

Finding Adam in random growing trees

Sébastien Bubeck ^{*} Luc Devroye [†] Gábor Lugosi [‡]

November 12, 2014

Abstract

We investigate algorithms to find the first vertex in large trees generated by either the uniform attachment or preferential attachment model. We require the algorithm to output a set of K vertices, such that, with probability at least $1 - \varepsilon$, the first vertex is in this set. We show that for any ε , there exist such algorithms with K independent of the size of the input tree. Moreover, we provide almost tight bounds for the best value of K as a function of ε . In the uniform attachment case we show that the optimal K is subpolynomial in $1/\varepsilon$, and that it has to be at least superpolylogarithmic. On the other hand, the preferential attachment case is exponentially harder, as we prove that the best K is polynomial in $1/\varepsilon$. We conclude the paper with several open problems.

1 Introduction

We consider one of the simplest models of a randomly growing graph: starting from a single node (referred to as the *root*), each arriving new node connects uniformly at random to one of the existing nodes. We are interested in the following question: given a large tree generated from this uniform attachment model, is it possible to find a small set of vertices for which we can certify with high probability that the root is in this set? Possible applications include finding the center of an epidemic, or the origin of a rumor.

In this paper we study *root-finding algorithms*: given a target accuracy $\varepsilon \in (0, 1)$ and a tree T of size n (for some $n \in \mathbb{N}$), a root-finding algorithm outputs a set $H(T, \varepsilon)$ of $K(\varepsilon)$ vertices, such that, with probability at least $1 - \varepsilon$ (with respect to the random generation of T from the uniform attachment model), the root is in $H(T, \varepsilon)$. An important aspect of the definition is that the size of the output set is allowed to depend on ε , but not on the size n of the input tree. Thus it is not obvious that root-finding algorithms exist at all. For instance a naive guess would be to output vertices of large degrees (indeed, the older a vertex, the larger is its expected degree), but for this to be correct with some constant probability one needs to output a logarithmic in n number of vertices. One of the main contributions of this paper is to show that root-finding algorithms indeed exist. Furthermore, we almost tightly characterize the best possible value for $K(\varepsilon)$, by showing that it can be subpolynomial in $1/\varepsilon$, and that it has to be superpolylogarithmic. More precisely the core of our contribution can be summarized by the following theorem.

^{*}Microsoft Research and Princeton University; sebubeck@microsoft.com.

[†]McGill University; lucdevroye@gmail.com.

[‡]ICREA and Pompeu Fabra University; gabor.lugosi@upf.edu. GL acknowledges support by the Spanish Ministry of Science and Technology grant MTM2012-37195.

Theorem 1 *There exist constants $c, c' > 0$ such that the following holds true in the uniform attachment model. Any root-finding algorithm must satisfy $K(\varepsilon) \geq \exp\left(c\sqrt{\log(1/\varepsilon)}\right)$. Furthermore, there exists a polynomial time root-finding algorithm with $K(\varepsilon) \leq \exp\left(c' \frac{\log(1/\varepsilon)}{\log \log(1/\varepsilon)}\right)$.*

We also investigate the existence of root-finding algorithms for the preferential attachment model, in which each arriving new node connects to an existing node with probability proportional to its degree. In this model the existence of root-finding algorithm is much less surprising, and, in fact, we show that considering vertices of large degrees works here. More interestingly, we prove that the preferential attachment model is *exponentially* more difficult than the uniform attachment model, in the sense that $K(\varepsilon)$ has to be at least polynomial in $1/\varepsilon$ (while it can be subpolynomial in uniform attachment). More precisely, we prove the following theorem.

Theorem 2 *There exist constants $c, c' > 0$ such that the following holds true in the preferential attachment model. Any root-finding algorithm must satisfy $K(\varepsilon) \geq \frac{c}{\varepsilon}$. Furthermore, there exists a polynomial time root-finding algorithm with $K(\varepsilon) \leq c' \frac{\log(1/\varepsilon)}{\varepsilon^2}$.*

1.1 Related work

There is a vast and rapidly growing literature on both uniform attachment and preferential attachment models. However, to the best of our knowledge, root-finding algorithms have not yet been investigated. On the other hand, a *local* variant of root-finding algorithms is studied for the preferential attachment model in Brautbar and Kearns [2010], Borgs et al. [2012], Frieze and Pegden [2014]. The setting in these papers is quite different from ours, as the algorithm only has a local access to the graph, and it is assumed that it can recognize the root once it attains it. In this context it is proved that the algorithm only has to visit a polylogarithmic (in n) number of vertices, which has to be contrasted with our condition that a root-finding algorithm has to output a set of size independent of n .

Another recent line of work intimately related to the question studied here is Bubeck et al. [2014b], Curien et al. [2014], Bubeck et al. [2014a]. In these papers the uniform attachment and preferential attachment models are initialized with some finite seed tree. It is then proved that different seeds lead to different distributions, even in the limit as n goes to infinity. In other words, the uniform attachment and preferential attachment trees are strongly influenced by their state after a finite number of steps, which gives hope that root-finding algorithms indeed exist¹.

1.2 Content of the paper

We start by presenting a simple root-finding algorithm for the uniform attachment model in Section 2. The idea is to rank vertices according to the size of their *largest subtree*. Using basic Pólya urn results, we show that by taking the smallest $K = \frac{\log(1/\varepsilon)}{\varepsilon}$ vertices (according to their largest subtree size), one obtains a root-finding algorithm.

We prove the impossibility result described in Theorem 1 in Section 3. Using a basic combinatorial argument we show that the optimal estimator for the root—the maximum likelihood estimator (MLE)—can be computed in polynomial time, and we use well-known results about the uniform attachment tree (such as the behavior of its height) to exhibit the limits of the MLE (and thus of any root-finding algorithm).

¹However, chronologically, the work presented here was done before Bubeck et al. [2014a] which proves the influence of the seed in the uniform attachment model.

In Section 4 we observe that the root-finding algorithm studied in Section 2 can be viewed as a relaxation of the MLE. We then propose a “tighter” relaxation which corresponds to the subpolynomial root-finding algorithm mentioned in Theorem 1. The analysis of this algorithm is the most technical part of the paper. It relies on a sharp concentration inequality for sums of Gamma random variables, as well as a beautiful result of Hardy and Ramanujan on the number of partitions of an integer, Hardy and Ramanujan [1918].

Observe that, since the MLE can be computed in polynomial-time, the relaxations of Section 2 and Section 4 are mainly introduced to simplify the analysis. However these relaxations also have computational advantages. While the worst-case complexity is of the same order of magnitude, the relaxed procedures allow for early stopping because the functions being minimized are unimodal over paths, which implies in particular that the set of vertices they output form a subtree of the original tree. On the other hand the MLE can output a set of disconnected vertices².

Finally the preferential attachment model is analyzed in Section 5, where Theorem 2 is proven. We conclude the paper with several open problems in Section 6.

1.3 Notation

For a labeled tree T we denote by T° the isomorphism class of T . In other words T° is an unlabeled copy of T . For notational convenience we denote vertices from T° using the labeling of T (formally one would need to chose an arbitrary labeling of T° , and to compare vertices of T° to those of T one would need to introduce the corresponding isomorphism). We denote by $V(T)$ the vertex set of a labeled tree, and again with a slight abuse of notation we extend this notation to unlabeled trees. The degree of a vertex v is denoted by $d_T(v)$, or simply by $d(v)$ if the underlying tree is clear from the context. An increasing labeling for a tree T with $V(T) = [n]$ is such that any path away from vertex 1 has increasing labels. A recursive tree is a tree labeled with an increasing labeling. A rooted tree (labeled or unlabeled) is denoted by (T, u) , with $u \in V(T)$. We sometimes denote a rooted tree simply as T , in which case we denote the root as \emptyset . In a recursive tree T it is understood that $\emptyset = 1$. In a rooted tree T we denote by $T_{v\downarrow}$ the subtree starting at v . In a rooted tree the descendants of a vertex are referred to as its children, and the set of vertices with no children (i.e., the leaves) is denoted $\mathcal{L}(T)$. A plane-oriented recursive tree is a recursive tree together with an ordering of the children of each vertex.

For $\alpha \in \mathbb{R}$ we define the random labeled tree $\mathcal{T}_\alpha(n)$ with vertex set $[n]$ by induction as follows: $\mathcal{T}_\alpha(2)$ is the unique tree on 2 vertices, and $\mathcal{T}_\alpha(n+1)$ is built from $\mathcal{T}_\alpha(n)$ by adding the vertex $n+1$ and an edge $\{i, n+1\}$, where $i \in [n]$ is chosen at random with probability proportional to the degrees in $\mathcal{T}_\alpha(n)$ raised to the power α . We focus on the case $\alpha = 0$, which we alternatively denote $\text{UA}(n)$ for *uniform attachment model*, and on $\alpha = 1$ which we denote $\text{PA}(n)$ for *preferential attachment model*. It is well known that $\text{UA}(n)$ is equivalently described as a uniformly chosen tree among all recursive trees, while $\text{PA}(n)$ can be described as a uniformly chosen tree among all plane-oriented trees, see e.g., Drmota [2009]. For this reason $\text{UA}(n)$ is also referred to as the uniform random recursive tree, and $\text{PA}(n)$ as the random plane-oriented recursive tree.

We can now formalize the problem introduced at the beginning of the introduction. Let $\alpha \in \{0, 1\}$, $\varepsilon \in (0, 1)$, and $K \in \mathbb{N}$. We are interested in mappings H from unlabeled trees to subsets of

²For instance, on a 5-vertex path, the two most likely vertices to be the root in the uniform attachment model are the two non-central vertices of degree 2.

K vertices with the property that

$$\liminf_{n \rightarrow +\infty} \mathbb{P}(1 \in H(\mathcal{T}_\alpha(n)^\circ)) \geq 1 - \varepsilon \quad (1)$$

In most instances the input tree will be clear from the context, and thus we often write H instead of $H(\mathcal{T}_\alpha(n)^\circ)$.

2 A simple root-finding algorithm

For a tree T we introduce the function $\psi_T : V(T) \rightarrow \mathbb{N}$, defined by

$$\psi_T(u) = \max_{v \in V(T) \setminus \{u\}} |(T, u)_{v\downarrow}|. \quad (2)$$

In words, viewing u as the root of T , $\psi_T(u)$ returns the size of the largest subtree starting at a child of u . We denote by H_ψ the mapping which returns the set of K vertices with smallest ψ values (ties are broken arbitrarily). The following theorem shows that H_ψ is a root-finding algorithm for the uniform attachment model. We also prove in Section 5 that it is a root-finding algorithm for the preferential attachment.

Theorem 3 *Let $K \geq 2.5 \frac{\log(1/\varepsilon)}{\varepsilon}$, then one has $\liminf_{n \rightarrow +\infty} \mathbb{P}(1 \in H_\psi(\text{UA}(n)^\circ)) \geq 1 - \frac{4\varepsilon}{1-\varepsilon}$.*

We observe that the theorem is optimal up to a logarithmic factor. Indeed, leaves clearly maximize the value of ψ and one can easily see that, with probability at least $\Omega(1/K)$, the true root is a leaf in $\text{UA}(K)$. Nonetheless, perhaps surprisingly, there is an exponentially better root-finding algorithm than H_ψ for the uniform attachment model.

Proof In this proof vertices are labeled by chronological order. We also introduce a notation which will prove useful; for $1 \leq i \leq k$ we denote by $T_{i,k}$ the tree containing vertex i in the forest obtained by removing in $\text{UA}(n)$ all edges between vertices $\{1, \dots, k\}$. In particular the vector $(|T_{1,k}|, \dots, |T_{k,k}|)$ follows a standard Pólya urn with k colors, and thus, using a classical result,

$$\frac{1}{n}(|T_{1,k}|, \dots, |T_{k,k}|)$$

converges, in distribution, to a Dirichlet distribution with parameters $(1, \dots, 1)$. Now observe first that

$$\mathbb{P}(1 \notin H_\psi) \leq \mathbb{P}(\exists i > K : \psi(i) \leq \psi(1)) \leq \mathbb{P}(\psi(1) \geq (1 - \varepsilon)n) + \mathbb{P}(\exists i > K : \psi(i) \leq (1 - \varepsilon)n).$$

Clearly

$$\psi(1) \leq \max(|T_{1,2}|, |T_{2,2}|),$$

and thus since $|T_{1,2}|/n$ and $|T_{2,2}|/n$ are identitically distributed and converge in distribution to a uniform random variable in $[0, 1]$,

$$\limsup_{n \rightarrow +\infty} \mathbb{P}(\psi(1) \geq (1 - \varepsilon)n) \leq 2 \lim_{n \rightarrow +\infty} \mathbb{P}(|T_{1,2}| \geq (1 - \varepsilon)n) = 2\varepsilon.$$

On the other hand, for any $i > K$,

$$\psi(i) \geq \min_{1 \leq k \leq K} \sum_{j=1, j \neq k}^K |T_{j,K}|,$$

and $\frac{1}{n} \sum_{j=1, j \neq k}^K |T_{j,K}|$ converges, in distribution, to the Beta($K - 1, 1$) distribution, which implies

$$\begin{aligned} \limsup_{n \rightarrow +\infty} \mathbb{P}(\exists i > K : \psi(i) \leq (1 - \varepsilon)n) &\leq \lim_{n \rightarrow +\infty} \mathbb{P} \left(\exists 1 \leq k \leq K : \sum_{j=1, j \neq k}^K |T_{j,K}| \leq (1 - \varepsilon)n \right) \\ &= K(1 - \varepsilon)^{K-1}. \end{aligned}$$

Thus we proved

$$\limsup_{n \rightarrow +\infty} \mathbb{P}(1 \notin H_\psi) \leq 2\varepsilon + K(1 - \varepsilon)^{K-1},$$

which clearly concludes the proof. ■

3 Maximum likelihood estimator for UA(n)

For a rooted tree T and a vertex v , we define $\text{Aut}(v, T)$ as follows. Let T_1, \dots, T_k be the subtrees of T rooted at the children of v (in particular $k \in \{d(v) - 1, d(v)\}$). Let S_1, \dots, S_L be the different isomorphism classes realized by these rooted subtrees. For $i \in [L]$, let $\ell_i = |\{j \in [k] : T_j^\circ = S_i\}|$. Finally we let $\text{Aut}(v, T) := \prod_{i=1}^L \ell_i!$.

Proposition 1 *Let T be an unlabeled rooted tree with n vertices, then*

$$|\{t \text{ recursive tree} : t^\circ = T\}| = \frac{n!}{\prod_{v \in V(T) \setminus \mathcal{L}(T)} (|T_{v\downarrow}| \cdot \text{Aut}(v, T))}.$$

Proof We prove this result by induction on n (the formula is clearly true for $n = 2$). Let T_1, \dots, T_k be the subtrees rooted at the children of \emptyset (in particular $k = d(\emptyset)$). An increasing labeling for T is obtained by partitioning $\{2, \dots, n\}$ into k subsets of sizes $|T_1|, \dots, |T_k|$, and then labeling correctly each subtree with the corresponding element of the partition. Clearly the number of possibilities for choosing the partition is $\frac{(n-1)!}{\prod_{i=1}^k |T_i|!}$, and given the partition the number of correct labelings is $\prod_{i=1}^k |\{t \text{ recursive tree} : t^\circ = T_i\}|$. In this calculation we have counted multiple times similar recursive trees (of size n). Indeed if T_i and T_j are isomorphic then we have considered separately the case where T_i is labeled with $S \subset \{2, \dots, n\}$ and T_j is labeled with S' , and the case where T_i is labeled with S' and T_j is labeled with S . In fact we have precisely overcounted by a factor $\text{Aut}(\emptyset, T)$. Thus we obtain the formula

$$|\{t \text{ recursive tree} : t^\circ = T\}| = \frac{(n-1)!}{\text{Aut}(\emptyset, T) \prod_{i=1}^k |T_i|!} \prod_{i=1}^k |\{t \text{ recursive tree} : t^\circ = T_i\}|,$$

which easily concludes the proof (by using the induction hypothesis). ■

Proposition 1 immediately implies that, given an observation T (i.e., an unlabeled tree on n vertices), the maximum likelihood estimator for the root in the uniform attachment model is the vertex minimizing the function

$$\zeta_T(u) = \prod_{v \in V(T) \setminus \mathcal{L}((T,u))} (|(T, u)_{v\downarrow}| \cdot \text{Aut}(v, (T, u))). \quad (3)$$

In fact it implies more generally that the optimal strategy to output a set of K vertices is to choose those with the smallest values of ζ . We denote the corresponding mapping by H_ζ . Using this representation for the optimal procedure we prove now the following impossibility result:

Theorem 4 *There exists $\varepsilon_0 > 0$ such that for all $\varepsilon \leq \varepsilon_0$, any procedure that satisfies (1) in the uniform attachment model must have $K \geq \exp\left(\sqrt{\frac{1}{30} \log \frac{1}{2\varepsilon}}\right)$.*

Proof Let³ $K = \exp\left(\sqrt{\frac{1}{30} \log \frac{1}{2\varepsilon}}\right)$, which we assume to be large enough (that is ε is small enough, see below). Since H_ζ is the optimal procedure, one clearly has that $\mathbb{P}(1 \notin H_\zeta(\text{UA}(n)^\circ))$ is non-decreasing with n , and thus to prove the theorem it is enough to show that $\mathbb{P}(1 \notin H_\zeta(\text{UA}(K+1)^\circ)) > \varepsilon$. Equivalently we need to show that there is a set \mathcal{T} of recursive trees on $K+1$ vertices such $\mathbb{P}(\text{UA}(K+1) \in \mathcal{T}) > \varepsilon$ and for any $T \in \mathcal{T}$ one has $\zeta_T(1) > \zeta_T(i), \forall i \in \{2, \dots, K+1\}$.

The set of recursive trees \mathcal{T} we consider consist of those where (i) vertices $\{1, \dots, 10 \log(K)\}$ form a path with 1 being an endpoint, (ii) all vertices in $\{10 \log(K) + 1, \dots, K+1\}$ are descendants of $10 \log(K)$, and (iii) the height of the subtree rooted at $10 \log(K)$ is smaller than $4 \log(K)$. Next we verify that this set of recursive trees has probability at least $\frac{1}{2} \exp(-30 \log^2(K)) = \varepsilon$. More precisely the probability that (i) happens is exactly $1/(10 \log(K) - 1)! \geq \exp(-10 \log^2(K))$; the probability that (iii) happens conditionally on (ii) is at least $1/2$ for K large enough (indeed it is well-known that the height of $\text{UA}(n)$ rescaled by $\log(n)$ converges in probability to e , see Devroye [1987]); and the probability that (ii) happens is equal to

$$\prod_{i=10 \log(K)}^K \left(1 - \frac{10 \log(K) - 1}{i}\right) \geq \exp(-20 \log^2(K)).$$

We show now that for trees in \mathcal{T} , the root 1 has the largest ζ value, thus concluding the proof. First observe that, in general for any tree T and vertices u, v , with (u_1, \dots, u_k) being the unique path with $u_1 = u$ and $u_k = v$,

$$\zeta_T(u) > \zeta_T(v) \Leftrightarrow \prod_{i=1}^k \left(|(T, u)_{u_i \downarrow}| \cdot \text{Aut}(u_i, (T, u))\right) > \prod_{i=1}^k \left(|(T, v)_{u_i \downarrow}| \cdot \text{Aut}(u_i, (T, v))\right).$$

Furthermore it is easy to verify that one always has

$$\text{Aut}(u_i, (T, v)) \leq |(T, v)_{u_i \downarrow}| \cdot \text{Aut}(u_i, (T, u)).$$

Indeed the computation of $\text{Aut}(u_i, (T, v))$ is based on a list of subtrees T_1, \dots, T_k , and only one of those subtrees is modified in the computation of $\text{Aut}(u_i, (T, u))$ (or possibly a subtree is added if $i = 1$), which results in a mutiplicative change of at most $k + 1 \leq |(T, v)_{u_i \downarrow}|$.

Putting together the two displays above, one obtains

$$\prod_{i=1}^k |(T, u)_{u_i \downarrow}| > \prod_{i=1}^k |(T, v)_{u_i \downarrow}|^2 \Rightarrow \zeta_T(u) > \zeta_T(v). \quad (4)$$

Now consider $T \in \mathcal{T}$, and $v \in \{10 \log(K) + 1, \dots, K+1\}$. Let (u_1, \dots, u_k) be the unique path with $u_1 = 1$ and $u_k = v$. Clearly

$$\prod_{i=1}^k |(T, v)_{u_i \downarrow}| \leq (K+1)^{4 \log(K)} \cdot (10 \log(K))! \leq (K+1)^{4 \log(K)} \cdot (10 \log(K))^{10 \log(K)},$$

³For clarity of the presentation we ignore integer rounding issues in this paper.

and

$$\prod_{i=1}^k |(T, u)_{u_i \downarrow}| \geq (K + 1 - 10 \log(K))^{10 \log(K)}.$$

Thus using (4) one has that for K large enough, $\zeta_T(1) > \zeta_T(v)$. Furthermore it is obvious that $\zeta_T(1) > \zeta_T(v)$ for $v \in \{2, \dots, 10 \log(K)\}$. This concludes the proof. \blacksquare

4 A subpolynomial root-finding algorithm for $\text{UA}(n)$

The performance of the maximum likelihood estimate is complicated to analyze due to the presence of the automorphism numbers in the expression (11). In order to circumvent this difficulty, we analyze algorithms that minimize modified versions of the maximum likelihood criterion. The function ψ defined in (2) can be viewed as such a “relaxation” of the likelihood function ζ defined in (11). In this section we analyze a tighter relaxation, which we denote φ and define as

$$\varphi_T(u) = \prod_{v \in V(T) \setminus \{u\}} |(T, u)_{v \downarrow}|. \quad (5)$$

We denote by H_φ the mapping which returns the set of K vertices with smallest φ values.

Theorem 5 *There exist universal constants $a, b > 0$ such that if $K \geq a \exp\left(b \frac{\log(1/\varepsilon)}{\log \log(1/\varepsilon)}\right)$, then one has $\liminf_{n \rightarrow +\infty} \mathbb{P}(1 \in H_\varphi(\text{UA}(n)^\circ)) \geq 1 - \varepsilon$.*

Before going into the proof of the above result, we start with a technical lemma on the concentration of sums of Gamma random variables.

Lemma 1 *Let $j_1, \dots, j_\ell \in \mathbb{N}$, and $s = \sum_{k=1}^{\ell} k j_k$. For $k \in [\ell]$, let⁴ $X_k \sim \text{Ga}(j_k, k)$, with X_1, \dots, X_ℓ being independent. Then for any $t \in (0, s)$,*

$$\mathbb{P}\left(\sum_{k=1}^{\ell} X_k < t\right) \leq \exp\left(-\sqrt{\frac{s}{2}} \log\left(\frac{s}{et}\right)\right).$$

Proof Recall that for any $\lambda \geq 0$, one has

$$\mathbb{E}_{X \sim \text{Ga}(a,b)} \exp(-\lambda X) = \frac{1}{(1 + \lambda b)^a}.$$

Thus using Chernoff’s method one obtains (using also that $x \mapsto \frac{\log(1+x)}{x}$ is non-increasing for the second inequality)

$$\begin{aligned} \mathbb{P}\left(\sum_{k=1}^{\ell} X_k < t\right) &\leq \exp(\lambda t) \mathbb{E} \exp\left(-\lambda \sum_{k=1}^{\ell} X_k\right) = \exp\left(\lambda t - \sum_{k=1}^{\ell} j_k \log(1 + \lambda k)\right) \\ &\leq \exp\left(\lambda t - \sum_{k=1}^{\ell} \lambda k j_k \frac{\log(1 + \lambda k)}{\lambda k}\right) \\ &= \exp\left(\lambda t - \frac{s}{\ell} \log(1 + \lambda \ell)\right). \end{aligned}$$

⁴We denote $\text{Ga}(a, b)$ for the Gamma distribution with density proportional to $x^{a-1} \exp(-x/b) \mathbf{1}_{\{x>0\}}$.

Taking $\lambda = (s/t - 1)/\ell > 0$ yields

$$\mathbb{P}\left(\sum_{k=1}^{\ell} X_k < t\right) \leq \exp\left(-\frac{s \log\left(\frac{s}{et}\right)}{\ell}\right).$$

It remains to observe that $s \geq \ell^2/2$ to conclude the proof. ■

We also recall Erdős' non-asymptotic version of the Hardy-Ramanujan formula on the number of partitions of an integer, Erdős [1942],

$$\left| \left\{ (j_1, \dots, j_\ell) \in \mathbb{N}^\ell \text{ s.t. } \ell \in \mathbb{N}, \sum_{k=1}^{\ell} k j_k = s \right\} \right| \leq \exp\left(\pi \sqrt{\frac{2}{3}s}\right). \quad (6)$$

Proof We decompose the proof into four steps. Denote $T = \text{UA}(n)^\circ$, and let $S \geq 1$ be a value appropriately chosen later.

Step 1. We first introduce an alternative labeling of the vertices in $\text{UA}(n)$. We define this labeling by induction as follows. The root is labeled with the empty set \emptyset . Node $(j_1, \dots, j_\ell) \in \mathbb{N}^\ell$ with $\ell \in \mathbb{N}$ is defined to be the j_ℓ^{th} children (in birth order) of node $(j_1, \dots, j_{\ell-1})$. Thus instead of labeling vertices with elements of \mathbb{N} , they are labeled with elements of $\mathbb{N}^* := \cup_{\ell=0}^{\infty} \mathbb{N}^\ell$. For any vertex $v \in \mathbb{N}^*$ we define $\ell(v)$ to be such that $v \in \mathbb{N}^{\ell(v)}$ (in other words $\ell(v)$ is the depth of v), and $s(v) = \sum_{k=1}^{\ell(v)} (\ell(v) + 1 - k) j_k$.

We observe the following important property: for any vertex v such that $s(v) > 3S$, one has either

- (i) there exists u such that $s(u) \in (S, 3S]$ and $v \in (T, \emptyset)_{u\downarrow}$, or,
- (ii) there exists u such that $s(u) \leq S$ and $v \in (T, \emptyset)_{(u,j)\downarrow}$ for some $j > S$.

We prove this property by induction on the depth $\ell(v)$ of v . For $\ell(v) = 1$, (ii) is clearly true with $u = \emptyset$. For $\ell(v) > 1$, let u be the parent of v . We now have three cases:

- (a) If $s(u) > 3S$, then one can apply the induction hypothesis on u .
- (b) If $s(u) \in (S, 3S]$, then (i) is true.
- (c) Finally if $s(u) \leq S$, then v is the j^{th} children of u with $j > S$ (this uses the fact that $s((u, j)) \leq j + 2S$, and $s(v) > 3S$), and thus (ii) is true.

This concludes the proof that either (i) or (ii) is always true. Using the fact that φ is monotone over paths, this proves that

$$\begin{aligned} & \mathbb{P}(\exists v : s(v) > 3S \text{ and } \varphi(v) \leq \varphi(1)) \\ & \leq \mathbb{P}(\exists v : s(v) \in (S, 3S] \text{ and } \varphi(v) \leq \varphi(1)) + \mathbb{P}(\exists v, j : s(v) \leq S, j > S, \text{ and } \varphi(v) \leq \varphi(1)). \end{aligned} \quad (7)$$

Step 2. We show here that, after a union bound on v , the second term in (7) is bounded by the first term. First observe that, for $v = (j_1, \dots, j_\ell)$, one has

$$\varphi(v) \leq \varphi(1) \Leftrightarrow \prod_{i=1}^{\ell(v)} |(T, \emptyset)_{(j_1, \dots, j_i)\downarrow}| \geq \prod_{i=0}^{\ell(v)-1} |(T, v)_{(j_1, \dots, j_i)\downarrow}| = \prod_{i=1}^{\ell(v)} (n - |(T, \emptyset)_{(j_1, \dots, j_i)\downarrow}|). \quad (8)$$

In particular, this implies that

$$\begin{aligned} & \exists j > S : \varphi((v, j)) \leq \varphi(1) \\ & \Rightarrow \prod_{i=1}^{\ell(v)} |(T, \emptyset)_{(j_1, \dots, j_i) \downarrow}| \cdot \left(\sum_{j=S+1}^{+\infty} |(T, \emptyset)_{(v, j) \downarrow}| \right) \geq \prod_{i=1}^{\ell(v)} (n - |(T, \emptyset)_{(j_1, \dots, j_i) \downarrow}|) \cdot \left(n - \sum_{j=S+1}^{+\infty} |(T, \emptyset)_{(v, j) \downarrow}| \right). \end{aligned}$$

Now it is easy to see that the random variable $\sum_{j=S+1}^{+\infty} |(T, \emptyset)_{(v, j) \downarrow}|$ is stochastically dominated by $|(T, \emptyset)_{(v, S) \downarrow}|$, which, together with the two above displays, show that

$$\mathbb{P}(\exists j > S : \varphi((v, j)) \leq \varphi(1)) \leq \mathbb{P}(\varphi((v, S)) \leq \varphi(1)).$$

In particular, putting this together with (7) and an union bound, one obtains:

$$\mathbb{P}(\exists v : s(v) > 3S \text{ and } \varphi(v) \leq \varphi(1)) \leq 2 \sum_{v: s(v) \in (S, 3S]} \mathbb{P}(\varphi(v) \leq \varphi(1)). \quad (9)$$

Step 3. This is the main step of the proof, where we show that for any vertex v with $s(v) \geq 10^{10}$,

$$\limsup_{n \rightarrow +\infty} \mathbb{P}(\varphi(v) \leq \varphi(1)) \leq 7 \exp\left(-0.21\sqrt{s(v)} \log(s(v))\right). \quad (10)$$

Observe that the random vector $(\frac{1}{n} |(T, \emptyset)_{(j_1, \dots, j_i) \downarrow}|)_{i=1, \dots, \ell(v)}$ evolves according to a "pile" of Pólya urns, and thus using a standard argument it converges in distribution to the random vector $(\prod_{k=1}^i U_{j_k, k})_{i=1, \dots, \ell(v)}$, where $U_{j,1}, U_{j,2}, \dots$ are (independent) products of j independent uniform random variables in $[0, 1]$. In particular, by (8) this gives

$$\begin{aligned} & \limsup_{n \rightarrow +\infty} \mathbb{P}(\varphi(v) \leq \varphi(1)) \\ & = \mathbb{P}\left(\prod_{i=1}^{\ell(v)} \prod_{k=1}^i U_{j_k, k} \geq \prod_{i=1}^{\ell(v)} \left(1 - \prod_{k=1}^i U_{j_k, k}\right)\right) \\ & = \mathbb{P}\left(\prod_{i=1}^{\ell(v)} U_{j_i, i}^{\ell(v)+1-i} \geq \prod_{i=1}^{\ell(v)} \left(1 - \prod_{k=1}^i U_{j_k, k}\right)\right) \\ & \leq \mathbb{P}\left(\prod_{i=1}^{\ell(v)} U_{j_i, i}^{\ell(v)+1-i} \geq \exp\left(-\frac{s(v)}{R}\right)\right) + \mathbb{P}\left(\prod_{i=1}^{\ell(v)} \left(1 - \prod_{k=1}^i U_{j_k, k}\right) \leq \exp\left(-\frac{s(v)}{R}\right)\right). \end{aligned}$$

where $R > 0$ will be appropriately chosen later. Using Lemma 1, one directly obtains

$$\mathbb{P}\left(\prod_{i=1}^{\ell(v)} U_{j_i, i}^{\ell(v)+1-i} \geq \exp\left(-\frac{s(v)}{R}\right)\right) \leq \exp\left(-\sqrt{\frac{s(v)}{2}} \log\left(\frac{R}{e}\right)\right).$$

Furthermore, using that $1 - \exp(-x) \geq \frac{1}{2} \min(x, 1)$ for any $x \geq 0$, one can check that

$$\prod_{i=1}^{\ell(v)} \left(1 - \prod_{k=1}^i U_{j_k, k}\right) \succeq \frac{1}{2^{\ell(v)}} \prod_{i=1}^{\ell(v)} \min\left(\sum_{k=1}^i \log(1/U_{j_k, k}), 1\right) \succeq \frac{1}{2^{\ell(v)}} X,$$

where $X := \prod_{i=1}^{+\infty} \min\left(\sum_{k=1}^i \log(1/U_{1,k}), 1\right)$. Thus, using Lemma 2 in the Appendix, one gets,

$$\mathbb{P}\left(\prod_{i=1}^{\ell(v)} \left(1 - \prod_{k=1}^i U_{j_k, k}\right) \leq \exp\left(-\frac{s(v)}{R}\right)\right) \leq 6 \cdot 2^{\frac{\ell(v)}{4}} \exp\left(-\frac{s(v)}{4R}\right) \leq 6 \exp\left(\frac{\sqrt{s(v)}}{4} - \frac{s(v)}{4R}\right),$$

where we used $s(v) \geq \ell(v)/2$ for the second inequality. Taking $R = e \cdot s(v)^{0.3}$ easily concludes the proof of (10).

Step 4. Putting together (9) and (10) one obtains that for $S \geq 10^{10}$,

$$\limsup_{n \rightarrow +\infty} \mathbb{P}(\exists v : s(v) > 3S \text{ and } \varphi(v) \leq \varphi(1)) \leq 14 \cdot |\{v : s(v) \leq 3S\}| \cdot \exp\left(-0.21\sqrt{S} \log(S)\right).$$

Furthermore by (6) one clearly has $|\{v : s(v) \leq 3S\}| \leq 3S \exp(\pi\sqrt{2S})$, and thus for $S \geq 10^{10}$ one can check that

$$\limsup_{n \rightarrow +\infty} \mathbb{P}(\exists v : s(v) > 3S \text{ and } \varphi(v) \leq \varphi(1)) \leq \exp\left(-\frac{1}{100}\sqrt{S} \log(S)\right),$$

which easily concludes the proof. ■

5 Root-finding algorithms for $\text{PA}(n)$

In this section we investigate the preferential attachment model. We first analyze the simple root-finding algorithm defined in Section 2.

Theorem 6 *If $K \geq 8 \frac{\log(1/\varepsilon)}{\varepsilon^2}$, then one has $\liminf_{n \rightarrow +\infty} \mathbb{P}(1 \in H_\psi(\text{PA}(n)^\circ)) \geq 1 - \frac{4\varepsilon}{1-\varepsilon}$.*

Similarly to the situation for uniform attachment we observe that the above theorem is optimal up to a logarithmic factor, since with probability at least $\Omega(1/\sqrt{K})$, the root is a leaf in $\text{PA}(K)$ (and thus for H_ψ to be correct with probability at least $1 - \varepsilon$ one needs $K = \Omega(1/\varepsilon^2)$). We also note that a similar result (up to a worse logarithmic factor) could be obtained by simply considering the vertices with largest degree. However the proof of this latter statement is more involved technically, because degrees behave as triangular Pólya urns which have more complicated limiting distributions, see, e.g., Janson [2006]. On the contrary, subtree sizes are Pólya urns with diagonal replacement matrix, which allows for the simplicity of the following proof.

Proof Similarly to the proof of Theorem 3 we label vertices by chronological order, and for $1 \leq i \leq k$ we denote by $T_{i,k}$ the tree containing vertex i in the forest obtained by removing in $\text{PA}(n)$ all edges between vertices $\{1, \dots, k\}$. Here the vector $(|T_{1,k}|, \dots, |T_{k,k}|)$ follows a Pólya urn with k colors and replacement matrix $2I_k$ (where I_k is the $k \times k$ identity matrix), and thus using again a classical result one has the following convergence in distribution:

$$\frac{1}{n}(|T_{1,k}|, \dots, |T_{k,k}|) \xrightarrow[n \rightarrow +\infty]{} \text{Dir}\left(\frac{1}{2}, \dots, \frac{1}{2}\right).$$

Let $\eta \in (0, 1)$ and observe that

$$\mathbb{P}(1 \notin H_\psi) \leq \mathbb{P}(\exists i > K : \psi(i) \leq \psi(1)) \leq \mathbb{P}(\psi(1) \geq (1 - \eta)n) + \mathbb{P}(\exists i > K : \psi(i) \leq (1 - \eta)n).$$

Using that $\psi(1) \leq \max(|T_{1,2}|, |T_{2,2}|)$, and since $|T_{1,2}|/n$ and $|T_{2,2}|/n$ are identitically distributed and converge in distribution to a Beta(1/2, 1/2), one has

$$\limsup_{n \rightarrow +\infty} \mathbb{P}(\psi(1) \geq (1 - \eta)n) \leq 2 \lim_{n \rightarrow +\infty} \mathbb{P}(|T_{1,2}| \geq (1 - \eta)n) = \frac{2}{\pi} \arcsin(\sqrt{\eta}) \leq \sqrt{\eta}.$$

On the other hand, for any $i > K$, $\psi(i) \geq \min_{1 \leq k \leq K} \sum_{j=1, j \neq k}^K |T_{j,K}|$, and $\frac{1}{n} \sum_{j=1, j \neq k}^K |T_{j,K}|$ converges in distribution to a Beta((K - 1)/2, 1/2), which implies

$$\begin{aligned} \limsup_{n \rightarrow +\infty} \mathbb{P}(\exists i > K : \psi(i) \leq (1 - \varepsilon)n) &\leq \lim_{n \rightarrow +\infty} \mathbb{P}\left(\exists 1 \leq k \leq K : \sum_{j=1, j \neq k}^K |T_{j,K}| \leq (1 - \varepsilon)n\right) \\ &= \frac{K}{\sqrt{\eta}} (1 - \eta)^{\frac{K-1}{2}}, \end{aligned}$$

where the last line follows from a rather easy calculation on the Beta((K - 1)/2, 1/2) distribution. By taking $\eta = \varepsilon^2$ we thus proved

$$\limsup_{n \rightarrow +\infty} \mathbb{P}(1 \notin H_\psi) \leq 2\varepsilon + \frac{K}{\varepsilon} (1 - \varepsilon^2)^{\frac{K-1}{2}},$$

which clearly concludes the proof. ■

Next we prove an impossibility result. On the contrary to the proof of Theorem 4, here we do not use the structure of the maximum likelihood estimator, as a simple symmetry argument suffices.

Theorem 7 *There exists $c > 0$ such that for $\varepsilon \in (0, 1)$, any procedure that satisfies (1) in the preferential attachment model must have $K \geq c/\varepsilon$.*

Proof As we observed in the proof of Theorem 4, the probability of error for the optimal procedure is non-decreasing with n , so it suffices to show that the optimal procedure must have a probability of error of at least ε for some finite n .

We show that there is some finite n for which, with probability at least 2ε , 1 is isomorphic to at least $2c/\varepsilon$ vertices in PA(n). This clearly implies a probability of error of at least ε for any procedure that outputs less than c/ε vertices. To simplify the proof we now use the Θ notation. First observe that the probability that 1 is a leaf is $\Theta(1/\sqrt{n})$, and thus for $n = \Theta(1/\varepsilon^2)$ this happens with probability at least $\Theta(\varepsilon)$. Furthermore it is an easy exercise⁵ to verify that, conditioned on 1 being a leaf, with probability at least $\Theta(1)$, 2 is connected to $\Theta(\sqrt{n}) = \Theta(1/\varepsilon)$ leaves, which are then isomorphic to 1. ■

For sake of completeness, we observe that the maximum likelihood estimator in the preferential attachment model is obtained by minimizing the function

$$\xi_T(u) = \frac{1}{d_T(u)} \prod_{v \in V(T) \setminus \mathcal{L}((T, u))} (|(T, u)_{v \downarrow}| \cdot \text{Aut}(v, (T, u))). \quad (11)$$

This is a consequence of the following observation.

⁵Simply recall that $d_{\text{PA}(n)}(1)/\sqrt{n}$ converges in distribution (this follows from Janson [2006]) and at least a constant fraction of the vertices in $\{n/2, \dots, n\}$ are leaves.

Proposition 2 *Let T be an unlabeled rooted tree with n vertices, then*

$$|\{t \text{ plane-oriented recursive tree} : t^\circ = T\}| = \frac{n! \cdot d_T(\emptyset)! \cdot \prod_{v \in V(T) \setminus \{\emptyset\}} (d_T(v) - 1)!}{\prod_{v \in V(T) \setminus \mathcal{L}(T)} (|T_{v\downarrow}| \cdot \text{Aut}(v, T))}.$$

Proof Recall that a plane-oriented recursive tree is a recursive tree together with an ordering of the children of each vertex. For the root there are $d_T(\emptyset)!$ possible orderings of the children, while for any other vertex $v \neq \emptyset$ there are only $(d_T(v) - 1)!$ possible orderings. Together with Proposition 1 this immediately yields the above formula. \blacksquare

6 Open problems

1. Our results leave gaps for both uniform attachment and preferential attachment models. In particular, is it possible to find a procedure with $K \leq \exp(c\sqrt{\log(1/\varepsilon)})$ (respectively $K \leq c/\varepsilon$) in the uniform attachment (respectively preferential attachment) model?
2. In the seeded models of Bubeck et al. [2014b,a], how large does K need to be to certify that the obtained set of vertices contains the seed with probability at least $1 - \varepsilon$?
3. In Section 1.3 we introduced the more general non-linear preferential attachment model $\mathcal{T}_\alpha(n)$. What can be said for this model?
4. What about growing graphs instead of trees? For instance one can investigate root-finding algorithms in the model $\mathcal{T}_{\alpha,m}(n)$, where each arriving new node sends m independent edges according to the same rule than in $\mathcal{T}_\alpha(n)$.
5. If one has to output only a single vertex, what is the best possible success probability? For instance our results imply a lower bound of at least 1% success probability in the uniform attachment model (the strategy is to pick at random a vertex in H_ψ with $K = 35$).

Appendix

Lemma 2 *Let E_1, E_2, \dots be an i.i.d. sequence of exponential random variables, and let*

$$X = \prod_{i=1}^{+\infty} \min \left(\sum_{k=1}^i E_k, 1 \right).$$

Then for any $t > 0$,

$$\mathbb{P}(X \leq t) \leq 6t^{1/4}.$$

Proof First observe that, almost surely, X is a finite product, and thus one can write for $t \in (0, 1)$,

$$\begin{aligned} \mathbb{P}(X \leq t) &= \mathbb{P} \left(\exists \ell : \sum_{k=1}^{\ell+1} E_k > 1, \text{ and } \sum_{k=1}^{\ell} E_k \leq 1, \text{ and } \prod_{i=1}^{\ell} \left(\sum_{k=1}^i E_k \right) \leq t \right) \\ &\leq \sum_{\ell=1}^{+\infty} \min \left(\mathbb{P} \left(\sum_{k=1}^{\ell} E_k \leq 1 \right), \mathbb{P} \left(\sum_{k=1}^{\ell+1} E_k > 1, \text{ and } \prod_{i=1}^{\ell} \left(\sum_{k=1}^i E_k \right) \leq t \right) \right). \quad (12) \end{aligned}$$

One has

$$\mathbb{P}\left(\sum_{k=1}^{\ell} E_k \leq 1\right) = \int_0^1 \frac{x^{\ell-1} \exp(-x)}{(\ell-1)!} dx \leq \frac{1}{\ell!}.$$

Furthermore we show below that

$$\mathbb{P}\left(\sum_{k=1}^{\ell+1} E_k > 1, \text{ and } \prod_{i=1}^{\ell} \left(\sum_{k=1}^i E_k\right) \leq t\right) \leq 2^\ell \sqrt{t}. \quad (13)$$

The two above displays together with (12) concludes the proof, since they imply (with $\min(a, b) \leq \sqrt{ab}$):

$$\mathbb{P}(X \leq t) \leq \sum_{\ell=1}^{+\infty} \sqrt{\frac{2^\ell \sqrt{t}}{\ell!}} \leq 6t^{1/4}.$$

To prove (13) we first write

$$\prod_{i=1}^{\ell} \left(\sum_{k=1}^i E_k\right) = \left(\sum_{k=1}^{\ell+1} E_k\right) \prod_{i=1}^{\ell} \left(\frac{\sum_{k=1}^i E_k}{\sum_{k=1}^{\ell+1} E_k}\right).$$

Now observe that the vector $\left(\sum_{k=1}^i E_k / \sum_{k=1}^{\ell+1} E_k\right)_{i=1, \dots, \ell}$ is equal in law to $(U_{(i)})_{i=1, \dots, \ell}$ where U_1, \dots, U_ℓ is an i.i.d. sequence of uniform random variables in $[0, 1]$. Thus one has

$$\mathbb{P}\left(\sum_{k=1}^{\ell+1} E_k > 1, \text{ and } \prod_{i=1}^{\ell} \left(\sum_{k=1}^i E_k\right) \leq t\right) \leq \mathbb{P}\left(\prod_{i=1}^{\ell} U_i \leq t\right) \leq \mathbb{E}\left(\frac{\sqrt{t}}{\sqrt{\prod_{i=1}^{\ell} U_i}}\right) \leq 2^\ell \sqrt{t},$$

which concludes the proof. ■

References

- Christian Borgs, Michael Brautbar, Jennifer Chayes, Sanjeev Khanna, and Brendan Lucier. The power of local information in social networks. In *Internet and Network Economics*, pages 406–419. Springer, 2012.
- Michael Brautbar and Michael J. Kearns. Local algorithms for finding interesting individuals in large networks. In *Innovations in Theoretical Computer Science (ITCS)*, 2010.
- Sébastien Bubeck, Ronen Eldan, Elchanan Mossel, and Miklós Z. Rácz. From trees to seeds: on the inference of the seed from large trees in the uniform attachment model. *arXiv preprint arXiv:1409.7685*, 2014a.
- Sébastien Bubeck, Elchanan Mossel, and Miklós Z. Rácz. On the influence of the seed graph in the preferential attachment model. *arXiv preprint arXiv:1401.4849*, 2014b.
- Nicolas Curien, Thomas Duquesne, Igor Kortchemski, and Ioan Manolescu. Scaling limits and influence of the seed graph in preferential attachment trees. *arXiv preprint arXiv:1406.1758*, 2014.

- Luc Devroye. Branching processes in the analysis of the heights of trees. *Acta Informatica*, 24(3): 277–298, 1987.
- Michael Drmota. *Random trees: an interplay between combinatorics and probability*. Springer, 2009.
- Paul Erdős. On an elementary proof of some asymptotic formulas in the theory of partitions. *Annals of Mathematics*, 43(3):pp. 437–450, 1942.
- Alan Frieze and Wesley Pegden. Looking for vertex number one. *arXiv preprint arXiv:1408.6821*, 2014.
- Godfrey H Hardy and Srinivasa Ramanujan. Asymptotic formulæ in combinatory analysis. *Proceedings of the London Mathematical Society*, 2(1):75–115, 1918.
- Svante Janson. Limit theorems for triangular urn schemes. *Probability Theory and Related Fields*, 134(3):417–452, 2006.