Stat 316 Stochastic Processes on Graphs

q-Colorings of Random Graphs

Andrea Montanari *Lecture 13 - 11/5/2007*

In this and the following lectures we shall consider a class of models that are related to random constraint satisfaction problems from computer science. Despite many efforts, and some fascinating conjectures from statistical physics, very little is known rigorously on their behavior. We shall focus mainly on q-coloring and k-satisfiability.

1 Definitions and broad picture

Given a graph $G = (V, E)$, recall that a proper q-coloring of G is an assignment of colors to the vertices of G such that no edge has both end-points of the same color. We shall consider the uniform distribution over proper q-colorings. If $x = \{x_i : in \in V\}$ denotes a q-coloring (here $x_i \in \{1, \dots, q\}$), this takes the form

$$\mu_G(x) = \frac{1}{Z_G} \prod_{(ij)\in E} \mathbb{I}(x_i \neq x_j)\,. \tag{1}$$

The graph will be random over the vertex set $V = [N]$. We have two examples in mind (and will switch between them):

1. G is uniformly random with $M = N\alpha$ edges (and therefore has average degree 2α).
2. G is a random regular graph with degree $(k + 1)$.

Heuristic statistical mechanics studies suggest a rich phase transition structure for the measure $\mu_G(\,\cdot\,)$. For any[1] $q \geq 4$, different regimes are separated by three distinct critical values of the average degree: $0 < \alpha_{\rm d}(q) < \alpha_{\rm c}(q) < \alpha_{\rm s}(q)$. Such regimes can be characterized as follows (all the statements below are understood to hold with high probability with respect to the graph choice):

I. $\alpha < \alpha_{\rm d}(q)$. The model does not undergo a phase transition. Roughly speaking, the set of proper q colorings forms a unique compact lump.

II. $\alpha_{\rm d}(q) < \alpha < \alpha_{\rm c}(q)$. The model is undergoes a phase transition. More precisely, there exists a partition of the space of configurations $\{1, \dots, q\}^V$ into $\mathcal{N}$ sets $\{\Omega_\alpha\}$, such that for any α

$$\frac{\mu(\partial_\epsilon \Omega_\alpha)}{\mu(\Omega_\alpha)[1 - \mu(\Omega_\alpha)]} \leq e^{-C\,N}\,. \tag{2}$$

Further $\mathcal{N} \doteq e^{N\Sigma}$.

III. $\alpha_{\rm c}(q) < \alpha < \alpha_{\rm s}(q)$. The situation is analogous to the previous one, but $\mathcal{N}$ is sub-exponential. More precisely, for any $\delta > 0$, a fraction $1 - \delta$ of the measure μ is comprised in $\mathcal{N}(\delta)$ 'lumps', whereby $\mathcal{N}(\delta)$ converges to a finite random variable.

IV. $\alpha_{\rm s}(q) < \alpha$. A random graph is, with high probability, uncolorable.

Statistical mechanics methods provide semi-explicit expressions for the values $\alpha_{\rm d}(q)$, $\alpha_{\rm c}(q)$, $\alpha_{\rm s}(q)$. Such expressions involve solving an equation for a probability distribution over the $(q - 1)$-dimensional simplex. The thresholds values are given in terms of the solution of this an equation.

[1] The case $q = 3$ is special in that $\alpha_{\rm d}(q) = \alpha_{\rm c}(q)$. The reader is invited to discuss the case $q = 2$.

2 The COL-UNCOL transition: A few simple results

Although the existence of a colorable-uncolorable transition is not established, Friedgut theory allows to make a first step in this direction.

Theorem 1. *Denote by $G_{N,\alpha}$ a uniformly random graph with N vertices and $N\alpha$ edges. For any $q \geq 3$ there exists $\alpha_{\rm s}(q;N)$ such that*

$$\lim_{N\to\infty} \mathbb{P}\{G_{N,\alpha_{\rm s}(q;N)(1-\delta)} \textit{ is q-colorable}\} = 1\,, \tag{3}$$

$$\lim_{N\to\infty} \mathbb{P}\{G_{N,\alpha_{\rm s}(q;N)(1+\delta)} \textit{ is q-colorable}\} = 0\,, \tag{4}$$

In the following we shall drop the N dependence from $\alpha_{\rm s}(q;N)$.

Proposition 2. *The COL-UNCOL threshold is upper bounded as*

$$\alpha_{\rm s}(q) \leq \overline{\alpha}_{\rm s}(q) \equiv \frac{\log q}{\log(1-1/q)}\,. \tag{5}$$

Proof The expected number of proper q-colorings is

$$\mathbb{E}\{Z_G\} = q^N\left(1-\frac{1}{q}\right)^M\,, \tag{6}$$

and $\mathbb{E}\{Z_G\} \to 0$ for $\alpha > \overline{\alpha}_{\rm s}(q)$. The thesis follows from Markov inequality. □

Notice that $\overline{\alpha}_{\rm s}(q) = q\log q[1+o(1)]$ as $q \to \infty$. This asymptotic behavior is known to be tight.

Recall that the k-core of a graph G is the maximal induced subgraph with minimal degree not smaller than k. A simple 'algorithmic' lower bound follows from the next remark.

Proposition 3. *If G does not contain a q-core, then it is q-colorable.*

Proof Given a graph G and a vertex i, denote by $G \setminus \{i\}$ the graph obtained by removing vertex i and all of the adjacent edges. If G does not contain a q-core, then there exists an ordering $i(1), i(2), \dots, i(N)$ of the vertices, such the following is true. If we define G_t by $G_0 = G$ and $G_t = G_{t-1} \setminus i(t)$, then, for any $t \leq N$, $i(t)$ has degree smaller than q in G_{t-1}.

The thesis follows from the observation that if $G \setminus \{i\}$ is q-colorable, and i has degree smaller than q, then G is q-colorable as well. □

The ordinary differential equations method can be used to establish a threshold for the appearance of a q-core in a random graph G.

Proposition 4. *Let $F(\rho;\alpha) = \mathbb{P}\{\text{Poisson}(2\alpha\rho) \geq q-1\}$, and define (for $q \geq 3$)*

$$\alpha_{\rm core}(q) = \sup\{\alpha :\ F(\rho;\alpha) \leq \rho\,\forall\rho \in [0,1]\}\,. \tag{7}$$

Then, with high probability, a uniformly random graph G with $M = N\alpha$ edges over N vertices has a q-core if $\alpha > \alpha_{\rm core}(q)$, and does not have one if $\alpha < \alpha_{\rm core}(q)$.

We omit the proof of this statement (due to Wormald, Spencer and Pittel) since it loosely follows the one we explained for the 2-core of a random hypergraph. The value of $\alpha_{\rm core}(q)$ can be derived by a simple heuristic argument. For a vertex $i \in V$ we call 'q-core induced by i' the largest induced subgraph that has minimum degree not smaller than q except (possibly) at i. Given a random edge (i,j), we denote by ρ the probability that i belongs to the q-core induced by j. It is then natural to write the following equation for ρ:

$$\rho = \mathbb{P}\{\text{Poisson}(2\alpha\rho) \geq q-1\}\,. \tag{8}$$

The threshold $\alpha_{\rm core}(q)$ corresponds to the appearance of a positive solution of this equation.

3 The clustering transition: The conjecture

The conjectured value for $\alpha_{\rm d}(q)$ has a particularly elegant interpretation in terms of phase transition on a tree. Let T be a Galton-Watson tree with Poisson offspring distribution of mean 2α. With an abuse of notation, let μ denote the free boundary measure over proper colorings of T. More explicitly, a proper coloring $x = \{x_i : i \in V\}$ is sampled from μ as follows. First sample the root color uniformly are random. Then recursively, for each colored node i, sample the colors of its direct descendants uniformly at random among the ones different from x_i.

We denote by r the root of T, and by $\overline{\mathsf{B}}(r,t)$ the set of vertices of T whose distance form the root is at least t. finally for any subset of vertices U, we let $\mu_U(\,\cdot\,)$ be the marginal distribution of the corresponding color assignments.

At small α the color at the root decorrelates from the far away ones in $\overline{\mathsf{B}}(r,t)$. At large α they remain correlated at any distance t. The 'reconstruction threshold' separates these two regimes.

Definition 5. *The reconstruction threshold $\alpha_{\rm r}(q)$ is the supremum value of α such that*

$$\lim_{t\to\infty} \mathbb{E}||\mu_{r,\overline{\mathsf{B}}(r,t)}(\,\cdot\,,\,\cdot\,) - \mu_r(\,\cdot\,)\mu_{\overline{\mathsf{B}}(r,t)}(\,\cdot\,)||_{\rm TV} = 0\,. \tag{9}$$

It is conjectured that $\alpha_{\rm d}(q)$ (the clustering threshold on random graphs) and $\alpha_{\rm r}(q)$ (the reconstruction threshold on random trees), do indeed coincide. We shall try to argue in favor of this conjecture in the following.

4 The clustering transition: A physicist's approach

We now present a statistical physics argument to derive the location of the clustering threshold. There are various versions of this argument and not all of them are necessarily equivalent. However, they all predict the same location for the threshold.

In trying to identify the existence of 'lumps,' the major difficulty is that we do not know, a priori, where the lumps are. However if x^* is a configuration sampled from $\mu(\,\cdot\,)$, it will fall inside one such lumps. The idea is to study how a second configuration x behaves when tilted towards the first one.

In practice we fix $x^* = \{x_i^* : i \in V\}$ and study the tilted measure $\mu(\,\cdot\,) = \mu_{G,x^*,\epsilon}(\,\cdot\,)$ defined by

$$\mu^*(x) = \frac{1}{Z}\prod_{(i,j)\in E}\mathbb{I}(x_i \neq x_j)\prod_{i\in V}\psi_\epsilon(x_i^*, x_i)\,. \tag{10}$$

Here ψ_ϵ is a tilting function depending on the continuous parameter ϵ. We assume $\psi_0(x,y) = 1$ identically (therefore μ reduces to the uniform measure over colorings in this case) and $\epsilon > 0$ to favor $x = y$. For instance we might take

$$\psi_\epsilon(x,y) = \exp\Big\{\epsilon\,\mathbb{I}(x=y)\Big\}\,. \tag{11}$$

Studying the above distribution is not an easy task, but we can hope Bethe approximation to work in this case. Messages will depend on the graph but also on x^* and ϵ. Bethe equations read

$$\nu_{i\to j}(x_i) \cong \psi_\epsilon(x_i^*, x_i)\prod_{l\in\partial i\setminus j}\sum_{x_l\neq x_i}\nu_{l\to i}(x_l)\,. \tag{12}$$

Introducing a shorthand for the right-hand side, we will write this equation as

$$\nu_{i\to j} = \mathsf{F}_\epsilon\{\nu_{l\to i} : l \in \partial i\setminus j\}\,. \tag{13}$$

Let us now assume that G is a regular graph of degree $k+1$ and x^* a uniformly random proper coloring of G. Then the message $\nu_{i\to j}$ is itself a random variable, taking values in the $(q-1)$-dimensional simplex.

For $x \in \{1, \dots, q\}$, let us denote by Q_x the conditional distribution of $\nu_{i\to j}$ given that $x_i^* = x$. In formulae, for a subset A of the simplex, we have

$$Q_x(A) \equiv \mathbb{P}\left\{\nu_{i\to j}(\,\cdot\,) \in A \middle| x_i^* = x\right\} . \tag{14}$$

It is then easy to write a recursion for Q, namely

$$Q_x(A) = \sum_{x_1 \dots x_k} \mu(x_1, \dots, x_k | x) \int \mathbb{I}(\mathsf{F}_\epsilon(\nu_1, \dots, \nu_k) \in A) \prod_{i=1}^{k} Q_{x_i}(\mathrm{d}\nu_i)\,. \tag{15}$$

If w no consider a random regular graph, we might hope that $\mu(x_1, \dots, x_k|x)$ converge to the analogous conditional distribution on a tree. If this is the case we obtain a fixed point equation for Q

$$Q_x(A) = \frac{1}{(q-1)^k} \sum_{x_1 \dots x_k \neq x} \int \mathbb{I}(\mathsf{F}_\epsilon(\nu_1, \dots, \nu_k) \in A) \prod_{i=1}^{k} Q_{x_i}(\mathrm{d}\nu_i)\,. \tag{16}$$

It is generally believed (at least by physicists!) that the measure μ undergoes a phase transition if and only if this equation admits a non-degenerate solution.