Some Half-Baked Thoughts on Graph Sampling and Parameter Estimation

T. Mitchell Roddenberry

Version 0.2 – August 18, 2021

Introduction

The goal of this document is to convey a broad view of how graph parameter estimation via random sampling might be studied. I use the word "might" because this note does *not* properly survey the literature, nor does it even claim to be correct. Rather, it reflects how I have come to view this subject based on some light reading, and can perhaps provide perspective for exploration of related problems. Thus, the ensuing discussion will be fairly high-level, rather than focusing in on particular (useful!) constructions and ideas in graph parameter estimation. In some sense, everything I say here may be quite obvious to the probabilistically-inclined, but maybe less trivial to those in machine learning or network science.

Essentially, this document will describe how random sampling procedures induce a pseudometric on the set of graph isomorphism classes, and how that pseudometric can be used to study the estimability of graph parameters. Although some of the developments might seem a bit tricky, the big idea is to show that a parameter being estimable under a random sampling model is equivalent to uniform continuity in the induced metric. Moreover, reasonable assumptions on the sampling model will result in this metric space being compact, so that uniform continuity is equivalent to continuity. I will assume a general understanding of how ϵ-bashing is done in basic analysis on metric spaces, but will provide gentle reminders of definitions for concepts such as compactness, sequential compactness, uniform continuity, and so on. One should of course consult their favorite resource on introductory real analysis for a better review of these topics.

Graph Parameters and the Estimation Thereof

We are concerned with computing summaries of large graphs, also known as *graph parameters*. That is to say,

Definition 1. *Let $\mathcal{G}$ be the set of all graphs. A **graph parameter** is a map $p : \mathcal{G} \to \mathbb{R}$ such that for any isomorphic $X, Y \in \mathcal{G}$, we have $p(X) = p(Y)$. That is to say, p is a function on the isomorphism classes in $\mathcal{G}$.*

Graph parameters are closely related to graph properties, such as the property of being k-regular or bipartite. A graph property can be envisioned as a graph parameter that happens to only take values $0, 1$. However, the nature of an estimate for a graph property requires a binary output (perhaps coupled with a probability), while graph parameter estimation allows for arbitrary real-valued output.

To motivate the *estimation* of graph parameters, let us consider Example 1: the number of connected components of a graph. For a graph on n nodes and m edges, the number of connected components can (pretty much) be counted in $\mathcal{O}(n + m)$ time. That is, if you allow for an arbitrarily large graph, an exact calculation of the number of connected components can take a corresponding arbitrarily long time. To get around this, one may wish to sample a subgraph of size much smaller than the original, perhaps on some constant number of nodes. Then, by examining this random subgraph, one forms an estimate of the number of connected components. If the sampled subgraph is indeed of fixed size, such an algorithm has a runtime independent of the size of the original graph. This clearly is not an exact computation, but the question remains: is it even a good approach? It isn't: to see this, observe that for some bound n on the number of nodes in the sampled subgraph, there are 2^{n^2} different graphs that can be sampled. However, the input graphs can be made arbitrarily large, with a corresponding arbitrarily large number of connected components. Since any sampling procedure will certainly not be injective, two graphs with significantly different numbers of connected components will necessarily be mapped to the same small graph. In this case, any estimator will fail to distinguish them, despite the vast difference in their parameters.

Other parameters are more reasonably estimated by constant-time sampling, though. In Example 2, one could quite easily envision a constant-sized subsample of a large graph being a reasonable way to estimate the average node degree, particularly under extra assumptions on the structure of the graph (for instance, bounded degree nodes). This points to some difference between parameters that resemble densities or averages across the graph, and those that are more fragile in the context of the graph's global structure. In the remainder of this document, we will form some ways to understand this difference, giving a precise meaning to what this fragility really is.

Example 1. *The number of **connected components** in a graph.*

Example 2. *The **average node degree**, given by twice the number of edges divided by the number of nodes. Similarly, the **edge density**, approximately given by the number of edges divided by the number of nodes squared.*

Example 3. *The **global clustering coefficient**, which measures how likely a set of three nodes connected by a "path" is to be closed. This is a measure of transitivity, in the form of a probability.*

What do We Mean by "Constant Time?"

When randomly subsampling a graph, one can consider a computational model for how subgraphs are drawn. In particular, such a model will involve some number of *queries*, each with an associated computational cost, such as:

1. Select a node v from the graph uniformly at random: $\mathcal{O}(1)$.
2. Given a node v, return the set of nodes in its neighborhood: $\mathcal{O}(d_v)$.
3. Given two nodes u, v, return whether or not they are adjacent: $\mathcal{O}(1)$.
4. Given a node v, randomly sample one of its neighbors: $\mathcal{O}(1)$.

A sampling algorithm, then, performs a sequence of such queries in a way designed to yield a subgraph that can be used to estimate graph parameters.

See Example 4 and Example 5 in the margin for examples of random sampling algorithms, with corresponding runtimes. In particular, we see that Example 4 has a runtime dependant on the maximum degree of the graph, which is strictly bounded by n in the set of all possible graphs (consider the complete graphs K_n). However, if one considers the set of graphs with bounded degree, then this runtime is asymptotically independent of the graph size: we say that it is a constant time algorithm in this case.

On the other hand, Example 5 has runtime determined by the depth k of the random walk. If k is chosen independently of the graph size n, then this approach is a constant time algorithm for any graph, not just those with some regularity condition such as having bounded degree.

Observe that under this computational model for sampling, constant time algorithms yield graphs of bounded size. Since we do not consider decorated graphs, *e.g.*, with weighted edges, the set of possible sampled graphs is finite in this case. For this reason, not only does the sampling procedure run in constant time, but so does any estimation procedure. That is to say, if we denote some constant time sampling procedure by $\mathcal{S}$, and a graph parameter estimator by $\widehat{p} : \mathrm{Ima}(\mathcal{S}) \to \mathbb{R}$, the procedure $\widehat{p} \circ \mathcal{S}$ has a runtime that is asymptotically independent of the input graph's size. For instance, we can allow a given graph G to be arbitrarily large, but $\mathcal{S}(G)$ will always be bounded in size. This allows for $\widehat{p}$ to run in constant time as well.

We can now state the object of study of this document. Let Σ be a subset of the isomorphism classes of graph, *e.g.*, all graphs of degree

Some examples of sampling algorithms.

Example 4. *Randomly sample a node v ($\mathcal{O}(1)$), then return the subgraph induced by v and its neighborhood ($\mathcal{O}(d_v)$ to get the neighborhood, then $\mathcal{O}(d_v^2)$ to check all pairs of nodes for edges). This algorithm has a worst-case runtime $\mathcal{O}(d_{\max}^2)$, where $d_{\max}$ is the maximum degree of the graph.*

Example 5. *Randomly sample a node v ($\mathcal{O}(1)$), then perform a random walk for k steps ($\mathcal{O}(k)$). Return the subgraph induced by the set of visited nodes ($\mathcal{O}(k^2)$). This algorithm has a worst-case runtime of $\mathcal{O}(k^2)$.*

bounded by some d. Let $\mathcal{S}$ be a random sampling algorithm that runs in constant time on graphs in Σ, and denote the set of possible outputs by Σ'. For some graph parameter $p : \Sigma \to \mathbb{R}$, we want to find an estimator $\widehat{p} : \Sigma' \to \mathbb{R}$ such that the diagram pictured in Figure 1 "approximately commutes."

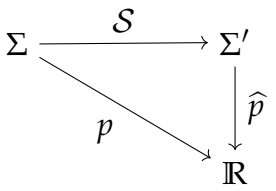

Figure 1: Diagram illustrating the estimation of a graph parameter p via a sampling algorithm $\mathcal{S}$ followed by an estimator $\widehat{p}$.

The Metric of a Sampling Algorithm

An implicit hypothesis in the application of a constant time sampling algorithm is that a graph parameter can be well understood by a small subgraph. Thus, two graphs with similar statistical properties in terms of their small subgraphs ought to have similar parameters as well. In this section, we will use this idea as a way to describe a *metric* on a class of graphs determined by a sampling algorithm. To ground this, we will first consider the sampling of bandlimited signals as a motivating analogy to the study of random graph sampling.

Shannon-Nyquist Sampling

Consider square-integrable real-valued functions on the interval: $L^2[0,1]$. Denote by $L^2_B[0,1]$ the subspace of $L^2[0,1]$ consisting of B-bandlimited functions. The sampling theorem says that taking equispaced samples of a function $f \in L^2_B[0,1]$ over $2B$ points will "preserve all of the information" of that function, so that f can be reconstructed from its samples. One can also interpret this by saying that any $f, g \in L^2_B[0,1]$ can be distinguished by the sampled vector, lying in $\mathbb{R}^{2B}$. In the class of bandlimited functions, if two functions have the same set of samples, then they must be equal.

Denote the Shannon-Nyquist sampling map by $\mathcal{S} : L^2_B[0,1] \to \mathbb{R}^{2B}$. For any function $p : L^2_B[0,1] \to \mathbb{R}$, there exists an estimator $\widehat{p} : \mathbb{R}^{2B} \to \mathbb{R}$ such that $\widehat{p} \circ \mathcal{S} = p$. This is great news: any parameter on the space of bandlimited functions can be computed exactly from its samples.

However, in the case of sub-Nyquist sampling, there exist two functions $f, g \in L^2_B[0,1]$ such that $f \neq g$ and $\mathcal{S}(f) = \mathcal{S}(g)$. If p is such that $p(f) \neq p(g)$, then it is *impossible* to compute p from the samples of a function.

In this case, the sampling map has induced a pseudometric d on the space $L^2_B[0,1]$. If p is such that $d(f,g) = 0$ implies $p(f) = p(g)$, then p can be recovered from the samples of a function. However, if this condition is not met, there will *always* be ambiguity present.

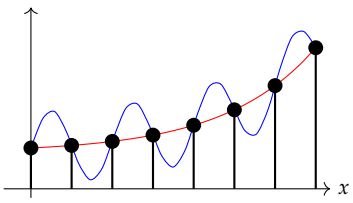

Figure 2: The ambiguity present in sub-Nyquist sampling. Observe that in this case, the two functions in $L^2[0,1]$ have the same set of samples, due to the fact that sub-Nyquist sampling is not bijective. Thus, the sampling map can't induce a proper metric on $L^2[0,1]$, since there are distinct points in the space that it fails to distinguish.

Graph Sampling

With a general idea of what sampling is, we are now ready to discuss how some of these ideas might port over to random graph sampling. Throughout, Σ will denote some set of isomorphism classes of graphs, such as those with maximum degree bounded by d. We will also assume that we have some constant-time random sampler $\mathcal{S} : \Sigma \to \Sigma'$, where Σ' is just the image of Σ under $\mathcal{S}$. It could very well be that Σ' is a subset of Σ, or perhaps Σ' is some abstract discrete set, or maybe each element of Σ' is some element of Σ' plus some extra information. Either way, the constant-time assumption of $\mathcal{S}$ will generally tell us that Σ' is a finite set.

Given that $\mathcal{S}$ is a *random* sampling algorithm, it is not quite correct to characterize it as a map from Σ to Σ'. Rather, it maps elements of Σ to *distributions* on Σ'. Denoting the set of all probability distributions on Σ' by $\Pi\Sigma'$, we define a random sampling map as follows:

Definition 2. *Let Σ be a set of graph isomorphism classes, and Σ' some finite set. A **random sampling map** $\mathcal{S}$ from Σ to Σ' is a function*

$$\mathcal{S} : \Sigma \to \Pi\Sigma'. \tag{1}$$

In practice, one would apply the map $\mathcal{S}$, then draw a sample from the resultant distribution.

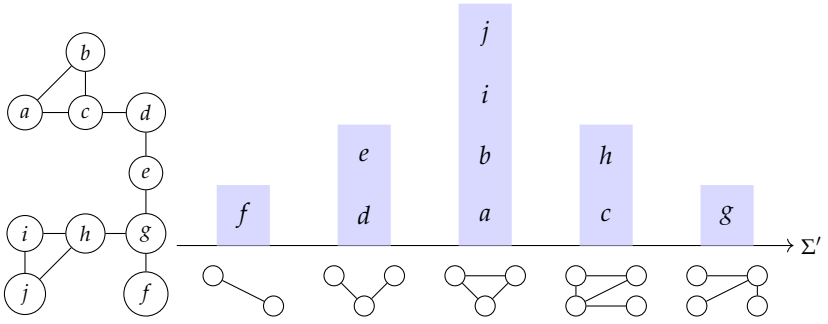

Figure 3: The distribution yielded by the following sampling algorithm: "pick a node uniformly at random, then return the induced subgraph of its neighborhood." For bounded-degree graphs, this algorithm runs in constant time, as discussed before.

By viewing sampling algorithms as functions that yield distributions, rather than a single random sample, we can evaluate how well a parameter can be estimated in a probabilistic sense. Moreover, we can use this space of distributions to induce a pseudometric on Σ.

We will proceed in the following way. First, we will define a metric on $\Pi\Sigma'$. Then, we will pull said metric back via $\mathcal{S}$, thus forming a pseudometric on Σ. This pseudometric will then establish the degree to which two graphs can be discerned by $\mathcal{S}$ in a probabilistic sense.

Since Σ' is a finite set, we can easily consider the total variation

Definition 3. *Let X be a set. Consider a function $d : X \times X \to \mathbb{R}$ such that for every $x, y, z \in X$,*

1. *If $x = y$, then $d(x, y) = 0$*
2. *$d(x, y) = d(y, x)$*
3. *$d(x, z) \leq d(x, y) + d(y, z)$*

*Under these conditions, we say that d is a **pseudometric** on X, and that (X, d) is a **pseudometric space**. Note that the definition of a pseudometric space is a slight relaxation of the familiar notion of a metric space, in which we allow distinct elements to be indistinguishable: that is, we allow $d(x, y) = 0$ for $x \neq y$. If we strengthen the first requirement to an if-and-only-if statement, then d is called a **metric** on X, and (X, d) is a **metric space**.*

distance on $\Pi\Sigma'$. Define the metric

$$\begin{aligned} d_1 : \Pi\Sigma' \times \Pi\Sigma' &\to \mathbb{R} \\ (\lambda, \mu) &\mapsto \frac{1}{2} \sum_{x \in \Sigma'} |\lambda(x) - \mu(x)|. \end{aligned} \tag{2}$$

One can check that d_1 is indeed a metric on $\Pi\Sigma'$. It is also a natural choice for a metric, as it describes the "overlap" between two distributions quite well. Of course, two identical distributions will have a total variation distance of zero. Additionally, two distributions will have a total variation distance of one if and only if they have disjoint support. In general, the total variation distance lies in the interval $[0, 1]$, and measures the total overlapping "area" between the two distributions[1].

Given a suitable metric on $\Pi\Sigma'$, we will now *pull it back* by $\mathcal{S}$ to define a pseudometric on Σ[2]. Abusing notation a bit, this yields a pseudometric on Σ:

$$\begin{aligned} \mathcal{S}^* d_1 : \Sigma \times \Sigma &\to \mathbb{R} \\ (K, L) &\mapsto d_1(\mathcal{S}(K), \mathcal{S}(L)). \end{aligned} \tag{3}$$

This is not necessarily a metric on Σ, since two graphs in Σ may have the same value under the map $\mathcal{S}$: these two graphs would then be indiscernible by $\mathcal{S}$. Referring back to our earlier example regarding Shannon-Nyquist sampling, this is analogous to two signals having the same set of samples due to sampling at sub-Nyquist rates. In the case of two graphs having a small distance $\mathcal{S}^* d_1$, their corresponding distributions on Σ' have substantial overlap. Thus, in a probabilistic sense, it will be difficult to form a good estimator to distinguish them as well, since the sampling map has a high chance of yielding the same object in Σ' for both of them.

That is to say, if two graphs have very small distance $\mathcal{S}^* d_1$, then the parameter p can only be estimated with reasonable accuracy if it takes a similar value on both graphs. We formalize this idea in the next section.

Continuity of Graph Parameters

LOOKING BACK AT GRAPH PARAMETERS: a graph parameter for Σ is simply a map $p : \Sigma \to \mathbb{R}$. Based on our developments, the use of of a random sampling algorithm induces a pseudometric on Σ via the pullback of the total variation distance by the sampling map. That is to say, p is a map from a pseudometric space to the real numbers: we can thus study its continuity properties, drawing conclusions about

[1] The observant reader might notice that the total variation distance resembles a 1-norm on $\mathbb{R}^n$. One can define similar metrics in analogy to p-norms, and even ∞-norms. However, all such metrics on $\Pi\Sigma'$ turn out to be strongly equivalent, and thus will turn out to not matter for our later discussion. Worry not: this will be noted again later!

[2] At this point I'm just being pompous in my vocabulary: pulling d_1 back by $\mathcal{S}$ is just a fancy way of saying that we compose the two functions.

its estimability along the way. One such result will pertain directly to Example 1: as a sneak peek, we state it here.

Proposition. *(Informal) There exists no constant time sampling algorithm on $\mathcal{G}$ such that the number of connected components can be estimated uniformly well.*

As we will see, the proof will be hardly graph-theoretic at all. Indeed, by forming an appropriate pseudometric on $\mathcal{G}$, it follows quite simply by basic arguments from analysis on metric spaces.

Let us continue with our investigation of the pseudometric space $(\Sigma, \mathcal{S}^* d_1)$. In particular, we have the following convenient fact at our disposal.

Proposition 1. *The pseudometric space $(\Sigma, \mathcal{S}^* d_1)$ is compact. Therefore, any ϵ-continuous map $p : \Sigma \to \mathbb{R}$ is also (ϵ, δ)-uniformly continuous, for some $\delta > 0$*[3].

With the proper definitions and basic results in place, we are now ready to put forth the primary expression of this document. We wish to characterize the existence of an estimator $\widehat{p} : \Sigma' \to \mathbb{R}$ such that $p \approx \widehat{p} \circ \mathcal{S}$ with sufficiently large probability. Dwell on the following definition, which formalizes this notion.

Definition 6. *A parameter $p : \Sigma \to \mathbb{R}$ is said to be **ϵ-estimable** under a random sampling map $\mathcal{S} : \Sigma \to \Pi\Sigma'$ if there exists an estimator $\widehat{p} : \Sigma' \to \mathbb{R}$ such that for all $K \in \Sigma$, we have*

$$\mathbb{P}\left\{|p(K) - \widehat{p}(\mathcal{S}(K))| < \epsilon\right\} > 1 - \epsilon. \tag{4}$$

*In this case, we say that $\widehat{p}$ is an **ϵ-estimator** for p under $\mathcal{S}$.*

After all of this trouble, we arrive at a nice statement characterizing ϵ-estimable functions under a random sampling map $\mathcal{S} : \Sigma \to \Pi\Sigma'$.

Theorem 1. *A parameter $p : \Sigma \to \mathbb{R}$ is ϵ-estimable under a random sampling map $\mathcal{S} : \Sigma \to \Pi\Sigma'$ only if it is $(2\epsilon, 1 - 2\epsilon)$-uniformly continuous with respect to the metric $\mathcal{S}^* d_1$.*

Proof. Suppose $\widehat{p} : \Sigma' \to \mathbb{R}$ is an ϵ-estimator for p under $\mathcal{S}$. Suppose, for the sake of contradiction, that p is not $(2\epsilon, 1 - 2\epsilon)$-uniformly continuous with respect to $\mathcal{S}^* d_1$, so that there exists $K, L \in \Sigma$ such that $\mathcal{S}^* d_1(K, L) < 1 - 2\epsilon$ and $|p(K) - p(L)| > 2\epsilon$.

By the hypothesis on $\widehat{p}$, we have that

$$\begin{aligned} &\mathbb{P}\left\{|p(K) - \widehat{p}(\mathcal{S}(K))| < \epsilon\right\} > 1 - \epsilon \\ &\mathbb{P}\left\{|p(L) - \widehat{p}(\mathcal{S}(L))| < \epsilon\right\} > 1 - \epsilon. \end{aligned} \tag{5}$$

[3] Two metrics on a set are strongly equivalent if uniform continuity of a function with respect to one metric always implies uniform continuity in the other. For any total variation distance based on a higher-powered norm, the resulting metric is equivalent to $\mathcal{S}^* d_1$, so we can safely ignore all of these cases.

The following definition can be found in any basic textbook on analysis on metric spaces, or perhaps in point-set topology.

Definition 4. *A pseudometric space (X, d) is **compact** if it is closed and bounded. Equivalently, every sequence of points $\{x_j\}_{j=1}^{\infty}$ in X has a subsequence $\{x_{j_k}\}_{k=1}^{\infty}$ that converges to a point $x \in X$.*

We define a notion of continuity that we find convenient for handling functions on pseudometric spaces.

Definition 5. *Let (X, d) be a pseudometric space. A map $p : X \to \mathbb{R}$ is **ϵ-continuous at a point** $x \in X$ if there exists a $\delta > 0$ such that for each $y \in X$, $d(x, y) < \delta$ implies $|p(x) - p(y)| < \epsilon$. If p is ϵ-continuous at all points in X, we say that p is **ϵ-continuous**.*

*If there exists a $\delta > 0$ such that for all $x \in X$, $d(x, y) < \delta$ implies $|p(x) - p(y)| < \epsilon$, then we say that p is (ϵ, δ)-**uniformly continuous**. Note that for pseudometric spaces, it is meaningful to speak of $(\epsilon, 0)$-uniformly continuous functions.*

Put $\Psi_K = \widehat{p}^{-1}(\mathcal{B}_\epsilon(p(K)))$ and $\Psi_L = \widehat{p}^{-1}(\mathcal{B}_\epsilon(p(K)))$[4]. One can see that $\mathcal{B}_\epsilon(p(K)) \cap \mathcal{B}_\epsilon(p(L)) = \emptyset$, so that $\Psi_K \cap \Psi_L = \emptyset$ also holds. Moreover, by (5), we have that $(\mathcal{S}(K))(\Psi_K) > 1 - \epsilon$ and $(\mathcal{S}(L))(\Psi_L) > 1 - \epsilon$, from which one can check that $d_1(\mathcal{S}(K), \mathcal{S}(L)) > 1 - 2\epsilon$. That is to say, $\mathcal{S}^* d_1(K, L) > 1 - 2\epsilon$, thus yielding a contradiction. □

[4] Here, $\mathcal{B}_\epsilon(p(K))$ denotes the open ball of radius ϵ centered about $p(K)$, with a similar description for $\mathcal{B}_\epsilon(p(L))$.

How nice! Given some random sampling model that runs in constant time, we can characterize what graph parameters can be estimated up to some accuracy with high probability. Let us look at some examples next.

Example: Connected Components

We revisit the example of counting connected components, which we characterize in the following proposition.

Proposition 2. *For any $\epsilon > 0$, there exists no constant time sampling algorithm $\mathcal{S}$ on $\mathcal{G}$ such that the number of connected components $p_{CC} : \Sigma \to \mathbb{R}$ is ϵ-estimable under $\mathcal{S}$.*

Proof. Let $\mathcal{S}$ be an arbitrary constant time random sampling map $\mathcal{S} : \Sigma \to \Pi\Sigma'$, so that Σ' is finite. Consider the sequence $\{E_n\}_{n=1}^\infty$ in $\mathcal{G}$ of empty graphs on n nodes. Since $(\mathcal{G}, \mathcal{S}^* d_1)$ is compact, there exists a subsequence $\{E_{n_j}\}_{j=1}^\infty$ that converges to some $X \in \mathcal{G}$ with respect to the metric $\mathcal{S}^* d_1$.

Despite this, the sequence of numbers $\{p_{CC}(E_{n_j})\}_{j=1}^\infty$ is divergent. Therefore, $p_{CC} : \mathcal{G} \to \mathbb{R}$ is not continuous, and is thus not (ϵ, ϵ)-uniformly continuous. By Theorem 1, this implies that p_{CC} is not uniformly estimable under $\mathcal{S}$. Since $\mathcal{S}$ was given arbitrarily, this holds for all constant time random sampling maps $\mathcal{S}$. □

As one can see, the non-estimability of the number of connected components in constant time was reduced down to a statement about convergent sequences, only appealing to an extremely simple sequence of graphs.

Remarks

This short note has discussed some basic ideas regarding the estimation of graph parameters from small subgraphs. In particular, it has demonstrated the idea of using a sampling algorithm to construct a pseudometric on the space of graphs, which can then be used to understand how well parameters can be estimated. In the case where

the parameter to be estimated is the number of connected components, this approach can be used to show that there does not exist a constant-time estimator.

The discussion of any sampling algorithms in particular has been largely omitted: however, one can use approaches of this sort to understand the estimability of certain parameters under particular sampling models. For instance, it has been shown[5] that the normalized Betti numbers for simplicial complexes with bounded degree can be estimated from a constant number of random balls sampled from the complex. For graphs, one might say that the number of connected components divided by the number of nodes is estimable with a constant number of samples, or even that the cardinality of a cycle basis for the graph divided by the number of nodes is estimable as well. This is not too surprising, especially given the chosen counterexample in the proof we constructed. Denoting $p_{CC} : \Sigma \to \mathbb{R}$ the parameter that counts the number of connected components for graphs of bounded degree, we constructed a convergent (with respect to the pseudometric $\mathcal{S}^* d_1$) sequence of empty graphs such that $\{p_{CC}(E_{n_j})\}_{j=1}^{\infty}$ diverges. However, if we take $p_N : \Sigma \to \mathbb{R}$ to be such that $p_N(K) = p_{CC}(K)/|K_0|$, where K_0 is the set of nodes in K, we see that $p_N(E_n) = 1$ for any n. That is to say, the sequence of $\{p_N(E_{n_j})\}_{j=1}^{\infty}$ converges immediately. Not to suggest that the removal of a counterexample constitutes a proof, but this variant gives a taste of what is going on here.

[5] Gábor Elek. Betti numbers are testable. In *Fete of Combinatorics and Computer Science*, pages 139–149. Springer, 2010

There are many other avenues for interesting reading here. One pleasant place to begin is to look at the work of Benjamini and Schramm[6], which talks about sampling "rooted balls" on graphs, and how convergence can be defined in that sense. This is an interesting approach, as convergence in the Benjamini-Schramm sense for bounded graphs can be shown to correspond to weak convergence of the graph spectral distribution, from which many graph parameters can be approximated, or at least bounded.

[6] Itai Benjamini and Oded Schramm. Recurrence of distributional limits of finite planar graphs. In *Selected Works of Oded Schramm*, pages 533–545. Springer, 2011

References

Itai Benjamini and Oded Schramm. Recurrence of distributional limits of finite planar graphs. In *Selected Works of Oded Schramm*, pages 533–545. Springer, 2011.

Gábor Elek. Betti numbers are testable. In *Fete of Combinatorics and Computer Science*, pages 139–149. Springer, 2010.