{T_n(\mathbb{P})}{\sqrt{n}}, \mu_n(\mathbb{P}) \right) \xrightarrow[\text{GHP}]{(d)} (c_{\mathbb{P}} \mathcal{T}_{\text{Br}}, \mu_{\text{Br}}), \quad c_{\mathbb{P}} \sim 1.491$$

where $\mu_n(\mathbb{P})$ denotes the uniform probability on the vertices of $T_n(\mathbb{P})$.

The proof of [59] uses connections with MB-trees, whereas that of [81] uses, still, connections with Galton–Watson trees. Let us first quickly discuss the MB point of view. It is easy to check that the sequence $(T_n(\mathbb{P}))$ is not Markov-Branching (this is left as an exercise!), however it is not far from being so. It is actually possible to

couple this sequence with a Markov-Branching sequence $(T'_n(\mathbb{P}))$ such that

$$\mathbb{E} \left[d_{\text{GHP}}(n^{-\varepsilon} T_n(\mathbb{P}), n^{-\varepsilon} T'_n(\mathbb{P})) \right] \xrightarrow{n \rightarrow \infty} 0, \quad \forall \varepsilon > 0$$

and $(T_n(\mathbb{P}))$ and $(T'_n(\mathbb{P}))$ have the same splitting probabilities (p_n) (by splitting probabilities for trees that are not MB, we mean the distribution of the sizes of the subtrees above the root). These splitting probabilities are given here by

$$p_{n-1}(\lambda) = \frac{\prod_{j=1}^{n-1} \#F_j(m_j(\lambda))}{\#\mathbb{T}_n}, \quad \text{for } \lambda \in \mathcal{P}_{n-1}$$

where $m_j(\lambda) = \{i : \lambda_i = j\}$, $\#\mathbb{T}_n$ is the number of rooted trees with n vertices and $F_j(k)$ denotes the set of multisets with k elements in \mathbb{T}_j (with the convention $F_j(0) := \{\emptyset\}$). It remains to check that these splitting probabilities satisfy (H) with appropriate parameters and to do this, we use the result of Otter [79]:

$$\#\mathbb{T}_n \underset{n \rightarrow \infty}{\sim} c \frac{\kappa^n}{n^{3/2}}, \quad \text{for some } c > 0, \kappa > 1.$$

Very roughly, this allows to conclude that the two following sums

$$\begin{aligned} & \sqrt{n} \sum_{\lambda \in \mathcal{P}_n} p_n(\lambda) \left(1 - \frac{\lambda_1}{n} \right) f \left(\frac{\lambda}{n} \right) \quad \text{and} \\ & \frac{c}{n} \sum_{\lambda_1 = \lceil (n-1)/2 \rceil}^{n-1} f \left(\frac{\lambda_1}{n}, \frac{n - \lambda_1}{n}, 0, \dots \right) \left(\frac{\lambda_1}{n} \right)^{-3/2} \left(\frac{n - \lambda_1}{n} \right)^{-3/2} \end{aligned}$$

are asymptotically equivalent, so that finally, using that the second sum is a Riemann sum, (H) holds with parameters $\gamma = 1/2$ and $\nu = \nu_{\text{Br}}/c_P$, with $c_P = \sqrt{2}/(c\sqrt{\pi})$. The method of [81] is different. It consists in showing that asymptotically $T_n(\mathbb{P})$ can be seen as a large finite-variance critical Galton–Watson tree of random size concentrated around a constant times n on which small subtrees of size $O(\log(n))$ are attached. The conclusion then follows from the classical result by Aldous on scaling limits of Galton–Watson trees.

Both methods can be adapted to Pólya trees with other degree constraints. In [59] uniform Pólya trees with n vertices having out-degree in $\{0, m\}$ for some fixed integer m , or out-degree at most m , are considered. More generally, in [81], uniform Pólya trees with n vertices having out-degree in a fixed set A (containing at least 0 and an integer larger than 2) are studied. In all cases, the trees rescaled by \sqrt{n} converge in distribution towards a multiple of the Brownian tree.

Further Result To complete the picture on combinatorial trees asymptotics, we mention a recent result by Stufler on *unrooted* trees, that was conjectured by Aldous, but remained open for a while.

Theorem 5.5 (Stufler [92]) *Let $T_n^*(\mathbb{P})$ be uniform amongst the set of unrooted trees with n vertices (unordered, unlabelled). Then,*

$$\frac{T_n^*(\mathbb{P})}{\sqrt{n}} \xrightarrow[\text{GH}]{(d)} c_{\mathbb{P}} \mathcal{T}_{\text{Br}}$$

(with the same $c_{\mathbb{P}}$ as in Theorem 5.4).

The main idea to prove this scaling limit of unrooted uniform trees consists in using a decomposition due to Bodirsky et al. [25] to approximate $T_n^*(\mathbb{P})$ by uniform rooted Pólya trees and then use Theorem 5.4. This result more generally holds for unrooted trees with very general degree constraints.

5.3 Dynamical Models of Tree Growth

As mentioned in Sect. 2.2, the prototype example of Rémy’s algorithm ($T_n(\mathbb{R})$, $n \geq 1$) is strongly connected to Galton–Watson trees since the shape of $T_n(\mathbb{R})$ (to which has been subtracted the edge between the root and the first branch point) is distributed as the shape of a binary critical Galton–Watson tree conditioned to have $2n - 1$ vertices. This implies that

$$\left(\frac{T_n(\mathbb{R})}{\sqrt{n}}, \mu_n(\mathbb{R}) \right) \xrightarrow[\text{GHP}]{(d)} \left(2\sqrt{2} \mathcal{T}_{\text{Br}}, \mu_{\text{Br}} \right).$$

Similar scaling limits results actually extends to most of the tree-growth models seen in Sect. 2.2. To see this, it suffices to check that their splitting probabilities satisfy Hypothesis (H). Technically, this mainly relies on Stirling’s formula and/or balls in urns schemes. Note however that the convergence in distribution is not fully satisfactory in these cases, since the trees are recursively built on a same probability space, and we may hope to have convergence in a stronger sense. We will see below that this is indeed the case.

5.3.1 Ford’s Alpha Model

For Ford’s α -model, with $\alpha \in (0, 1)$, it is easy to check (see [63]) that the splitting probabilities $q_n^{\text{Ford}, \alpha}$ satisfy hypothesis (H) with $\gamma = \alpha$ and $\nu = \nu_{\text{Ford}, \alpha}$, where $\nu_{\text{Ford}, \alpha}$ is a binary measure on $\mathcal{S}^{\downarrow}(\nu_{\text{Ford}, \alpha}(s_1 + s_2 < 1) = 0)$ defined by

$$\nu_{\text{Ford}, \alpha}(s_1 \in dx) = \frac{\mathbb{1}_{\{1/2 \leq x \leq 1\}}}{\Gamma(1 - \alpha)} \left(\alpha(x(1 - x))^{-\alpha-1} + (2 - 4\alpha)(x(1 - x))^{-\alpha} \right) dx.$$

This, together with Theorem 4.2 leads for $\alpha \in (0, 1)$ to the convergence:

Theorem 5.6 ([63] and [59]) *For all $\alpha \in (0, 1)$,*

$$\left(\frac{T_n(\alpha)}{n^\alpha}, \mu_n(\alpha) \right) \xrightarrow[\text{GHP}]{(d)} (\mathcal{T}_{\alpha, \nu_{\text{Ford}, \alpha}}, \mu_{\alpha, \nu_{\text{Ford}, \alpha}}).$$

This result was actually first proved in [63], using the fact that the sequence $(T_n(\alpha))$ is Markov-Branching and consistent. Chen and Winkel [32] then improved this result by showing that the convergence holds in probability.

For $\alpha = 1/2$ (Rémy's algorithm), note that we recover the result obtained via the Galton–Watson approach. Note also that the case $\alpha = 1$ is not included in the hypotheses of the above theorem, however the trees $T_n(\alpha)$ are then deterministic (comb trees) and it is clear that they converge after rescaling by n to a segment of length 1, equipped with the Lebesgue measure. This tree is a *general fragmentation tree* as introduced by Stephenson [91], with pure erosion (and no dislocation). When $\alpha = 0$, we observe a different regime, the height of a typical leaf in the tree growth logarithmically, and there is no convergence in the GH-sense of the whole tree.

5.3.2 k -Ary Growing Trees

Observing the asymptotic behavior of the sequence of trees constructed via Rémy's algorithm, it is natural to wonder how this may change when deciding to branch at each step $k - 1$ branches on the pre-existing tree, instead of one. For this k -ary model, it was shown in [60] that q_n^k satisfies (H) with $\gamma = 1/k$ and $\nu = \nu_k$ where

$$\nu_k(\mathbf{ds}) = \frac{(k-1)!}{k(\Gamma(\frac{1}{k}))^{k-1}} \prod_{i=1}^k s_i^{-(1-1/k)} \left(\sum_{i=1}^k \frac{1}{1-s_i} \right) \mathbb{1}_{\{s_1 \geq s_2 \geq \dots \geq s_k\}} \mathbf{ds},$$

is supported on the simplex of dimension $k - 1$. Together with Theorem 4.1 this gives the limit in distribution of the sequence $(T_n(k), n \geq 1)$. Besides, using some connections with the Chinese Restaurant Processes of Dubins and Pitman (see [86, Chapter 3] for a definition) and more general urns schemes, it was shown that these models converge in probability (however this second approach did not give the distribution of the limiting tree.) Together, these two methods lead to:

Theorem 5.7 ([60]) *Let $\mu_n(k)$ be the uniform measure on the leaves of $T_n(k)$. Then,*

$$\left(\frac{T_n(k)}{n^{1/k}}, \mu_n(k) \right) \xrightarrow[\text{GHP}]{\mathbb{P}} (\mathcal{T}_k, \mu_k)$$

where (\mathcal{T}_k, μ_k) is a self-similar fragmentation tree, with index of self-similarity $1/k$ and dislocation measure ν_k .

Interestingly, using the approximation by discrete trees, it is possible to show that randomized versions of the limiting trees \mathcal{T}_k , $k \geq 2$ —note that \mathcal{T}_2 is the Brownian tree up to a scaling factor—can be embedded into each other so as to form an increasing (in k) sequence of trees [60, Section 5].

In Sect. 6.1 we will discuss a generalization of this model. Here, we glue at each step star-trees with $k - 1$ branches. However more general results are available when deciding to glue more general tree structures, with possibly a random number of leaves.

5.3.3 Marginals of Stable Trees

For $\beta \in (1, 2]$, the sequence $(T_n(\beta), n \geq 1)$ built by Marchal’s algorithm provides, for each n , a tree that is distributed as the shape of the subtree of the stable tree \mathcal{T}_β spanned by n leaves taken independently according to μ_β . Duquesne and Le Gall [46] showed that $T_n(\beta)$ is distributed as a Galton–Watson tree whose offspring distribution has probability generating function $z + \beta^{-1}(1 - z)^\beta$, conditioned to have n leaves. As so, it is not surprising that appropriately rescaled it should converge to the β -stable tree. Marchal [73] proved an almost-sure *finite-dimensional convergence*, whereas the results of [63] give the convergence in probability for the GHP-topology. Additional manipulations even lead to an almost-sure convergence for the GHP-topology:

Theorem 5.8 ([35]) *Let $\mu_n(\beta)$ be the uniform measure on the leaves of $T_n(\beta)$. Then*

$$\left(\frac{T_n(\beta)}{n^\beta}, \mu_n(\beta) \right) \xrightarrow[\text{GHP}]{\text{a.s.}} (\beta \mathcal{T}_\beta, \mu_\beta)$$

Using this convergence, it was shown in [35] that randomized versions of the stable trees \mathcal{T}_β , $1 < \beta \leq 2$ can be embedded into each other so as to form a decreasing (in β) sequence of trees.

To complete these results, we mention that Chen et al. [33] propose a model that interpolate between the α -model of Ford and Marchal’s recursive construction of the marginals of stable trees, and determine there scaling limits, relying on the results of [63].

5.4 Cut-Trees

The notion of the cut-trees was introduced in Example 5 of Sect. 2.3.

Cut-Tree of a Uniform Cayley Tree We use the notation of Example 5, Sect. 2.3 and let C_n be a uniform Cayley tree and T_n^{cut} its cut-tree. Relying essentially on Stirling’s formula, one gets that $q_n^{\text{Cut,Cayley}}$ satisfies (H) with $\gamma = 1/2$ and $\nu =$

$\nu_{\text{Br}}/2$, which shows that the rescaled cut-tree $T_n^{\text{cut}}/\sqrt{n}$ endowed with the uniform measure on its leaves converges in distribution to $(2\mathcal{T}_{\text{Br}}, \mu_{\text{Br}})$. This was noticed in [15] and used to determine the scaling limits of the number of steps needed to isolated by edges delation a fixed number of vertices in C_n . Actually, Bertoin and Miermont [19] improve this result by showing the joint convergence

Theorem 5.9 (Bertoin–Miermont [19])

$$\left(\frac{C_n}{\sqrt{n}}, \frac{T_n^{\text{cut}}}{\sqrt{n}} \right) \xrightarrow[\text{GHP}]{(d)} (2\mathcal{T}_{\text{Br}}, 2\overline{\mathcal{T}}_{\text{Br}})$$

where $\overline{\mathcal{T}}_{\text{Br}}$ is a tree constructed from \mathcal{T}_{Br} , that can be interpreted as its cut-tree, and that is distributed as \mathcal{T}_{Br} .

Bertoin and Miermont [19] actually more generally extend this result to cut-trees of Galton–Watson trees with a critical offspring distribution with finite variance. This in turn was generalized by Dieuleveut [41] to Galton–Watson trees with a critical offspring distribution in the domain of attraction of a stable law. See also [5, 28, 29] for related results.

Cut-Tree of a Uniform Recursive Tree On the other hand, note that $q_n^{\text{Cut,Recursive}}$ does not satisfy (H). However, Bertoin showed in [17] that in this case, the cut-tree T_n rescaled by $n/\ln(n)$ converges for the GHP-topology to a segment of length 1, equipped with the Lebesgue measure.

6 Further Perspectives

6.1 Multi-Type Markov-Branching Trees and Applications

It is possible to enrich trees with *types*, by deciding that each vertex of a tree carries a type, which is an element of a finite or countable set. This multi-type setting is often used in the context of branching processes, where individuals with different types may evolve differently, and had been widely studied. For the trees point of view, scaling limits of multi-type Galton–Watson trees conditioned to have a given number of vertices have been studied by Miermont [77] when both the set of types and the covariance matrix of the offspring distributions are finite, by Berzunza [21] when the set of types is finite with offspring distributions in the domain of attraction of a stable distribution and by de Raphélis [40] when the number of types is infinite, under a finite variance-type assumption. The Brownian and stable trees appear in the scaling limits.

One may more generally be interested in multi-type Markov-Branching trees, which are sequences of trees with vertices carrying types, where, roughly, the subtrees above the root are independent and with distributions that only depend on their size and on the type of their root. In a work in progress [62], results similar

to Theorems 4.1 and 4.2 are set up for multi-type MB-trees, when the set of types is finite. Interestingly, different regimes appear in the scaling limits (multi-type or standard fragmentation trees), according to whether the rate of type change is faster than or equal to or slower than the rate of macroscopic branchings.

This should lead to new proofs of the results obtained in [21, 77]. This should also lead to other interesting applications, in particular to dynamical models of tree growth. In these growing models one starts from a finite alphabet of trees and then glues recursively trees by choosing at each step one tree at random in the alphabet and grafting it uniformly on an edge of the pre-existing tree. This generalizes the k -ary construction studied in Sects. 2.2 and 5.3.2, and is connected to multi-type MB-trees. In this general setting multi-type fragmentation trees will appear in the scaling limits.

6.2 Local Limits

This survey deals with scaling limits of random trees. There is another classical way to consider limits of sequences of trees (or graphs), that of *local limits*. This approach is quite different and provides other information on the asymptotics of the trees (e.g. on the limiting behavior of the degrees of vertices). Roughly, a sequence of finite rooted trees (t_n) is said to converge locally to a limit t if for all $R > 0$, the restriction of t_n to a ball of radius R centered at the root converges to the restriction of t to a ball of radius R centered at the root. The trees are therefore not rescaled and the limit is still a discrete object.

For results on the local limits of random models related to the ones considered here, we refer to: Abraham and Delmas [1, 2] and the references therein for Galton–Watson trees, Stefánsson [90] for Ford’s α -model and Pagnard [80] for general MB-sequences and the study of the volume growth of their local limits. We also mention the related work by Broutin and Mailler [27] that uses local limits of some models of MB-trees to study asymptotics of And/Or trees, that code boolean functions.

6.3 Related Random Geometric Structures

The discrete trees form a subclass of graphs and are generally simpler to study. There exist however several models of graphs (that are not trees) whose asymptotic study can be conducted by using trees. Different approaches are possible and it is not our purpose to present them here. However we still give some references that are related to some models of trees presented here (in particular Galton–Watson trees) to the interested reader (the list is not exhaustive):

- on random graphs converging to the Brownian tree: [7, 22, 31, 39, 67, 82, 93]
- on random graphs converging to tree-like structures: [36–38]

- on the Erdős–Rényi random graph in the critical window and application to the minimum spanning tree of the complete graph: [4, 6]
- on random maps (which are strongly connected to labeled trees): [78] and all the references therein.

In most of these works the Brownian tree intervenes in the construction of the continuous limit.

Acknowledgements I warmly thank the organizers of the *XII Simposio de Probabilidad y Procesos Estocásticos*, held in Mérida, as well as my collaborators on the topic of Markov-Branching trees, Grégory Miermont, Jim Pitman, Robin Stephenson and Matthias Winkel. Many thanks also to the referee for a very careful reading of a first version of this manuscript that led to many improvements in the presentation.

This work was partially supported by the ANR GRAAL ANR-14-CE25-0014.

References

1. R. Abraham, J.-F. Delmas, Local limits of conditioned Galton-Watson trees: the condensation case. *Electron. J. Probab.* **19**, 1–29 (2014)
2. R. Abraham, J.-F. Delmas, Local limits of conditioned Galton-Watson trees: the infinite spine case. *Electron. J. Probab.* **19**, 1–19 (2014)
3. R. Abraham, J.-F. Delmas, P. Hoscheit, A note on the Gromov-Hausdorff-Prokhorov distance between (locally) compact metric measure spaces. *Electron. J. Probab.* **18**(14), 1–21 (2013)
4. L. Addario-Berry, N. Broutin, C. Goldschmidt, The continuum limit of critical random graphs. *Probab. Theory Relat. Fields* **152**, 367–406 (2012)
5. L. Addario-Berry, N. Broutin, C. Holmgren, Cutting down trees with a Markov chainsaw. *Ann. Appl. Probab.* **24**, 2297–2339 (2014)
6. L. Addario-Berry, N. Broutin, C. Goldschmidt, G. Miermont, The scaling limit of the minimum spanning tree of the complete graph. *Ann. Probab.* **45**(5), 3075–3144 (2017)
7. M. Albenque, J.-F. Marckert, Some families of increasing planar maps. *Electron. J. Probab.* **13**(56), 1624–1671 (2008)
8. D. Aldous, The continuum random tree. I. *Ann. Probab.* **19**, 1–28 (1991)
9. D. Aldous, The continuum random tree. II. An overview, in *Stochastic Analysis (Durham, 1990)*. London Mathematical Society Lecture Note Series, vol. 167 (Cambridge University Press, Cambridge, 1991), pp. 23–70
10. D. Aldous, The continuum random tree III. *Ann. Probab.* **21**, 248–289 (1993)
11. D. Aldous, Probability distributions on cladograms, in *Random Discrete Structures (Minneapolis, MN, 1993)*. The IMA Volumes in Mathematics and its Applications, vol. 76 (Springer, New York, 1996), pp. 1–18
12. J. Bertoin, *Lévy Processes*. Cambridge Tracts in Mathematics, vol. 121 (Cambridge University Press, Cambridge, 1996)
13. J. Bertoin, Self-similar fragmentations. *Ann. Inst. Henri Poincaré Probab. Stat.* **38**, 319–340 (2002)
14. J. Bertoin, *Random Fragmentation and Coagulation Processes*. Cambridge Studies in Advanced Mathematics, vol. 102 (Cambridge University Press, Cambridge, 2006)
15. J. Bertoin, Fires on trees. *Ann. Inst. Henri Poincaré Probab. Stat.* **48**, 909–921 (2012)
16. J. Bertoin, Sizes of the largest clusters for supercritical percolation on random recursive trees. *Random Struct. Algorithm* **44**, 29–44 (2014)
17. J. Bertoin, The cut-tree of large recursive trees. *Ann. Inst. Henri Poincaré Probab. Stat.* **51**, 478–488 (2015)

18. J. Bertoin, I. Kortchemski, Self-similar scaling limits of Markov chains on the positive integers. *Ann. Appl. Probab.* **26**, 2556–2595 (2016)
19. J. Bertoin, G. Miermont, The cut-tree of large Galton-Watson trees and the Brownian CRT. *Ann. Appl. Probab.* **23**, 1469–1493 (2013)
20. J. Bertoin, N. Curien, I. Kortchemski, Random planar maps & growth-fragmentations. *Ann. Probab.* **46**(1), 207–260 (2018)
21. G. Berzunza, On scaling limits of multitype Galton-Watson trees with possibly infinite variance. Preprint (2016). arXiv:1605.04810
22. J. Bettinelli, Scaling limit of random planar quadrangulations with a boundary. *Ann. Inst. Henri Poincaré Probab. Stat.* **51**, 432–477 (2015)
23. P. Billingsley, *Convergence of Probability Measures*. Wiley Series in Probability and Statistics: Probability and Statistics, 2nd edn. (Wiley, New York, 1999). A Wiley-Interscience Publication
24. N.H. Bingham, C.M. Goldie, J.L. Teugels, *Regular Variation*. Encyclopedia of Mathematics and its Applications, vol. 27 (Cambridge University Press, Cambridge, 1989)
25. M. Bodirsky, É. Fusy, M. Kang, S. Vigerske, Boltzmann samplers, Pólya theory, and cycle pointing. *SIAM J. Comput.* **40**, 721–769 (2011)
26. N. Broutin, P. Flajolet, The distribution of height and diameter in random non-plane binary trees. *Random Struct. Algorithm* **41**, 215–252 (2012)
27. N. Broutin, C. Mailler, And/or trees: a local limit point of view. *Random Struct. Algorithm* (to appear)
28. N. Broutin, M. Wang, Reversing the cut tree of the Brownian continuum random tree. *Electron. J. Probab.* **22**(80), 1–23 (2017)
29. N. Broutin, M. Wang, Cutting down p-trees and inhomogeneous continuum random trees. *Bernoulli*, 23(4A), 2380–2433 (2017)
30. N. Broutin, L. Devroye, E. McLeish, M. de la Salle, The height of increasing trees. *Random Struct. Algorithm* **32**, 494–518 (2008)
31. A. Caraceni, The scaling limit of random outerplanar maps. *Ann. Inst. Henri Poincaré Probab. Stat.* **52**, 1667–1686 (2016)
32. B. Chen, M. Winkel, Restricted exchangeable partitions and embedding of associated hierarchies in continuum random trees. *Ann. Inst. Henri Poincaré Probab. Stat.* **49**, 839–872 (2013)
33. B. Chen, D. Ford, M. Winkel, A new family of Markov branching trees: the alpha-gamma model. *Electron. J. Probab.* **14**(15), 400–430 (2009)
34. D. Croydon, B. Hambly, Spectral asymptotics for stable trees. *Electron. J. Probab.* **15**(57), 1772–1801 (2010)
35. N. Curien, B. Haas, The stable trees are nested. *Probab. Theory Relat. Fields* **157**(1), 847–883 (2013)
36. N. Curien, I. Kortchemski, Random stable looptrees. *Electron. J. Probab.* **19**(108), 1–35 (2014)
37. N. Curien, I. Kortchemski, Percolation on random triangulations and stable looptrees. *Probab. Theory Relat. Fields* **163**, 303–337 (2015)
38. N. Curien, T. Duquesne, I. Kortchemski, I. Manolescu, Scaling limits and influence of the seed graph in preferential attachment trees. *J. Éc. Polytech. Math.* **2**, 1–34 (2015)
39. N. Curien, B. Haas, I. Kortchemski, The CRT is the scaling limit of random dissections. *Random Struct. Algorithms* **47**, 304–327 (2015)
40. L. de Raphélis, Scaling limit of multitype Galton-Watson trees with infinitely many types. *Ann. Inst. Henri Poincaré Probab. Stat.* **53**, 200–225 (2017)
41. D. Dieuleveut, The vertex-cut-tree of Galton-Watson trees converging to a stable tree. *Ann. Appl. Probab.* **25**, 2215–2262 (2015)
42. M. Drmota, *Random Trees: An Interplay Between Combinatorics and Probability* (Springer, Vienna, 2009)
43. M. Drmota, B. Gittenberger, The shape of unlabeled rooted random trees. *Eur. J. Comb.* **31**, 2028–2063 (2010)
44. T. Duquesne, A limit theorem for the contour process of conditioned Galton-Watson trees. *Ann. Probab.* **31**, 996–1027 (2003)
45. T. Duquesne, The exact packing measure of Lévy trees. *Stoch. Process. Appl.* **122**, 968–1002 (2012)

46. T. Duquesne, J.-F. Le Gall, *Random Trees, Lévy Processes and Spatial Branching Processes*. Astérisque (Société Mathématique de France, Paris, 2002), pp. vi+147
47. T. Duquesne, J.-F. Le Gall, Probabilistic and fractal aspects of Lévy trees. *Probab. Theory Relat. Fields* **131**, 553–603 (2005)
48. T. Duquesne, J.-F. Le Gall, The Hausdorff measure of stable trees. *ALEA Lat. Am. J. Probab. Math. Stat.* **1**, 393–415 (2006)
49. T. Duquesne, J.-F. Le Gall, On the re-rooting invariance property of Lévy trees. *Electron. Commun. Probab.* **14**, 317–326 (2009)
50. T. Duquesne, M. Winkel, Growth of Lévy trees. *Probab. Theory Relat. Fields* **139**, 313–371 (2007)
51. S.N. Evans, *Probability and Real Trees. Lectures from the 35th Summer School on Probability Theory Held in Saint-Flour, July 6–23, 2005*. Lecture Notes in Mathematics, vol. 1920 (Springer, Berlin, 2008)
52. S. Evans, J. Pitman, A. Winter, Rayleigh processes, real trees, and root growth with re-grafting. *Probab. Theory Relat. Fields* **134**(1), 918–961 (2006)
53. K. Falconer, *Fractal Geometry: Mathematical Foundations and Applications*, 3rd ed. (John Wiley & Sons, Chichester, 2014)
54. P. Flajolet, R. Sedgewick, *Analytic Combinatorics* (Cambridge University Press, Cambridge, 2009)
55. D. Ford, Probabilities on cladograms: introduction to the alpha model. Preprint (2005). arXiv:math/0511246
56. C. Goldschmidt, B. Haas, A line-breaking construction of the stable trees. *Electron. J. Probab.* **20**, 1–24 (2015)
57. B. Haas, G. Miermont, The genealogy of self-similar fragmentations with negative index as a continuum random tree. *Electron. J. Probab.* **9**(4), 57–97 (2004)
58. B. Haas, G. Miermont, Self-similar scaling limits of non-increasing Markov chains. *Bernoulli* **J. 17**, 1217–1247 (2011)
59. B. Haas, G. Miermont, Scaling limits of Markov branching trees with applications to Galton-Watson and random unordered trees. *Ann. Probab.* **40**, 2589–2666 (2012)
60. B. Haas, R. Stephenson, Scaling limits of k -ary growing trees. *Ann. Inst. Henri Poincaré Probab. Stat.* **51**, 1314–1341 (2015)
61. B. Haas, R. Stephenson, Bivariate Markov chains converging to Lamperti transform Markov Additive Processes. *Stoch. Process. Appl.* (to appear)
62. B. Haas, R. Stephenson, Multitype Markov Branching trees (In preparation)
63. B. Haas, G. Miermont, J. Pitman, M. Winkel, Continuum tree asymptotics of discrete fragmentations and applications to phylogenetic models. *Ann. Probab.* **36**, 1790–1837 (2008)
64. B. Haas, J. Pitman, M. Winkel, Spinal partitions and invariance under re-rooting of continuum random trees. *Ann. Probab.* **37**, 1381–1411 (2009)
65. S. Janson, Random cutting and records in deterministic and random trees. *Random Struct. Algorithms* **29**, 139–179 (2006)
66. S. Janson, Simply generated trees, conditioned Galton-Watson trees, random allocations and condensation: extended abstract, in *23rd International Meeting on Probabilistic, Combinatorial, and Asymptotic Methods for the Analysis of Algorithms (AofA'12). Discrete Mathematics and Theoretical Computer Science Proceedings, AQ* (Association of Discrete Mathematics and Theoretical Computer Science, Nancy, 2012), pp. 479–490
67. S. Janson, S.Ö. Stefánsson, Scaling limits of random planar maps with a unique large face. *Ann. Probab.* **43**, 1045–1081 (2015)
68. A. Kolmogorov, Über das logarithmisch normale Verteilungsgesetz der Dimensionen der Teilchen bei Zerstückelung. *C.R. Acad. Sci. U.R.S.S.* **31**, 99–101 (1941)
69. I. Kortchemski, Invariance principles for Galton-Watson trees conditioned on the number of leaves. *Stoch. Process. Appl.* **122**, 3126–3172 (2012)
70. I. Kortchemski, A simple proof of Duquesne’s theorem on contour processes of conditioned Galton-Watson trees, in *Séminaire de Probabilités XLV*. Lecture Notes in Mathematics, vol. 2078 (Springer, Cham, 2013), pp. 537–558

71. J.-F. Le Gall, Random real trees. *Ann. Fac. Sci. Toulouse Math.* (6) **15**, 35–62 (2006)
72. J.-F. Le Gall, Y. Le Jan, Branching processes in Lévy processes: the exploration process. *Ann. Probab.* **26**, 213–252 (1998)
73. P. Marchal, A note on the fragmentation of a stable tree, in *Fifth Colloquium on Mathematics and Computer Science. Discrete Mathematics and Theoretical Computer Science Proceedings, AI* (Association of Discrete Mathematics and Theoretical Computer Science, Nancy, 2008), pp. 489–499
74. J.-F. Marckert, G. Miermont, The CRT is the scaling limit of unordered binary trees. *Random Struct. Algorithms* **38**, 467–501 (2011)
75. G. Miermont, Self-similar fragmentations derived from the stable tree. I. Splitting at heights. *Probab. Theory Relat. Fields* **127**, 423–454 (2003)
76. G. Miermont, Self-similar fragmentations derived from the stable tree. II. Splitting at nodes. *Probab. Theory Relat. Fields* **131**, 341–375 (2005)
77. G. Miermont, Invariance principles for spatial multitype Galton-Watson trees. *Ann. Inst. Henri Poincaré Probab. Stat.* **44**, 1128–1161 (2008)
78. G. Miermont, *Aspects of Random Maps, Lecture Notes of the 2014 Saint-Flour Probability Summer School*. Preliminary draft: <http://perso.ens-lyon.fr/gregory.miermont/coursSaint-Flour.pdf>
79. R. Otter, The number of trees. *Ann. Math.* (2) **49**, 583–599 (1948)
80. C. Pagnard, Local limits of Markov-Branching trees and their volume growth. *Electron. J. Probab.* **22**(95), 1–53 (2017)
81. K. Panagiotou, B. Stuffer, Scaling limits of random Pólya trees. *Probab. Theory Relat. Fields* **170**, 801–820 (2018)
82. K. Panagiotou, B. Stuffer, K. Weller, Scaling limits of random graphs from subcritical classes. *Ann. Probab.* **44**, 3291–3334 (2016)
83. A. Panholzer, Cutting down very simple trees. *Quaest. Math.* **29**, 211–227 (2006)
84. J.L. Pavlov, The asymptotic distribution of the maximum size of trees in a random forest. *Teor. Veroyatnost. i Primenen.* **22**, 523–533 (1977)
85. J. Pitman, Coalescent random forests. *J. Combin. Theory Ser. A* **85**, 165–193 (1999)
86. J. Pitman, *Combinatorial Stochastic Processes. Lectures from the 32nd Summer School on Probability Theory Held in Saint-Flour, July 7–24, 2002*. *Lecture Notes in Mathematics*, vol. 1875 (Springer, Berlin, 2006)
87. G. Pólya, Kombinatorische Anzahlbestimmungen für Gruppen, Graphen und chemische Verbindungen. *Acta Math.* **68**, 145–254 (1937)
88. J.-L. Rémy, *Un procédé itératif de dénombrement d'arbres binaires et son application à leur génération aléatoire*. *RAIRO Inform. Théor.* **19**, 179–195 (1985)
89. D. Rizzolo, Scaling limits of Markov branching trees and Galton-Watson trees conditioned on the number of vertices with out-degree in a given set. *Ann. Inst. Henri Poincaré Probab. Stat.* **51**, 512–532 (2015)
90. S. Stefánsson, The infinite volume limit of Ford's alpha model. *Acta Phys. Pol. B Proc. Suppl.* **2**(3), 555–560 (2009)
91. R. Stephenson, General fragmentation trees. *Electron. J. Probab.* **18**(101), 1–45 (2013)
92. B. Stuffer, The continuum random tree is the scaling limit of unlabelled unrooted trees. Preprint (2016). arXiv:1412.6333
93. B. Stuffer, Random enriched trees with applications to random graphs. Preprint (2016). arXiv:1504.02006