

What is a Graphon?

Daniel Glasscock, June 2013

These notes complement a talk given for the What is ... ? seminar at the Ohio State University. The block images in this PDF should be sharp; if they appear fuzzy or smoothed out, change the preferences on your PDF reader. Updated: June 24, 2013.

Introduction

Graphons were introduced in 2006 by Lovász and Szegedy as limits of graph sequences. Graphon theory not only draws on graph theory (graphs are special types of graphons), it also employs measure theory, probability, and functional analysis. At only a few years old, the theory is developing quickly and finding new applications. The material for this talk, and most of the notation, was taken exclusively from Lovász's new book [1].

Roughly speaking, the set of finite graphs endowed with the cut metric gives rise to a metric space, and the completion of this space is the space of graphons. *Graphons, then, are limits of graph sequences in the cut metric.* More explicitly, these objects may be realized as symmetric, Lebesgue measurable functions from $[0, 1]^2$ to $[0, 1]$; such functions are weighted graphs on the vertex set $[0, 1]$. The space of graphons with the cut metric is compact, a fact equivalent to strong forms of Szemerédi's Regularity Lemma. Among other things, this compactness guarantees the existence of solutions in the space of graphons to certain problems in extremal graph theory. The aim of this talk is to explore these ideas.

A *graph* G is a set of vertices $V(G)$ (usually, $n = |V(G)|$) and a set of edges $E(G)$ between the vertices. The graphs in this talk will be simple, without loops or multiple edges, and finite unless otherwise specified. Weights (real numbers) will sometimes be given to the edges of a graph to make it an *edge-weighted* graph.

Graph homomorphisms and a motivating problem

We begin with the notions of graph homomorphism numbers and densities. These formalisms play an important role in graphon theory but are useful outside the context of graphons as well.

Definitions Let G and H be graphs. A map φ from $V(H)$ to $V(G)$ is a *homomorphism* if it preserves edge adjacency, that is, if for every edge $\{v, w\}$ in $E(H)$, $\{\varphi(v), \varphi(w)\}$ is an edge in $E(G)$. Denote by $\text{hom}(H, G)$ the number of homomorphisms from H to G . Normalizing by the total number of possible maps, we get the *density* of homomorphisms from H to G ,

$$t(H, G) = \frac{\text{hom}(H, G)}{|V(G)|^{|V(H)|}}.$$

For example, $\text{hom}(\bullet, G) = |V(G)|$, $\text{hom}(\bullet\bullet, G) = 2|E(G)|$, and $\text{hom}(\blacktriangle, G)$ is 6 times the number of triangles in G . In this talk, as in these examples, H will be thought of as much smaller than G , but homomorphism numbers are important the other way, too. For example, if K_q is the complete graph on q vertices, then $\text{hom}(H, K_q)$ is the number of q colorings of H .

We will be more concerned with homomorphism densities. The number $t(H, G)$ lies in $[0, 1]$ and represents the probability that a randomly chosen map from $V(H)$ to $V(G)$ preserves edge adjacency. It also represents (asymptotically as $n = |V(G)| \rightarrow \infty$) the density of H in G . For example, $t(\bullet\bullet, G) = 2|E(G)|/n^2$ while the density of edges in G is $2|E(G)|/n(n-1)$; these two expressions are nearly the same when n is large.

Let's now look at a problem from extremal graph theory as motivation for graphons.

How many 4-cycles must a graph with edge density at least $1/2$ have?

So, suppose G has n vertices and at least $n(n-1)/4$ edges, half as many as are possible. Can you avoid having many 4-cycles? It is an interesting and worthwhile exercise to try to find as many as you can; start with trying to find at least one. It is not hard to see that there are *at most* on the order of n^4 4-cycles (in fact, there are $3\binom{n}{4}$ possible). The following result of Erdős tells us that there must be very many 4-cycles, in fact, on the order of n^4 of them.

Theorem (Erdős) For any graph G ,

$$t(\square, G) \geq t(\bullet\bullet, G)^4.$$

In particular, if $t(\bullet\bullet, G) \geq 1/2$, then $t(\square, G) \geq 1/16$.

In light of the theorem, it would be best to reformulate our problem as follows.

Minimize $t(\square, G)$ over all finite graphs G satisfying $t(\bullet\bullet, G) \geq 1/2$.

It is beneficial at this point to draw an analogy with a problem familiar from elementary calculus.

Minimize $x^3 - 6x$ over all real numbers x satisfying $x \geq 0$.

The minimum here is attained at $x = \sqrt{2}$, which, though our polynomial has rational coefficients, is irrational. The best we can do in the rational numbers is find a sequence limiting to $\sqrt{2}$ at which the polynomial achieves values approaching the minimum. Completing the rational numbers to the real numbers allows us to objectify the limit, which algebra then allows us to realize and work with as $\sqrt{2}$.

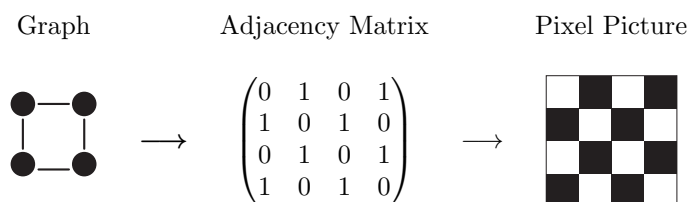
It turns out that we are in an analogous situation with our graph problem. Erdős' theorem tells us that the minimum of $t(\square, G)$ is greater than or equal to $1/16$, and with a little extra work, it can be shown that that minimum is *not* achieved by any finite graph. There is, however, a sequence of finite graphs $(R_n)_n$ with edge density at least $1/2$ and 4-cycle density approaching $1/16$. Indeed, for each $n \geq 1$, let R_n be an instance of a random graph on n vertices where the existence of each possible edge is decided independently with probability $1/2$. By throwing those R_n 's away for which $t(\bullet\bullet, R_n) < 1/2$, the 4-cycle density in the remaining graphs almost surely limits to $1/16$.

The situation is now primed for us to seek to, in pure analogy, complete the space of graphs, realize the limit of $(R_n)_n$ as workable object, and understand the way in which that object achieves the minimum of $1/16$ in our problem above.


Graphons

Let's speculate as to the possible limits of the graph sequence $(R_n)_n$, where R_n is an instance of a random graph with edge probability $1/2$. One real possibility is the *Rado graph*, the random graph with vertex set \mathbb{N} and edge probability $1/2$. (I write "the" random graph since any two instances of such a graph are almost surely isomorphic.) This and many other possible limits are explored in [1] but are not examples of graphons.

Exploring an idea that at first sight is a bit more naive, consider the following three representations of a graph.



The pixel picture is simply a graphic constructed from the adjacency matrix by turning the 1's into black boxes and erasing the 0's. It's important to note that the adjacency matrix and pixel picture correspond to specific labelings of the original graph, and so an unlabeled graph has many different representations as such. Now consider scaling pixel pictures for each of the graphs in the sequence $(R_n)_n$ to the unit square $[0, 1]^2$. Here are example pixel pictures for R_{10} , R_{50} , and R_{100} .




Standing way back, or imagining some sort of averaging process, the limit of this sequence of graphs is suggested by the pixel pictures to be the constant $1/2$ function on $[0, 1]^2$. The constant $1/2$ function on $[0, 1]^2$ is an example of a *graphon* (short for *graph function*). Before defining graphons and making this notion of convergence precise, let's consider some more examples.

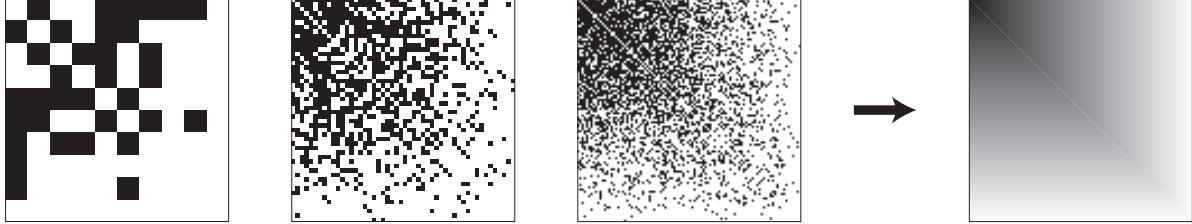
Let T_n be the graph on the vertex set $\{1, 2, \dots, n\}$ where vertices i, j are connected if $i + j \leq n$. Such graphs are called *threshold* graphs. The limit of this sequence of graphs as indicated by their pixel pictures is the graphon taking value 1 on the set $\{x + y \leq 1\}$ and 0 elsewhere. (Since matrixes are indexed with $(0, 0)$ in the top left corner, so too will our unit square be.)



The complete bipartite graph $K_{n,n}$ seems to have two different possible limits, depending on the way it is labeled. The limit of $(K_{n,n})_n$ is in fact unique and is represented by the graphon drawn below. The reader is encouraged to return to this example after digesting the metric determining this convergence.



Finally, consider the following inductively defined sequence of graphs $(G_n)_n$. Let $G_1 = \bullet$. For $n \geq 2$, construct G_n from G_{n-1} by adding one new vertex, then, considering each pair of non-adjacent vertices in turn, drawing an edge between them with probability $1/n$. This is called a *growing uniform attachment* graph sequence, and the pixel pictures below come from one particular instance of a such a sequence. This sequence of graphs almost surely limits to the graphon $1 - \max(x, y)$.



It is finally time to define graphons properly.

Definitions A *labeled graphon* is a symmetric, Lebesgue-measurable function from $[0, 1]^2$ to $[0, 1]$ (modulo the usual identification almost everywhere). An *unlabeled graphon* is a graphon up to relabeling, where a relabeling is given by an invertible, measure preserving transformation of the $[0, 1]$ interval. More formally, a labeled graphon W determines the equivalence class of graphons

$$[W] = \left\{ W^\varphi : (x, y) \mapsto W(\varphi(x), \varphi(y)) \mid \begin{array}{l} \varphi \text{ an invertible, measure} \\ \text{preserving transformation of } [0, 1] \end{array} \right\}.$$

Such equivalence classes are called unlabeled graphons.

It is helpful to think of graphons as edge-weighted graphs on the vertex set $[0, 1]$. In this sense, the sequence $(R_n)_n$ of instances of random graphs with edge probability $1/2$ almost surely limits to the complete graph on a continuum of vertices, each edge with weight $1/2$. Also, note that any graph gives rise to several labeled graphons via its various pixel pictures and that each of these graphons correspond to the same unlabeled graphon.

This viewpoint also allows us to extend homomorphism densities to graphons in an intuitive way. This will allow us to see how the limit of the graph sequence $(R_n)_n$, the constant $1/2$ graphon, solves the minimization problem from the previous section.

For a finite graph G , the value $t(\bullet\!\!\!\rightarrow\!\!\!\bullet, G)$ may be computed by giving each vertex of G a mass of $1/n$ and integrating the edge indicator function over all ordered pairs of vertices. In complete analogy, the edge density of a graphon W is given by the expression

$$t(\bullet\!\!\!\rightarrow\!\!\!\bullet, W) = \int_{[0,1]^2} W(x, y) \, dx dy.$$

It is not hard to see then that

$$t(\square, W) = \int_{[0,1]^4} W(x_1, x_2)W(x_2, x_3)W(x_3, x_4)W(x_4, x_1) \, dx_1 dx_2 dx_3 dx_4.$$

It is straightforward from here to write down the formula for the homomorphism density $t(H, W)$ of a finite graph H into a graphon W .

Finally, in the case of $W \equiv 1/2$ as the limit graphon of $(R_n)_n$, we see that $t(\bullet\!\!\!\rightarrow\!\!\!\bullet, W) = 1/2$ and $t(\square, W) = 1/16$, solving the minimization problem from the previous section elegantly.

The cut distance and convergence of graph sequences

It remains to formalize this “pixel picture” convergence and understand in what sense the space of graphons completes the space of finite graphs. In order to do this, we define a metric on the set of labeled graphons as follows.

Definition The *cut distance* between two labeled graphons W and U is given by

$$\delta_{\square}(W, U) = \inf_{\substack{\varphi, \psi \text{ mpt} \\ \text{of } [0,1]}} \sup_{\substack{S, T \subseteq [0,1] \\ \text{measurable}}} \left| \int_{S \times T} W^{\varphi}(x, y) - U^{\psi}(x, y) \, dx dy \right|.$$

This function measures the maximum discrepancy between the integrals of two labeled graphons over measurable boxes (hence the \square) of $[0, 1]$, then minimizes that maximum discrepancy over all possible relabelings. The infimum makes this function well defined on the space of unlabeled graphons (and could be just as well taken over all single *invertible* measure preserving transformations φ of $[0, 1]$). As an interesting side note, it is known that both the infimum and supremum in the definition above are attained.

Since finite graphs are graphons, it is meaningful to compute the cut distance between two finite graphs. In this case, there is a finite description of the distance as a maximum normalized edge discrepancy over subsets of vertices of blowups of the two graphs in question, minimized over all possible relabelings. The definition is quite involved (see 8.1 of [1]); the point is here that it is possible to define this distance between finite graphs without any analysis.

The infimum/supremum in the definition of the cut distance makes it difficult to do much concretely. It is clear that $\delta_{\square}(W, U) \leq \|W - U\|_1$; if $W \equiv w$ and $U \equiv u$ are constant, then the cut distance is simply $|w - u|$. One tool, which we introduce presently, will simultaneously allow us to show the convergence of the graph sequences above and prove that every graphon is the limit of some sequence of finite graphs.

Definitions A weighted graph H (with edge weights in $[0, 1]$) gives rise to a random simple graph $\mathbb{G}(H)$ by including the edge $\{h_i, h_j\}$ with probability equal to its weight. A graphon W and a finite subset S of $[0, 1]$ give rise to a weighted graph $\mathbb{H}(S, W)$ on $|S|$ nodes by giving the edge $\{s_i, s_j\}$ weight $W(s_i, s_j)$. Let $\mathbb{G}(n, W) = \mathbb{G}(\mathbb{H}(U, W))$ where U is a set of n points chosen uniformly and independently from $[0, 1]$, and let $\mathbb{G}(S, W) = \mathbb{G}(\mathbb{H}(S, W))$.

In other words, for each instance of $\mathbb{G}(n, W)$, a set $S = \{s_1, \dots, s_n\}$ is chosen at random from $[0, 1]$, then a graph is constructed on S where the edge $\{s_i, s_j\}$ is included with probability $W(s_i, s_j)$. The idea is that, with high probability, instances of the graph $\mathbb{G}(n, W)$ will approximate W well for large values of n . This is indeed true, as indicated by the following proposition.

Proposition *Let W be a graphon, and for every $n \geq 1$, let R_n be an instance of $\mathbb{G}(n, W)$. Then with probability 1, $(R_n)_n$ converges to W in the cut distance.*


This proposition tells us that every graphon is a limit of a sequence of finite graphs. The result relies on the subsets U_n being chosen uniformly and independently from $[0, 1]$. If our graphon is (nearly) continuous, then we can remove this reliance on randomness.

Proposition *Let W be a graphon which is continuous almost everywhere, and for every $n \geq 1$, let R_n be an instance of $\mathbb{G}(\{\frac{1}{n}, \dots, \frac{n}{n}\}, W)$. Then with probability 1, $(R_n)_n$ converges to W in the cut distance.*

With a little bit of work, this proposition allows us to prove the convergence of all of the graph sequence limits pictured in the previous section.

The cut metric and the space of graphons

The distance function δ_{\square} is actually only a pseudometric on the space of unlabeled graphons. That means that there is at least one pair of distinct unlabeled graphons which are a δ_{\square} distance of zero apart. Take, for example, the two graphons pictured below; the second is the first under the non-invertible measure preserving transformation $x \mapsto 2x \pmod{1}$.



The fact that δ_{\square} is a pseudometric is generally safely gone unmentioned by implicitly modding out the space of unlabeled graphons by the equivalence relation $W \sim U$ if and only if $\delta_{\square}(W, U) = 0$. It is worth mentioning in this case because those pairs a distance zero apart are exactly the weakly isomorphic pairs. Graphons W and U are *weakly isomorphic* if for every finite graph H , $t(H, W) = t(H, U)$; weak isomorphism is a topic treated in depth in [1].

Let \mathcal{G} be the space of unlabeled graphons, modded out by the equivalence relation above. Then $(\mathcal{G}, \delta_{\square})$ is a metric space. Since every graph is an graphon (by drawing a pixel picture) and every graphon is the limit of a sequence of finite graphs (as explained in the previous section), the space of graphons lies within the completion of the space of finite graphs with the cut metric. The following non-trivial fact is part of what makes graphons such an interesting and exciting tool.

Theorem (Lovász-Szegedy) *The space $(\mathcal{G}, \delta_{\square})$ is compact.*

Since a metric space is compact if and only if it is complete and totally bounded, this means that \mathcal{G} is indeed the completion of the space of finite graphs with the cut metric!

This compactness is equivalent to strong forms of Szemerédi’s Regularity Lemma, a result describing the structure of very large graphs. A weak form of the Regularity Lemma tells us that a graphon is well approximated in the cut metric by its averaged step functions; compactness of \mathcal{G} follows relatively easily from this fact. Conversely, assuming the compactness of \mathcal{G} , a strong form of the Regularity Lemma can be proven with a little effort. In this way, graphons provide a bridge between forms of the Regularity Lemma.

Another interesting fact brings the homomorphism density functions back into play.

Theorem (Lovász-Szegedy) *For any finite graph H , the function $t(H, \cdot) : \mathcal{G} \rightarrow [0, 1]$ is Lipschitz continuous.*

Combined with the fact that \mathcal{G} is compact, this theorem shows that minimization problems of the form discussed at the beginning of this note always have solutions in the space of graphons!

So, for example, it is certain that the minimum of $t(\triangle, W)$ over all graphons W satisfying $t(\bullet\bullet, W) \geq 1/2$ is attained. Graphons W solving this minimization problem provide a “template” for approximate solutions in the finite graphs; indeed, for large n , $\mathbb{G}(n, W)$ is a finite graph which is likely to have triangle density close W ’s (since it will likely be close to W in the cut metric.) Many classical problems from extremal graph theory may be reformulated in these terms and have solutions in the space of graphons.

This is just a baby consequence of the compactness of the space of graphons. The reader interested in more examples, applications, and directions is strongly encouraged to dive into [1].

References

- [1] László Lovász. *Large networks and graph limits*, volume 60 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2012.