Random Graphs — Assignment 3

Yuval Filmus

January 9, 2020

Question 1 (Alternative proof of zero-one law)**.** The k-round Ehrenfeucht–Fraïssé game is played on a pair of graphs (G_1, G_2) by two players, Spoiler and Duplicator. We require both graphs to contain at least k vertices.

In the first round, Spoiler chooses a vertex on one of the graphs, and then Duplicator chooses a vertex on the other graph. We let a_1 be the vertex chosen in G_1, and b_1 be the vertex chosen in G_2.

In the second round, Spoiler again chooses a vertex on one of the graphs, and Duplicator chooses a vertex on the other graph. Both players are forced to choose vertices not already chosen (different from a_1, b_1). We let a_2, b_2 be the vertices chosen in G_1, G_2 (respectively).

All subsequent rounds proceed in the same way. After k rounds, we end up with k vertices $a_1, \ldots, a_k$ in G_1 and k vertices $b_1, \ldots, b_k$ in G_2. Duplicator's goal is that the following property is satisfied: for all i, j, there is an edge (a_i, a_j) in G_1 iff there is an edge (b_i, b_j) in G_2. We say that *Duplicator wins* if Duplicator has a strategy which guarantees that her goal is fulfilled.

For every first-order formula ϕ in the language of graphs, there is a constant k such that if Duplicator wins (G_1, G_2) then $G_1 \vdash \phi$ (that is, ϕ is satisfied for G_1) iff $G_2 \vdash \phi$.

(a) Show that for each k there is a function $e(n) = o(1)$ such that if $G_1 \sim G(n_1, 1/2)$ and $G_2 \sim G(n_2, 1/2)$ then Duplicator wins (G_1, G_2) with probability at least $1 - e(\min(n_1, n_2))$.

(b) Show that for each ϕ there is a function $e'(n) = o(1)$ such that for each n, either $\Pr[G(n, 1/2) \vdash \phi] \le e'(n)$ or $\Pr[G(n, 1/2) \vdash \phi] \ge 1 - e'(n)$.

(c) Show that for each ϕ, either $\Pr[G(n, 1/2) \vdash \phi] \to 0$ or $\Pr[G(n, 1/2) \vdash \phi] \to 1$.

Question 2 (Failure of zero-one law for colored graphs)**.** A *colored graph* is a graph in which each vertex v has a color $c(v) \in \mathbb{N}$. Given a distribution π on $\mathbb{N}$, let $G_\pi(n, 1/2)$ be the colored graph obtained by coloring each vertex in $G(n, 1/2)$ according to π independently.

We say that two colored graphs G_1, G_2 are *isomorphic* if there is an isomorphism f of graphs between G_1 and G_2 that respects the coloring, that is, $c(v) = c(f(v))$.

(a) Show that if $G_1, G_2 \sim G_\pi(\aleph_0, 1/2)$ then almost surely, G_1 and G_2 are isomorphic (as colored graphs).

(b) Let π have a Poisson distribution with expectation 1: $\Pr[\pi = k] = e^{-1}/k!$. Show that

$$\Pr[G_\pi(k!, 1/2) \text{ contains a color appearing exactly once}] \to e^{-1-e^{-1}}.$$

(c) The *first-order language of colored graphs* is defined similarly to the first-order language of graphs, together with the additional basic predicate $c(x) = c(y)$. Show that the zero-one law doesn't hold for the first-order language of colored graphs with respect to the sequence G_π, where π is the distribution from the preceding item.

Choose one out of Question 3 and Question 4 (bonus if you do both).

Question 3 (Quasirandom permutations[1])**.** The symmetric group S_n consists of all permutations of $[n] := \{1, \ldots, n\}$. We think of permutations as sequences of length n.

For a permutation $\pi \in S_n$ and a "pattern" $\tau \in S_k$, the density $t(\pi, \tau)$ is the probability that if we sample k distinct indices $i_1, \ldots, i_k \in [n]$ then the relative order of $\pi(i_1), \ldots, \pi(i_k)$ is the same as τ. For example,

$$t(13245, 123) = \frac{7}{10}, \quad t(13245, 213) = \frac{2}{10}, \quad t(13245, 132) = \frac{1}{10},$$

since $134, 135, 124, 125, 145, 345, 245$ have relative order 123; $324, 325$ have relative order 213; and 132 has relative order 132.

For each n, let $\pi_n \in S_n$. The sequence $\vec{\pi}$ is *k-quasirandom* if for all $\tau \in S_k$,

$$t(\pi_n, \tau) \to \frac{1}{k!}.$$

The sequence $\vec{\pi}$ is *quasirandom* if it is k-quasirandom for each k. As an example, if π_n is chosen uniformly random permutation for each n, then $\vec{\pi}$ is quasirandom almost surely.

(a) Show that if $\vec{\pi}$ is $(k+1)$-quasirandom then it is k-quasirandom.

(b) Give an example of a 2-quasirandom sequence which is not 3-quasirandom.[2]

A *permuton* is a probability distribution μ over $[0,1]^2$ such that if $(x, y) \sim \mu$ then the marginal distributions of x and y are uniform over $[0,1]$. Given a permuton μ, for each n we can draw a random permutation $\pi \sim P(n, \mu)$ as follows. Let $(x_1, y_1), \ldots, (x_n, y_n)$ be n independent samples of μ. We arrange the x_i in order, and let π consist of the relative order of the y_i. (Since the marginal distributions are uniform over $[0,1]$, almost surely all x_i and all y_i are distinct.) For example, if μ is the uniform distribution over $[0,1]^2$ then $P(n, \mu)$ is a uniformly random permutation in S_n.

For $\tau \in S_k$, let $t(\mu, \tau)$ be the probability that if we take k samples (x_i, y_i) from μ and arrange the x_i in order, then the relative order of the y_i is τ.

[1] After Král' and Pikhurko, Quasirandom permutations are characterized by 4-point densities, GAFA vol. 23, pp. 570–579, 2013.

[2] There are also examples of 3-quasirandom sequences which are not 4-quasirandom, but they are more complicated. One example is described in the paper of Král' and Pikhurko mentioned above, and another one in Cooper and Petrarca, Symmetric and asymptotically symmetric permutations.

(c) Show that $\mathbb{E}_{\pi\sim P(n,\mu)}[t(\pi,\tau)] = t(\mu,\tau)$. (In fact, more is true: if $\pi_n \sim P(n,\mu)$ for each n independently, then almost surely $t(\pi,\tau) \to t(\mu,\tau)$.)

A permuton μ is *k-quasirandom* if for each $\tau \in S_k$, $t(\mu,\tau) = 1/k!$. A permuton μ is quasirandom if it is k-quasirandom for all k. As in the case of individual distributions, it is not hard to show that a $(k+1)$-quasirandom permuton is also k-quasirandom. Furthermore, if μ is a (k-)quasirandom permuton and for each n we sample $\pi_n \sim P(n,\mu)$, then almost surely $\vec{\pi}$ is (k-)quasirandom (where we think of π_n as a constant random variable).

In the rest of this exercise, we show that if μ is a 4-quasirandom permuton, then it is in fact quasirandom (the constant 4 is optimal). This implies that if $\vec{\boldsymbol{\pi}}$ is a 4-quasirandom sequence of random permutations, then it is in fact quasirandom.

(d) Let $F_\mu(X,Y) = \Pr_{(x,y)\sim\mu}[x \le X, y \le Y]$ be the CDF of μ. Show that

$$\mathop{\mathbb{E}}_{(X,Y)\sim\mu}[F_\mu(X,Y)^2] = \Pr_{(x_i,y_i)\sim\mu}[x_1,x_2 \le x_3;\ y_1,y_2 \le y_3].$$

(e) Deduce that $\mathbb{E}_{(X,Y)\sim\mu}[F_\mu(X,Y)^2] = 1/9$, using only the fact that μ is 3-quasirandom.

(f) Show that

$$\mathop{\mathbb{E}}_{(X,Y)\sim\mu}[F_\mu(X,Y)XY] = \Pr_{(x_i,y_i)\sim\mu}[x_1,x_2 \le x_4;\ y_1,y_3 \le y_4].$$

Hint: if $(x,y) \sim \mu$ then since the marginals x and y are uniform over $[0,1]$, then $\Pr[x \le X] = X$ and $\Pr[y \le Y] = Y$.

(g) Deduce that $\mathbb{E}_{(X,Y)\sim\mu}[F_\mu(X,Y)XY] = 1/9$, using the fact that μ is 4-quasirandom.

(h) Let λ be the permuton corresponding to two independent samples of the uniform distribution over $[0,1]$. Show that λ is quasirandom and $F_\lambda(X,Y) = XY$.

(i) Show that

$$\mathop{\mathbb{E}}_{(Z,W)\sim\lambda}[F_\mu(Z,W)^2] = \Pr_{(X_i,Y_i)\sim\mu}[x_1,x_2 \le x_3;\ y_1,y_2 \le y_4].$$

Hint: use two samples of μ to generate one sample of λ.

(j) Deduce that $\mathbb{E}_{(Z,W)\sim\lambda}[F_\mu(Z,W)^2] = 1/9$, using the fact that μ is 4-quasirandom.

(k) Show that

$$\mathop{\mathbb{E}}_{(Z,W)\sim\lambda}[F_\mu(Z,W)ZW] = \Pr_{\substack{(x,y)\sim\mu\\(z_i,w_i)\sim\lambda}}[x,z_1 \le z_2;\ y,w_1 \le w_2] = \frac{1}{4}\mathop{\mathbb{E}}_{(X,Y)\sim\mu}[(1-X^2)(1-Y^2)].$$

(l) Show that

$$A := \mathop{\mathbb{E}}_{(X,Y)\sim\mu}[F_\mu(X,Y)XY]^2 = \frac{1}{81}.$$

(m) Show that

$$B := \mathop{\mathbb{E}}_{(X,Y)\sim\mu}[F_\mu(X,Y)^2] \mathop{\mathbb{E}}_{(X,Y)\sim\mu}[X^2Y^2] = \frac{4}{9} \mathop{\mathbb{E}}_{(Z,W)\sim\lambda}[F_\mu(Z,W)ZW] - \frac{1}{27}.$$

Hint: if $(X,Y) \sim \mu$ then $\mathbb{E}[X^2] = \mathbb{E}[Y^2] = 1/3$ since X, Y are individually uniform over $[0,1]$.

(n) Show that

$$C := \frac{4}{9}\sqrt{\mathop{\mathbb{E}}_{(Z,W)\sim\lambda}[F_\mu(Z,W)^2]}\sqrt{\mathop{\mathbb{E}}_{(Z,W)\sim\lambda}[Z^2W^2]} - \frac{1}{27} = \frac{1}{81}.$$

(o) Explain why always $A \le B \le C$.

(p) Since $A = C = 1/81$, both inequalities are tight. Show that this implies that $F_\mu(XY) = XY$ and so $\mu = \lambda$ is quasirandom.

Question 4 (Dijkstra's algorithm on the uniform weight distribution)**.** In this question we will generalize the analysis of Dijkstra's algorithm from exponential weights to uniform weights. We will use $U([0,1])$ to denote the uniform distribution over $[0,1]$, and $\mathrm{Exp}(1)$ to denote the unit mean exponential distribution, given by $\Pr[\mathrm{Exp}(1) \ge t] = e^{-t}$.

(a) Let $X \sim U([0,1])$ and $Y = \log \frac{1}{1-X}$. Show that $Y \sim \mathrm{Exp}(1)$.

(b) Consider the coupling (X,Y) from the preceding item. Show that

$$Y(1 - Y/2) \le X \le Y.$$

(c) Suppose that w_1, w_2 are two sets of edge weights that satisfy $w_1(e) \le Cw_2(e)$ for all edges e. Let $d_1(x,y), d_2(x,y)$ be the shortest distance from x to y according to the two sets of edge weights. Show that $d_1(x,y) \le Cd_2(x,y)$.

(d) For a distribution $\mathcal{D}$ supported on $\mathbb{R}_+$, let $T_\mathcal{D}$ be the expected distance from vertex 1 to the farthest vertex, when weights are chosen according to $\mathcal{D}$ independently. Show that

$$T_{U([0,1])} \le T_{\mathrm{Exp}(1)}.$$

Proving a bound in the other direction is more tricky. For each edge e, we construct three different weights $w_1(e), w_2(e), w_3(e)$ as follows. We choose $w_1(e) \sim U([0,1])$ and let $w_2(e) = \log \frac{1}{1-w_1(e)}$, so that $w_2(e) \sim \mathrm{Exp}(1)$. For an $\epsilon \in (0, 1/2)$ to be determined, we choose $w_3(e) = \max(w_1(e), (1-\epsilon)w_2(e))$.

(e) We showed in class that $T_{\mathrm{Exp}(1)} \sim \frac{2\log n}{n}$. Let $\delta = \omega\left(\frac{\log n}{n}\right)$. Show that with high probability, the shortest path tree for vertex 1 under edge weights w_1 only contains edges of weight at most δ.

(f) Show that if $\epsilon = \omega\left(\frac{\log n}{n}\right)$ then with high probability, the shortest path trees for vertex 1 under w_1 and under w_3 are identical.

(g) Let τ_1, τ_2, τ_3 be the distances from vertex 1 to the farthest vertex under w_1, w_2, w_3, respectively. We showed in class that $\mathbb{V}[\tau_2] = O(1/n^2)$. Using Cauchy–Schwarz, show that if $\delta = \omega\left(\frac{\log n}{n}\right)$ then

$$\mathbb{E}[\tau_2 1_{\tau_2 > \delta}] = o(\mathbb{E}[\tau_2]).$$

(Here $1_{\tau_2 > \delta}$ is the indicator variable for the event $\tau_2 > \delta$.)

(h) Show that if $\epsilon = \omega\left(\frac{\log n}{n}\right)$ then $T_{U([0,1])} \geq (1 - \epsilon - o(1))T_{\text{Exp}(1)}$, and conclude that $T_{U([0,1])} \sim \frac{2 \log n}{n}$.