Random Graphs — Week 10

Yuval Filmus

January 1, 2020

1 Quasirandom graphs

Let H be a graph on v vertices and e edges. The *density of* H in a graph G is the probability that if we choose v vertices of G at random, then the corresponding edges of H appear in the graph. In a random graph $G(n, 1/2)$, the density of H is $2^{-e} \pm o(1)$ (with high probability). We say that a sequence of graphs G_i is *quasirandom* if the number of vertices of G_i tends to infinity, and for each fixed H, the density of H is $2^{-e(H)} \pm o(1)$.

The standard example of a quasirandom graph is the *Paley graph* P_p, which we have already encountered in the context of Ramsey graphs. Given a prime $p = 4m + 1$, the Paley graph has as vertices $0, \ldots, p-1$, and (i, j) is an edge if $i - j$ is a quadratic residue (has a square root modulo p), which holds exactly when $(i-j)^{2m} \equiv 1 \pmod{p}$ (otherwise, $(i-j)^{2m} \equiv -1 \pmod{p}$). The graph is undirected since $(-1)^{2m} \equiv 1 \pmod{p}$.

It is easy to check that exactly half of the numbers in $1, \ldots, p-1$ are quadratic residues. Given i, the values $i - 0, \ldots, i - (p-1)$, except $i - i$, go (modulo p) over all values in $1, \ldots, p-1$. This shows that the density of edges in P_p is exactly half.

What about other graphs? As an example, let us consider $H = \square$. The problem reduces to counting the number of common neighbors of any two vertices i, j. Let $N(i), N(j)$ denote the neighbors of i, j (respectively). A vertex k is in the symmetric difference of the neighborhoods if $\frac{(i-k)^{2m}}{(j-k)^{2m}} \equiv -1 \pmod{p}$, that is, iff $\frac{i-k}{j-k}$ is a quadratic nonresidue. If $\frac{i-k}{j-k} = \ell$ then $\frac{i-j}{j-k} = \ell - 1$ and so $k = j + \frac{j-i}{\ell-1}$. In particular, each value of ℓ appears at most once, with two missing values: $\ell = 0$ and $\ell = 1$. Roughly $p/2$ of these values are quadratic nonresidues, and so the symmetric difference of $N(i)$ and $N(j)$ contains roughly $p/2$ vertices. Since each neighborhood on its own is of size roughly $p/2$, it follows that the intersection of the neighborhoods has size roughly $p/4$, and so the density of squares is roughly $1/16$ (since given any $i \neq j$, there are roughly $(p/4)^2$ choice k_1, k_2 such that (i, k_1, j, k_2) is a square).

To show that the Paley graphs are quasirandom, we shall use *graphons*.

2 Graphons

Graphons parametrize generalizations of the $G(n, p)$ model. Let $W \colon [0,1]^2 \to [0,1]$ be a measurable symmetric function. We sample a graph according to $G(n, W)$ as follows: choose n random points $x_1, \ldots, x_n \sim U([0,1])$, and add the edge (i, j) with probability $W(x_i, x_j)$. If $W \equiv p$ then this is just the $G(n, p)$ model.

As another example, consider the following function:

$$W(x,y) = \begin{cases} 0 & \text{if } x,y \le 1/2 \text{ or } 1/2 < x,y, \\ 1 & \text{otherwise.} \end{cases}$$

This graphon generates a complete balanced bipartite graph (each side consists of $\mathrm{Bin}(n, 1/2)$ many vertices). If we change 1 to 1/2, we get a random balanced bipartite graph.

What is the expected density of edges in $G(n, W)$? Choosing two random vertices x, y is the same as choosing two random indices $i \neq j$ and then sampling $x_i, x_j \sim U([0, 1])$, which is the same as just sampling $x, y \sim U([0, 1])$. Hence the expected density of edges is

$$d_{-}(W) = \mathop{\mathbb{E}}_{x,y}[W(x, y)].$$

Similarly, the expected density of squares is

$$d_{\square} = \mathop{\mathbb{E}}_{x,y,z,w}[W(x, y)W(y, z)W(z, w)W(w, x)].$$

Routine calculations show that for any graph H, the density of H in $G(n, W)$ is with high probability close to its expectation, up to an error of $o(1)$. Furthermore, if $G_n \sim G(n, W)$, then the density of H in G_n almost surely converges to the expected density.

3 Finite forcing

Suppose that $d_{-}(W) = p$. What can we say about $d_{\square}(W)$?

$$d_{\square} = \mathop{\mathbb{E}}_{z,x}[\mathop{\mathbb{E}}_{y}[W(x, y)W(y, z)] \mathop{\mathbb{E}}_{w}[W(z, w)W(w, x)]].$$

The two factors are actually identical, and so

$$d_{\square} = \mathop{\mathbb{E}}_{z,x}[\mathop{\mathbb{E}}_{y}[W(x, y)W(y, z)]^2] \ge \mathop{\mathbb{E}}_{z,x,y}[W(x, y)W(y, z)]^2,$$

applying the Cauchy–Schwarz inequality (or convexity of $t \mapsto t^2$). Doing the same trick again,

$$\mathop{\mathbb{E}}_{z,x,y}[W(x, y)W(y, z)] = \mathop{\mathbb{E}}_{y}[\mathop{\mathbb{E}}_{x}[W(x, y)]^2] \ge \mathop{\mathbb{E}}_{x,y}[W(x, y)]^2,$$

and so

$$d_{\square} \ge d_{-}^4.$$

If $d_{\square} = d_{-}^4$, then the Cauchy–Schwarz inequality is tight throughout. In particular, for almost all y it is the case that $W(x, y)$ is almost constant, that is, there is a function $f(y)$ such that $W(x, y) = f(y)$ for almost all x, y. Since W is symmetric, for almost all x, y we have $f(y) = W(x, y) = W(y, x) = f(x)$, hence f itself is almost constant. Thus there is a constant C such that $W(x, y) = C$ for almost all x, y. Since $d_{-}(W) = \mathbb{E}[W(x, y)] = C$, we deduce finally that (up to measure zero) $W \equiv p$. In other words, out of all graphons with edge density p, the uniform graphon has the least density of squares.

4 Convergence

Given a graph G on n vertices, we can construct a graphon out of it by dividing $[0,1]^2$ into n^2 squares of size $1/n \times 1/n$, and setting the value of each square to 1 or 0 according to the presence or absence of the corresponding edge (respectively). If $G_n \sim G(n,p)$, then one can show that these graphons W_n converge, almost surely, to the uniform graphon $W \equiv p$, in the following sense:

$$\max_{\substack{0 \le a \le b \le 1 \\ 0 \le c \le d \le 1}} \left| \int_a^b \int_c^d (W_n(x,y) - W(x,y)) dx dy \right| \to 0.$$

This is convergence in the *cut norm*. In general, we allow rearrangements of $[0,1]$ by measure-preserving bijections. Different rearrangements are allowed for the two graphons, and the goal is to minimize the cut norm of the difference. The resulting notion of distance is the *cut metric*. (Rearrangements allow treating two graphons as identical if they only differ by a rearrangement.)

The regularity lemma, a deep result in graph theory, implies that the space of graphons under the cut metric is *compact*. This implies that any sequence of graphons has a convergent subsequence. In particular, if we start with any sequence G_i of graphs, then there is a subsequence whose corresponding graphons W_i converge to some graphon W.

Graph densities are continuous in the cut metric. This means that if $W_i \to W$ then $d_H(W_i) \to d_H(W)$. (The converse is also true: if $d_H(W_i) \to d_H(W)$ for all H, then W_i converges to W.) Furthermore, if W_i is formed from a graph G_i, then $d_H(W_i) = d_H(G_i) + o(1)$, where the error tends to zero with the number of vertices of G_i.

We now put everything together to prove a classic result of Chung, Graham and Wilson. Let G_i be a sequence of graphs such that $d_-(G_i) \to 1/2$ and $d_\square(G_i) \to 1/16$. By compactness, the sequence has at least one limit point. Any limit point W must satisfy $d_-(W) = 1/2$ and $d_\square(W) = 1/16$. As we have seen above, necessarily $W \equiv 1/2$. Thus there is a unique limit point, and so the sequence actually converges to W, implying that it is pseudorandom.

In particular, we can conclude that the sequence of Paley graphs is quasirandom, since we have shown that $d_-(P_p) \to 1/2$ and $d_\square(P_p) \to 1/16$.

5 Other limit objects

Graphons are the appropriate limit objects for dense graphs. The appropriate limit object for sparse graphs are called *graphings*. Instead of describing them, let us describe the corresponding notion of convergence.

Given a graph G, consider the following process. Choose a random vertex $v \in G$, and look at the neighborhood of v at radius r. This gives a random variable ranging over radius-r neighborhoods. A sequence of graphs G_n converges in the sense of Benjamini–Schramm if for each r, the distribution of radius-r neighborhoods converges. For example, finite paths converge to the infinite path, finite grids or tori converge to the infinite grid, and finite random d-regular graphs converge almost surely to the infinite d-regular tree.

The limit object for permutations or distributions over permutations is *permutons*, and is considered in the assignment.