

The Asymptotics of Large Constrained Graphs

Charles Radin* Kui Ren[†] Lorenzo Sadun[‡]

March 25, 2014

Abstract

We show, through local estimates and simulation, that if one constrains simple graphs by their densities ϵ of edges and τ of triangles, then asymptotically (in the number of vertices) for over 95% of the possible range of those densities there is a well-defined typical graph, and it has a very simple structure: the vertices are decomposed into two subsets V_1 and V_2 of fixed relative size c and $1 - c$, and there are well-defined probabilities of edges, g_{jk} , between $v_j \in V_j$, and $v_k \in V_k$. Furthermore the four parameters c, g_{11}, g_{22} and g_{12} are smooth functions of (ϵ, τ) except at two smooth ‘phase transition’ curves.

1 Introduction

We consider large, simple graphs subject to two constraints: fixed values of the density of edges, ϵ , and the density of triangles, τ , where the densities are normalized so as to have value 1 for a complete graph. Our goal is to obtain a qualitative understanding of the asymptotic structure of such graphs, as the vertex number diverges. (A graph is ‘simple’ if the vertices are labelled, the edges are undirected, there is at most one edge between any two vertices, and there are no edges from a vertex to itself.)

We show that asymptotically there is a ‘unique typical graph’ characterized by a small number of parameters each of which is a function of ϵ and τ , and that the parameters vary smoothly in ϵ and τ except across certain smooth (‘phase transition’) curves. In particular we show that more than 95% of the ‘phase space’ of possible pairs (ϵ, τ) consists of three phases separated by two smooth transition curves, and that within these three regions the typical graph requires at most four parameters for its description. Our strategy, and evidence, is a combination of local estimates and numerical simulation. The two parts are presented in separate sections but they were intertwined in obtaining the results and we do not see how to obtain them otherwise. In particular, in Section 2 and the beginning of Section 3 we

*Department of Mathematics, University of Texas, Austin, TX 78712; radin@math.utexas.edu

[†]Department of Mathematics, University of Texas, Austin, TX 78712; ren@math.utexas.edu

[‡]Department of Mathematics, University of Texas, Austin, TX 78712; sadun@math.utexas.edu

present evidence (not proof) that typical graphs have a very simple structure which we call multipodal. In the remainder of Section 3 we **assume** this multipodal structure and derive two things: the boundary between two of the resultant phases, and the behavior of typical graphs near another phase boundary.

The precise definition of ‘unique typical graph’ requires some technical clarification, given below, for two reasons: to allow for the degeneracy associated with relabelling of vertices, and to use a probabilistic description of graphs appropriate for asymptotics. (This formalism of constrained graphs is natural if one is motivated to obtain the sort of emergent behavior one sees in thermodynamics; the constrained graphs are analogous to the configurations in a microcanonical ensemble of a mean-field statistical mechanical system, with ϵ and τ corresponding to mass and energy densities [RS1, RS2].) We begin with a review of what is known about our constraint problem from extremal graph theory. Important recent results are in [Ra, PR], which are also useful for background on earlier results.




Figure 1: The phase space R , outlined in solid lines. Features are exaggerated for clarity.

Within the parameter space $\{(\epsilon, \tau)\} \subset [0, 1]^2$, for $1/2 \leq \epsilon < 1$ the lowest possible triangle density τ lies above the parabola $\tau = \epsilon(2\epsilon - 1)$, except for $\epsilon_k = (k-1)/k$, $k \geq 2$ when τ_k lies on the parabola and is uniquely achieved by the complete balanced k -partite graph. It is convenient to add two more points to the above sequence $\{(\epsilon_k, \tau_k)\}$, namely $(\epsilon_1, \tau_1) = (0, 0)$ corresponding to the empty graph, and $(\epsilon_\infty, \tau_\infty) = (1, 1)$ corresponding to the complete graph. Furthermore, for ϵ satisfying $\epsilon_{k-1} < \epsilon < \epsilon_k$, $k \geq 1$, the lowest possible triangle density τ lies on a known curve (a ‘scallop’), and the optimal graphs have a known structure [Ra, PR, RS1, RS2]. For all $0 < \epsilon < 1$ the maximization of triangle density τ for given ϵ is simpler than the above minimization: the maximum lies on $\tau = \epsilon^{3/2}$ and is uniquely achieved by a clique on enough vertices to give the appropriate value of ϵ . So extremal graph

theory has determined the shape of the achievable part of the parameter space as the region R in Fig. 1, and determined the structure of the graphs with densities on the boundary.

Our goal is to extend the study to the interior of R in a limited sense: for (ϵ, τ) in the interior we wish to determine, asymptotically in the size of the graph, what most graphs are (or what a typical graph is) with these densities. (We clarify ‘most’ and ‘typical’ below.) The typical graph associated with (ϵ, τ) will be described by probabilities of edges between, and within, a finite number of vertex sets, parameters which vary smoothly in (ϵ, τ) except across phase transition curves. Phases are the maximal connected regions in which the parameters vary smoothly. One transition curve was determined in [RS2], namely $(\epsilon, \tau) = (\epsilon, \epsilon^3)$, $0 \leq \epsilon \leq 1$. On this curve the typical graph corresponds to edges chosen independently with probability ϵ . This (‘Erdős-Rényi’ or ‘ER’) curve separates R into a high τ region, a single phase, and a low τ region consisting of infinitely many phases, at least one related to each scallop. We will present simulation evidence for our determination of the typical graphs of the high τ phase (phase I) and for the two phases, II and III, which are associated with the first scallop; see Fig. 2.




Figure 2: Boundary of phases I, II and III in the phase space. The blue dash-dotted line in the middle is the ER curve, the lower boundary of phase I.

There are two key features of these results: that by use of probability one can capture the structure of most (ϵ, τ) -constrained graphs uniquely; and that their structure is remarkably simple, for instance requiring at most four parameters for over 95% of the phase space.

We conclude this section with some notation which we use to clarify our use above of the terms ‘most’ and ‘typical’ graphs.

Consider simple graphs G with vertex set $V(G)$ of (labeled) vertices, edge set $E(G)$ and triangle set $T(G)$, with cardinality $|V(G)| = n$. Our main tool is $Z_{\epsilon, \tau}^{n, \alpha}$, the number of graphs

with densities:

$$e(G) \equiv \frac{|E(G)|}{\binom{n}{2}} \in (\epsilon - \alpha, \epsilon + \alpha); \quad t(G) \equiv \frac{|T(G)|}{\binom{n}{3}} \in (\tau - \alpha, \tau + \alpha). \quad (1)$$

From this we define the entropy density, the exponential rate of growth of $Z_{e,t}^{n,\alpha}$ as a function of n . First consider

$$s_{\epsilon,\tau}^{n,\alpha} = \frac{\ln(Z_{\epsilon,\tau}^{n,\alpha})}{n^2}, \text{ and then } s(\epsilon, \tau) = \lim_{\alpha \downarrow 0} \lim_{n \rightarrow \infty} s_{\epsilon,\tau}^{n,\alpha}. \quad (2)$$

The double limit defining the entropy density $s(\epsilon, \tau)$ was proven to exist in [RS1]. The objects of interest for us are the qualitative features of $s(\epsilon, \tau)$ in the interior of R . To analyze them we make use of a variational characterization of $s(\epsilon, \tau)$.

Let us first introduce some notation. Assume our vertex set V is decomposed into M subsets: V_1, \dots, V_M . We consider (probabilistic) ‘multipodal graphs’ G_A described by matrices $A = \{A_{ij} \mid i, j = 1, \dots, M\}$ such that there is probability A_{ij} of an edge between any $v_i \in V_i$ and any $v_j \in V_j$. Special cases in which each $A_{ij} \in \{0, 1\}$ allow one to include nonprobabilistic graphs in this setting.

This setting has been further generalized to analyze limits of graphs as $n \rightarrow \infty$. (This work was recently developed in [LS1, LS2, BCLSV, BCL, LS3]; see also the recent book [Lov].) The (symmetric) matrices A_{ij} are replaced by symmetric, measurable functions $g : (x, y) \in [0, 1]^2 \rightarrow g(x, y) \in [0, 1]$; the former are recovered by using a partition of $[0, 1]$ into consecutive subintervals to represent the partition of V into V_1, \dots, V_M . The functions g are called graphons, and using them the following was proven in [RS1] (adapting a proof in [CV]):

Variational Principle. For any possible pair (ϵ, τ) , $s(\epsilon, \tau) = \max[-I(g)]$, where the maximum is over all graphons g with $e(g) = \epsilon$ and $t(g) = \tau$, where

$$e(g) = \int_{[0,1]^2} g(x, y) dx dy, \quad t(g) = \int_{[0,1]^3} g(x, y) g(y, z) g(z, x) dx dy dz \quad (3)$$

and the rate function is

$$I(g) = \int_{[0,1]^2} I_0[g(x, y)] dx dy, \quad (4)$$

with the function $I_0(u) = \frac{1}{2} [u \ln(u) + (1 - u) \ln(1 - u)]$. The existence of a maximizing graphon $g = g_{\epsilon,\tau}$ for any (ϵ, τ) was proven in [RS1], again adapting a proof in [CV].

We want to consider two graphs *equivalent* if they are obtained from one another by relabelling the vertices. There is a generalized version of this for graphons, with the relabelling replaced by measure-preserving maps of $[0, 1]$ into itself [Lov]. The equivalence classes of graphons are called reduced graphons, and on this space there is a natural ‘cut metric’ [Lov].

We can now clarify the notions of ‘most’ and ‘typical’ in the introduction: if $g_{\epsilon,\tau}$ is the only reduced graphon maximizing $s(\epsilon, \tau)$, then as the number n of vertices diverges and $\alpha_n \rightarrow 0$, exponentially most graphs with densities $e(G) \in (\epsilon - \alpha_n, \epsilon + \alpha_n)$ and $t(G) \in (\tau - \alpha_n, \tau + \alpha_n)$ will have reduced graphon close to $g_{\epsilon,\tau}$ [RS1].

2 Numerical Computation

To perform the study we outlined in the previous section, we combine numerical computation with local analysis. Let us first introduce our computational framework.

We need to solve the following constrained minimization problem:

$$\min_g I(g) \tag{5}$$

subject to the constraints

$$e(g) = \epsilon, \quad \text{and} \quad t(g) = \tau, \tag{6}$$

where the functionals $I(g)$, $e(g)$ and $t(g)$ are defined in (4) and (3).

To solve this minimization problem numerically, we need to represent the continuous functions with discrete values. We restrict ourselves to the class of piecewise constant functions.

For each integer $N \geq 1$ let $P_N : 0 = p_0 < p_1 < p_2 < \dots < p_N = 1$ be a partition of the interval $[0, 1]$. We denote by $c_i = p_i - p_{i-1}$ ($1 \leq i \leq N$) the sizes of the subintervals in the partition. Then we can form a partition of the square $[0, 1]^2$ using the (Cartesian) product $P_N \times P_N$. We define the class, $G(P_N)$, of those symmetric functions $g : [0, 1]^2 \rightarrow [0, 1]$ which are piecewise constant on the subsets $P_N \times P_N$, and introduce the notation:

$$g(x, y) = g_{ij}, \quad (x, y) \in (p_{i-1}, p_i) \times (p_{j-1}, p_j), \quad 1 \leq i, j \leq N, \tag{7}$$

with $g_{ij} = g_{ji}$. They are probabilistic generalizations of multipartite graphs, which we call multipodal: bipodal for $N = 2$, tripodal for $N = 3$, etc. We will call such a multipodal graphon ‘symmetric’ if all c_j are equal and all g_{jj} are equal, and otherwise ‘asymmetric’.

It is easy to check that with this type of g , the functionals I , e and t become respectively

$$I(g) = \sum_{1 \leq i, j \leq N} I_0(g_{ij}) c_i c_j = \frac{1}{2} \sum_{1 \leq i, j \leq N} [g_{ij} \ln g_{ij} + (1 - g_{ij}) \ln(1 - g_{ij})] c_i c_j, \tag{8}$$

$$e(g) = \sum_{1 \leq i, j \leq N} g_{ij} c_i c_j, \quad t(g) = \sum_{1 \leq i, j, k \leq N} g_{ij} g_{jk} g_{ki} c_i c_j c_k. \tag{9}$$

Using the fact that $\cup_N G(P_N)$ is dense in the space of graphons, our objective is to solve the minimization problem in $G(P_N)$, for all N :

$$\min_{\{c_j\}_{1 \leq j \leq N}, \{g_{ij}\}_{1 \leq i, j \leq N}} I(g), \tag{10}$$

with the constraints (6) and

$$0 \leq g_{ij}, c_j \leq 1, \quad \sum_{1 \leq j \leq N} c_j = 1, \quad \text{and} \quad g_{ij} = g_{ji} \tag{11}$$

for any given pair of ϵ and τ values.

Minimum without the ϵ and τ constraints. Let us first look at the minimization problem without the $e(g) = \epsilon$ and $t(g) = \tau$ constraints. In this case, we can easily calculate the first-order variations of the objective function:

$$I_{g_{ij}} = I'_0(g_{ij})c_i c_j = \frac{1}{2} \ln[g_{ij}/(1 - g_{ij})]c_i c_j, \quad (12)$$

$$I_{c_p} = \sum_{1 \leq i, j \leq N} I_0(g_{ij}) \left(\frac{\partial c_i}{\partial c_p} c_j + c_i \frac{\partial c_j}{\partial c_p} \right), \quad \text{with} \quad \sum_{1 \leq j \leq N} \frac{\partial c_j}{\partial c_p} = 0. \quad (13)$$

where we used the constraint $\sum_{1 \leq j \leq N} c_j = 1$ to get the last equation.

We observe from (12) that when no constraint is imposed, the minimum of $I(g)$ is located at the constant graphon with $g_{ij} = 1/2$, $1 \leq i, j \leq N$. The value of the minimum is $-(\ln 2)/2 = -0.34657359027997$. At the minimum, $\epsilon = 1/2$ and $\tau = 1/8$ (on the ER curve $\tau = \epsilon^3$). We can check also from (13) that indeed $I_{c_p} = 0$ at the minimum for arbitrary partition $\{c_j\}$ that satisfies $\sum_{1 \leq j \leq N} c_j = 1$. This is obvious since the graphon is constant and thus the partition does not play a role here.

The second-order variation can be calculated as well:

$$I_{g_{ij}g_{lm}} = I''_0(g_{ij})c_i c_j \delta_{il} \delta_{jm} = \frac{1}{2} \frac{c_i c_j}{g_{ij}(1 - g_{ij})} \delta_{il} \delta_{jm}. \quad (14)$$

$$I_{g_{ij}c_p} = I_{c_p g_{ij}} = I'_0(g_{ij}) \left(\frac{\partial c_i}{\partial c_p} c_j + c_i \frac{\partial c_j}{\partial c_p} \right), \quad \text{with} \quad \sum_{1 \leq i \leq N} \frac{\partial c_i}{\partial c_p} = 0. \quad (15)$$

$$I_{c_p c_q} = \sum_{1 \leq i, j \leq N} I_0(g_{ij}) \left(\frac{\partial^2 c_i}{\partial c_p \partial c_q} c_j + \frac{\partial c_i}{\partial c_p} \frac{\partial c_j}{\partial c_q} + \frac{\partial c_i}{\partial c_q} \frac{\partial c_j}{\partial c_p} + c_i \frac{\partial^2 c_j}{\partial c_p \partial c_q} \right),$$

$$\text{with} \quad \sum_{1 \leq i \leq N} \frac{\partial c_i}{\partial c_p} = 0, \quad \sum_{1 \leq i \leq N} \frac{\partial^2 c_i}{\partial c_p \partial c_q} = 0. \quad (16)$$

Thus the unconstrained minimizer is stable with respect to perturbation in g values since the second variation $I_{g_{ij}g_{lm}}$ is positive definite. At the minimizer, however, the second variation with respect to c is zero! This is consistent with what we know already.

We now propose two different algorithms to solve the constrained minimization problem. The first algorithm is based on Monte Carlo sampling while the second algorithm is a variant of Newton's method for nonlinear minimization. All the numerical results we present later are obtained with the sampling algorithm and confirmed with the Newton's algorithm; see more discussion in Section 2.3.

2.1 Sampling Algorithm

In the sampling algorithm, we construct random samples of values of I in the parameter space. We then take the minimum of the sampled values. The algorithm works as follows.

[0] Set the number N and the number of samples to be constructed (L); Set counter $\ell = 1$;

[1] Generate the sizes of the blocks: $\{c_j^\ell\} \in [0, 1]^N$ and normalize so that $\sum_j c_j^\ell = 1$;

[2] Generate a sample $\{g_{ij}^\ell\} \in [0, 1]^{N \times N}$ with $g_{ij}^\ell = g_{ji}^\ell$, $1 \leq i, j \leq N$ and rescale such that:

$$[A] \quad \sum_{1 \leq i, j \leq N} g_{ij}^\ell c_i^\ell c_j^\ell = \epsilon;$$

$$[B] \quad \sum_{1 \leq i, j, k \leq N} g_{ij}^\ell g_{jk}^\ell g_{ki}^\ell c_i^\ell c_j^\ell c_k^\ell = \tau;$$

[3] Evaluate $I_\ell = \sum_{1 \leq i, j \leq N} I_0(g_{ij}^\ell) c_i^\ell c_j^\ell$;

[4] Set $\ell = \ell + 1$; Go back to [1] if $\ell \leq L$;

[5] Evaluate $I_N = \min\{I_\ell\}_{\ell=1}^L$.

We consider two versions of Step [2] which work equally well (beside a slight difference in computational cost) in practice. In the first version, we generate a sample $\{g_{ij}^\ell\} \in [0, 1]^{N \times N}$ with $g_{ij}^\ell = g_{ji}^\ell$, $1 \leq i, j \leq N$. We then rescale the sample using the relation $g_{ij}^\ell = \gamma g_{ij}^\ell + \tilde{\gamma}$. The relations [A] and [B] then give two equations for the parameter γ and $\tilde{\gamma}$. We solve the equations for γ and $\tilde{\gamma}$. If at least one of $\{g_{ij}^\ell\}$ violate the condition $g_{ij}^\ell \in [0, 1]$ after rescaling, we re-generate a sample and repeat the process until we find a sample that satisfies $g_{ij}^\ell \in [0, 1]$ ($1 \leq i, j \leq N$) after rescaling. In the second version, we simply find a sample $\{g_{ij}^\ell\}$ by solving [A] and [B] as a third-order algebraic system using a multivariate root-finding algorithm, with the constraint that $\{g_{ij}^\ell\} \in [0, 1]^{N \times N}$. When multiple roots are found, we take them as different qualified samples.

This sampling algorithm is a global method in the sense that the algorithm will find a good approximation to the global minimum of the functional I when sufficient samples are constructed. The algorithm will not be trapped in a local minimum. The algorithm is computationally expensive. However, it can be parallelized in a straightforward way. We will discuss the issues of accuracy and computational cost in Section 2.3.

2.2 SQP Algorithm

The second algorithm we used to solve the minimization problem is a sequential quadratic programming (SQP) method for constrained optimization. To briefly describe the algorithm, let us denote the unknown by $x = (c_1, \dots, c_N, g_{11}, \dots, g_{1N}, g_{21}, \dots, g_{2N}, \dots, g_{N1}, \dots, g_{NN})$. Following [GMSW], we rewrite the optimization problem as

$$\min_{x \in [0, 1]^{N+N^2}} I(x), \quad \text{subject to} \quad l \leq r(x) \leq u, \quad r(x) = \begin{pmatrix} x \\ \mathcal{A}x \\ \mathcal{C}(x) \end{pmatrix} \quad (17)$$

where l and u are the lower and upper bounds of $r(x)$ respectively, $\mathcal{C}_1(x) = e(x)$ and $\mathcal{C}_2(x) = t(x)$. The matrix \mathcal{A} is used to represent the last two linear constraints in (11):

$$\mathcal{A} = \begin{pmatrix} \Sigma & \mathbf{0}_{1 \times N^2} \\ \mathbf{0}_{N^2 \times N} & \mathcal{S} \end{pmatrix} \quad (18)$$

where $\Sigma = (1, \dots, 1) \in \mathbb{R}^{1 \times N}$ is used to represent the constraint $\sum c_j = 1$, $\mathbf{0}_{1 \times N^2} \in \mathbb{R}^{1 \times N^2}$ is a matrix with all zero elements, and $\mathcal{S} \in \mathbb{R}^{N^2 \times N^2}$ is a matrix used to represent the symmetry constraint $g_{ij} = g_{ji}$. The elements of \mathcal{S} are given as follows. For any index k , we define the conjugate index k' as $k' = (r-1)N + q$ with q and r the unique integers such that $k = (q-1)N + r$ ($0 \leq r < N$). Then for all $1 \leq k \leq N^2$, $\mathcal{S}_{kk} = -\mathcal{S}_{kk'} = 1$ if $k \neq (j-1)N + 1$ for some j , and $\mathcal{S}_{kk} = 0$ if $k \neq (j-1)N + 1$ for some j . All other elements of \mathcal{S} are zero.

The SQP algorithm is characterized by the following iteration

$$x_{\ell+1} = x_\ell + \alpha_\ell p_\ell, \quad \ell \geq 0 \quad (19)$$

where p_ℓ is the search direction of the algorithm at step ℓ and α_ℓ is the step length. The search direction p_ℓ in SQP is obtained by solving the following constrained quadratic problem

$$\min I(x_\ell) + \mathbf{g}(x_\ell)^\top p_\ell + \frac{1}{2} p_\ell^\top H_\ell p_\ell, \quad \text{subject to} \quad l \leq r(x_\ell) + J(x_\ell) p_\ell \leq u. \quad (20)$$

Here $\mathbf{g}(x_\ell) = \nabla_x I$, whose components are given analytically in (12) and (13), $J(x_\ell) = \nabla_x r(x_\ell)$, whose components are given by

$$e_{g_{ij}} = c_i^\ell c_j^\ell, \quad (21)$$

$$t_{g_{ij}} = 3 \sum_{1 \leq k \leq N} g_{jk}^\ell g_{ki}^\ell c_i^\ell c_j^\ell c_k^\ell, \quad (22)$$

$$e_{c_p} = \sum_{1 \leq i, j \leq N} g_{ij}^\ell \left(\frac{\partial c_i^\ell}{\partial c_p^\ell} c_j^\ell + c_i^\ell \frac{\partial c_j^\ell}{\partial c_p^\ell} \right), \quad \text{with} \quad \sum_{1 \leq i \leq N} \frac{\partial c_i^\ell}{\partial c_p^\ell} = 0 \quad (23)$$

$$t_{c_p} = \sum_{1 \leq i, j, k \leq N} g_{ij}^\ell g_{jk}^\ell g_{ki}^\ell \left(\frac{\partial c_i^\ell}{\partial c_p^\ell} c_j^\ell c_k^\ell + \frac{\partial c_j^\ell}{\partial c_p^\ell} c_i^\ell c_k^\ell + \frac{\partial c_k^\ell}{\partial c_p^\ell} c_i^\ell c_j^\ell \right). \quad (24)$$

$H(x_\ell)$ is a positive-definite quasi-Newton approximation to the Hessian of the objective function. We take the BFGS updating rule to form $H(x_\ell)$ starting from the identity matrix at the initial step [NW].

We implemented the SQP minimization algorithm using the software package given in [GMSW] which we benchmarked with the `fmincon` package in MATLAB R2012b.

2.3 Computational Strategy

The objective of our calculation is to minimize the rate function I for a fixed (ϵ, τ) pair over the space of all graphons. Our computational strategy is to first minimize for $g \in G(P_N)$,

for a fixed number of blocks N , and then minimize over the number of blocks. Let I_N be the minimum achieved by the graphon $g_N \in G(P_N)$, then the minimum of the original problem is $I = \min_N \{I_N\}$. Due to limitations on computational power we can only solve up to $N = 16$. The algorithms, however, are not limited by this.

By construction we know that $I_2 \geq I_3 \geq \dots \geq I_N$. What is surprising is that our computation suggests that the minimum is always achieved with bipodal graphons in the phases that we are considering. In other words, $I_2 = \min\{I_N\}_{N=2}^{N=16}$.

To find the minimizers of I_N for a fixed N , we run both the sampling algorithm and the SQP algorithm. We observe that both algorithms give the same results (to precision 10^{-6}) in all cases we have simulated. In the sampling algorithm, we observe from (12) and (13) that the changes in I caused by perturbations in g_{ij} and c_p are given respectively by $\delta I / \delta g_{ij} \sim \ln[\delta g_{ij} / (1 - \delta g_{ij})] c_i c_j$ and $\delta I / \delta c_p \sim I_0(g_{ij}) \delta c_p$. These mean that to get an accuracy of order η , our samples have to cover a grid of parameters with mesh size $\delta g_{ij} \sim e^{2\eta/c_i c_j} / (1 + e^{2\eta/c_i c_j}) \sim 1$ in the g_{ij} direction and $\delta c_p \sim \eta / I_0(g_{ij})$ in the c_p direction. Since $I_0(g_{ij}) \in [-(\ln 2)/2, 0]$, we have $\eta / I_0(g_{ij}) > \eta$. It is thus enough to sample on a grid of size η . Similar analysis following (21), (22), (23) and (24) shows that, to achieve an accuracy η on the ϵ and τ constraints, we need to sample on grids of size at most on the order of η in both the g_{ij} and c_p directions. The total computational complexity is thus $\sim (1/\eta)^{N+N^2}$ in terms of function evaluations.

To run the SQP algorithm, for each (ϵ, τ) constraint, we start from a collection of $L_c^N \times L_g^{N^2}$ initial guesses. These initial guesses are generated on the uniform grid of L_c intervals in each c_p direction and L_g intervals in each g_{ij} direction. They are then rescaled linearly (if necessary) to satisfy the ϵ and τ constraints. The results of the algorithm after convergence are collected to be compared with the sampling algorithm. We observe that in most cases, the algorithm converges to identical results starting from different initial guesses.

Finally we note that there are not many parameters we need to tune to get the results that we need. We observe that the SQP algorithm is very robust using the general default algorithmic parameters. The only parameters that we can adjust are L_c and L_g which control how many initial guesses we want to run. Our calculation shows that $L_c = L_g = 10$ is enough for all the cases we studied. When we increased L_c and L_g to get more initial guesses, we did not gain any new minimizers.

2.4 Benchmarking the Computations

Before using the computational algorithms to explore the regions of the phase space that we plan to explore, we first benchmark our codes by reproducing some theoretically known results.

On the upper boundary. We first reproduce minimizing graphons on the curve $\tau = \epsilon^{\frac{3}{2}}$. It is known [RS2] that the minimizing graphons are equivalent to the bipodal graphon with $c = \sqrt{\epsilon}$, $g_{11} = 1$ and $g_{12} = g_{21} = g_{22} = 0$. The minimum value of the rate function is $I = 0$. In Fig. 3 we show the difference between the simulated minimizing graphons and the

true minimizing graphons given by the theory. We observe that the difference is always well below 10^{-6} for all the results on the curve.




Figure 3: Difference between numerically computed (with superscript *num*) and true (with superscript *true*) minimizing graphons on the curve $\tau = \epsilon^{\frac{3}{2}}$. From left to right: $c^{true} - c^{num}$, $g_{11}^{true} - g_{11}^{num}$, $g_{22}^{true} - g_{22}^{num}$ and $g_{12}^{true} - g_{12}^{num}$.

On the segment $(\epsilon, \tau) = (0, 0.5) \times \{0\}$. It is known that on the segment $(\epsilon, \tau) = (0, 0.5) \times \{0\}$ the minimizing graphons are symmetric bipodal with $g_{11} = g_{22} = 0$ and $g_{12} = g_{21} = 2\epsilon$.

In Fig. 4 we show the difference between the simulated minimizing graphons and the true minimizing graphons given by the theory [RS2]. The differences are again very small, $< 10^{-6}$. This shows again that our numerical computations are fairly accurate.




Figure 4: Difference between numerically computed (with superscript *num*) and true (with superscript *true*) minimizing graphons on the line segment $(\epsilon, \tau) = (0, 0.5) \times \{0\}$. From left to right: $c^{true} - c^{num}$, $g_{11}^{true} - g_{11}^{num}$, $g_{22}^{true} - g_{22}^{num}$ and $g_{12}^{true} - g_{12}^{num}$.

On the segment $(\epsilon, \tau) = \{0.5\} \times (0, 0.5^3)$. It is again known from the theory in [RS2] that on this segment the minimizing graphons are symmetric bipodal with $g_{11} = g_{22} = 0.5 - (0.5^3 - \tau)$ and $g_{12} = g_{21} = 0.5 + (0.5^3 - \tau)$. In Fig. 5 we show the difference between the simulated minimizing graphons and the true minimizing graphons. The last point on this segment is $(0.5, 0.125)$. This is the point where the global minimum of the rate function is achieved. Our algorithms produce a minimizer that gives the value $I = -0.34657359$ which is less than 10^{-8} away from the true minimum value of $-(\ln 2)/2$.




Figure 5: Difference between numerically computed (with superscript *num*) and true (with superscript *true*) minimizing graphons on the line segment $(\epsilon, \tau) = \{0.5\} \times (0, 0.5^3)$. From left to right: $c^{true} - c^{num}$, $g_{11}^{true} - g_{11}^{num}$, $g_{22}^{true} - g_{22}^{num}$ and $g_{12}^{true} - g_{12}^{num}$.

Symmetric bipodal graphons. In Section 3.3 we find a formula for the optimizing graphons for phase II, namely the following symmetric bipodal graphons:

$$g(x, y) = \begin{cases} \epsilon - (\epsilon^3 - \tau)^{1/3} & x, y < 1/2 \text{ or } x, y > 1/2 \\ \epsilon + (\epsilon^3 - \tau)^{1/3} & x < \frac{1}{2} < y \text{ or } y < \frac{1}{2} < x \end{cases} \quad (25)$$

Our computational algorithms also find these minimizing symmetric bipodal graphons; see discussions in the next section.

2.5 Numerical Experiments




Figure 6: The minimal value of I at different (ϵ, τ) values.

We now present numerical explorations of the minimizing graphons in some subregions of the phase space. We focus on two main subregions. The first is the subregion above the

ER curve and below the upper boundary $\tau = \epsilon^{3/2}$. This is the phase I in Fig. 2. The second subregion is the region below the ER curve and above the lower boundary of the region where bipodal graphons exist. We further split this region into two separate phases, II and III; see Fig. 2. Phase II is the region where the minimizing graphons are symmetric bipodal while in phase III asymmetric bipodal graphons are the minimizers. We first show in Fig. 6

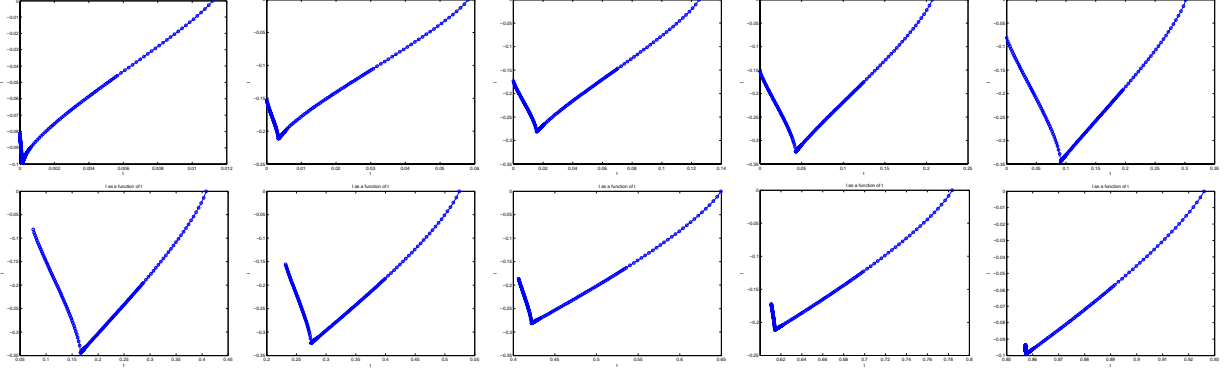


Figure 7: Cross-sections of the minimal value of the rate function I as a function of τ along lines $\epsilon = a_k$ ($k = 1, \dots, 10$).

the minimal values of I in the whole region in which we are interested: $I \cup II \cup III$. The minimum of I for fixed ϵ is achieved on the ER curve. The cross-sections in Fig. 7 along the lines $\epsilon = a_k = (k - 1) * 0.1 + 0.05$ give a better visualization of the landscape of the I .

The main feature of the computational results is that the minimal values of the rate function I in the region we show in Fig. 6 are achieved by bipodal graphons. In other words, the minimizing graphons in the whole subregion are bipodal. On the ER curve, the graphons are constant. Thus the partition does not matter anymore. We consider them as bipodal purely from the continuity perspective. We show in Fig. 8 minimizing graphons at some typical points in the phase space. Shown are minimizing graphons at $\epsilon = 0.3$ (left column), $\epsilon = 0.6$ (middle column) and $\epsilon = 0.8$ (right column) respectively. For the $\epsilon = 0.3$ (resp. $\epsilon = 0.6$ and $\epsilon = 0.8$) column, the τ values for the corresponding graphons from top to bottom are $\tau = 0.09565838$ (resp. $\tau = 0.34037900$ and $\tau = 0.61784171$) which is in the middle of phase I, $\tau = 0.03070755$ (resp. $\tau = 0.22010451$ and $\tau = 0.55677919$) which is just above the ER curve, $\tau = 0.02700000$ (resp. $\tau = 0.21600000$ and $\tau = 0.51200000$) which is on the ER curve, $\tau = 0.02646000$ (resp. $\tau = 0.21472000$ and $\tau = 0.51104000$) which is just below the ER curve, and $\tau = 0.01296000$ (resp. $\tau = 0.18400000$ and $\tau = 0.50800000$) which is in the middle of phase II (resp. phase II and phase III) respectively.

A glance at the graphons in Fig. 8 gives the impression that the minimizing graphons are very different at different values of (ϵ, τ) . A prominent feature is that the sizes of the two blocks vary dramatically. Since there is only one parameter that controls the sizes of the blocks, that is, if the first block has size $c_1 = c$ then the second block has size $c_2 = 1 - c$, we can easily visualize the change of the sizes of the two blocks in the minimizing graphons in the phase space through c , the size of the smaller block as a function of (ϵ, τ) . We show in Fig. 9 the minimal c values associated with the minimizing bipodal graphons. The cross-sections along the lines $\epsilon = a_k$ ($k = 1, \dots, 10$) are shown in Fig. 10.




Figure 8: Minimizing graphons at $\epsilon = 0.3$ (left column), $\epsilon = 0.6$ (middle column) and $\epsilon = 0.8$ (right column). For each column τ values decrease from top to bottom. The second row represents points just above the ER curve, the third row points on the ER curve, and the fourth row points just below the ER curve. The exact τ values are given in Section 2.5.

Among the minimizing bipodal graphons some are symmetric and the others are asymmetric. We show in Fig. 11 the region where the minimizing bipodal graphons are symmetric (brown) and the region where they are asymmetric (blue). This is done by checking the conditions $c_1 = c_2 = 0.5$ and $g_{11} = g_{22}$ (both with accuracy up to 10^{-7}) for symmetric bipodal graphons. Our numerical computation here agrees with the theoretical results we obtain later in Section 3.3.




Figure 9: The minimal c values of the minimizing bipodal graphons as a function of (ϵ, τ) .

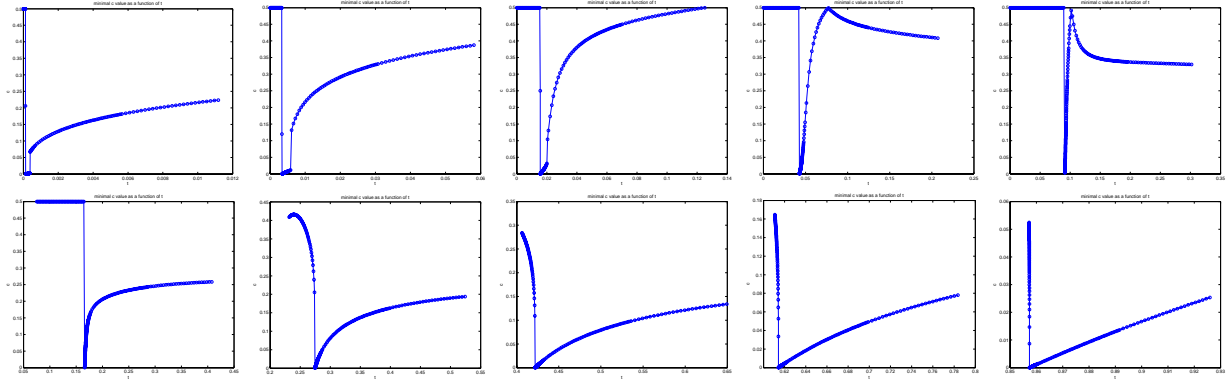


Figure 10: Cross-sections of the function $c(\epsilon, \tau)$ along lines of $\epsilon = 0.05, \epsilon = 0.15, \dots, \epsilon = 0.95$. The 4th and 5th graphs display coordinate singularities as the size of a cluster passes through 0.5, since we have chosen c to always be 0.5 or less. The right-most kinks in those graphs are not phase transitions.

3 Local Analysis

In this section we do a perturbative analysis of all three phases near the ER curve. This gives a qualitative explanation for why these three phases appear in the form they do, and we compute exactly the boundary between phases II and III, assuming the multipodal structure of the phases.




Figure 11: The regions of symmetric bipodal and asymmetric bipodal optima.

3.1 General Considerations

We treat the graphon $g(x, y)$ as the integral kernel of an operator on $L^2([0, 1])$. Let $|\phi_1\rangle \in L^2([0, 1])$ be the constant function $\phi_1(x) = 1$. Then the edge density is $e(g) = \epsilon = \langle \phi_1 | g | \phi_1 \rangle$ and the triangle density is $t(g) = \tau = \text{Tr}(g^3)$. On the ER curve the optimal graphon is $g_0 = \epsilon |\phi_1\rangle \langle \phi_1|$. Near the ER curve we take $g = g_0 + \delta g$, where $\langle \phi_1 | \delta g | \phi_1 \rangle = 0$. For fixed ϵ , the rate function is minimized at the ER curve, where $\tau = \epsilon^3$. For nearby values of $t(g) = \tau = \epsilon^3 + \delta\tau$, minimizing the rate function involves solving two simpler optimization problems:

1. We want to minimize the size of δg , as measured in the L^2 norm, for a given $\delta\tau$.
2. We want to choose the form of δg that minimizes δI for a given $\|\delta g\|_{L^2}^2 = \iint \delta g(x, y)^2 dx dy$.

If we could solve both optimization problems simultaneously, we would have a rigorously derived optimum graphon. This is indeed what happens when $\epsilon = 1/2$ and $\tau < \epsilon^3$ [RS2]. For other values of (ϵ, τ) , the solutions to the two optimizations disagree somewhat. In phases I and III, the actual best graphon appears to be a compromise between the two optima, while in phase II it appears to be a solution to the first optimization problem, but of a form suggested by the second problem.

By playing the two problems off of one another, we derive candidates for the optimal graphon and gain insight into the numerical results. By measuring the extent to which these candidates fail to solve each problem, we can derive rigorous estimates on the behavior of the entropy. However, we do not claim to prove that our bipodal ansatz is correct.

A simple expansion of $g^3 = (g_0 + \delta g)^3$ shows that

$$\delta\tau = 3\epsilon\langle\phi_1|\delta g^2|\phi_1\rangle + \text{Tr}(\delta g^3). \quad (26)$$

The first term is positive-definite, while the second is indefinite.

For $\delta\tau$ negative, the solution to the first optimization problem is to have $\delta g = \nu|\psi\rangle\langle\psi|$, where ψ is an arbitrary normalized vector in $L^2([0, 1])$ such that $\langle\phi_1|\psi\rangle = 0$. This eliminates the positive-definite term and makes the indefinite term as negative as possible, namely equal to ν^3 .

For $\delta\tau$ positive, the solution to the first optimization problem is of the form

$$\delta g = \frac{\mu}{\sqrt{2}}|\psi\rangle\langle\phi_1| + \frac{\mu}{\sqrt{2}}|\phi_1\rangle\langle\psi| + \nu|\psi\rangle\langle\psi|, \quad (27)$$

with $\langle\phi_1|\psi\rangle = 0$. We then have

$$\|\delta g\|_{L^2}^2 = \mu^2 + \nu^2; \quad \delta\tau = \nu^3 + \frac{3\mu^2}{2}(\epsilon + \nu). \quad (28)$$

Maximizing $\delta\tau$ for fixed $\mu^2 + \nu^2$ is then a Lagrange multiplier problem in two variables with one constraint. There are minima for $\delta\tau$ at $\mu = 0$ and maxima at $\nu = \mu^2/(2\epsilon)$.

Note that the form of the vector ψ did not enter into this calculation. The form of ψ is determined **entirely** from considerations of δI *versus* $\|\delta g\|_{L^2}$. We now turn to this part of the problem.

The rate function $I_0(u)$ is an even function of $u - 1/2$. Furthermore, all even derivatives of this function at $u = 1/2$ are positive, so the function can be written as a power series $I_0(u) = -(\ln 2)/2 + \sum_{n=1}^{\infty} c_n (u - 1/2)^{2n}$, where the coefficients c_n are all positive. As a function of $(u - 1/2)^2$, $I_0(u)$ is then concave up. Thus the way to minimize $I(g) = \iint I[g(x, y)] dx dy$ for fixed $\iint (g(x, y) - \frac{1}{2})^2 dx dy$ is to have $[g(x, y) - 1/2]^2$ constant; in other words to have $g(x, y)$ only take on two values, and for those values to sum to 1.

This is exactly our second minimization problem. We have a fixed value of $\|\delta g\|_{L^2}^2$, and hence a fixed value of

$$\begin{aligned} \iint \left(g(x, y) - \frac{1}{2}\right)^2 dx dy &= \iint \left(\epsilon - \frac{1}{2} + \delta g(x, y)\right)^2 dx dy \\ &= \left(\epsilon - \frac{1}{2}\right)^2 + 2\left(\epsilon - \frac{1}{2}\right) \iint \delta g(x, y) dx dy + \iint \delta g(x, y)^2 dx dy \\ &= \left(\epsilon - \frac{1}{2}\right)^2 + \|\delta g\|_{L^2}^2. \end{aligned} \quad (29)$$

In order for $g(x, y)$ to only take on two values, we need ψ to only take on two values. After doing a measure-preserving transformation of $[0, 1]$, we can assume that

$$\psi(x) = \psi_c(x) := \begin{cases} \sqrt{\frac{1-c}{c}} & x < c \\ -\sqrt{\frac{c}{1-c}} & x > c \end{cases} \quad (30)$$

for some constant $c \leq \frac{1}{2}$.

We henceforth restrict our attention to the ansatz (27) with the function (30), and vary the parameters μ , ν and c to minimize $I(g)$ while preserving the fixed values of ϵ and τ .

3.2 Phase I: $\tau > \epsilon^3$

Suppose that τ is slightly greater than ϵ^3 . By our previous analysis, we want δg to be small in an L^2 sense. Maximizing $\delta\tau$ for fixed $\|\delta g\|_{L^2}$ means taking $\nu = \mu^2/(2\epsilon) \ll \mu$, so

$$\delta\tau = \frac{3\mu^2\epsilon}{2} + \frac{3\mu^4}{4} + \frac{\mu^6}{8\epsilon^3}; \quad \|\delta g\|_{L^2}^2 = \mu^2 + \frac{\mu^4}{4\epsilon^2}. \quad (31)$$

If $\epsilon \neq 1/2$, then there is no way to make $|g(x, y) - 1/2|$ constant while keeping δg pointwise small. Instead, we take $\delta g(x, y)$ to be large in a small region. We compute

$$g(x, y) = \begin{cases} \epsilon - \mu\sqrt{2}\sqrt{\frac{c}{1-c}} + \nu\frac{c}{1-c} & x, y > c \\ \epsilon + \frac{\mu}{\sqrt{2}}\left(\sqrt{\frac{1-c}{c}} - \sqrt{\frac{c}{1-c}}\right) - \nu & x < c < y \text{ or } y < c < x \\ \epsilon + \mu\sqrt{2}\sqrt{\frac{1-c}{c}} + \nu\frac{1-c}{c} & x, y < c \end{cases} \quad (32)$$

By taking $c = \mu^2/2(2\epsilon - 1)^2 + O(\mu^4)$, we can get the values of g in the rectangles $x < c < y$ (or $y < c < x$) and the large square $x, y > c$ to sum to 1. This makes $I_0[g(x, y)]$ constant except on a small square of area $c^2 = O(\mu^4)$. Since $g(x, y)$ is equal to $\epsilon - \mu^2/(1 - 2\epsilon) + O(\mu^4)$ when $x, y > c$, we have

$$\begin{aligned} I(g) &= I_0\left(\epsilon - \frac{\mu^2}{1 - 2\epsilon}\right) + O(\mu^4) \\ &= I_0(\epsilon) - I'_0(\epsilon)\frac{\mu^2}{1 - 2\epsilon} + O(\mu^4) \\ &= I_0(\epsilon) - \frac{2I'_0(\epsilon)\delta\tau}{3\epsilon(1 - 2\epsilon)} + O(\delta\tau^2) \end{aligned} \quad (33)$$

Thus

$$\frac{\partial s(\epsilon, \tau)}{\partial \tau} = \frac{\ln\left(\frac{\epsilon}{1-\epsilon}\right)}{3\epsilon(1 - 2\epsilon)} + O(\delta\tau) \quad (34)$$

just above the ER curve.

This ansatz gives an accurate description of our minimizing graphons near the ER curve, and these extend continuously all the way to the upper boundary $\tau = \epsilon^{3/2}$. Although the behavior in a first-order neighborhood of the ER curve changes discontinuously when ϵ passes through $1/2$, the behavior a finite distance above the ER curve appears to be analytic in ϵ and τ , so that the entire region between the ER curve and the upper boundary corresponds to a single phase. As can be seen from these formulas, or from the numerics of Section 2, phase I has the following qualitative features:

1. The minimizing graphon is bipodal. The smallest value of $g(x, y)$ is found on one of the two square regions $x, y < c$ or $x, y > c$, and the largest value is found on the other, with the value on the rectangles $x < c < y$ and $y < c < x$ being intermediate.
2. If $\epsilon \neq 1/2$, then as τ approaches ϵ^3 from above, c goes to 0. The value of $g(x, y)$ is then approximately ϵ on the large square $x, y > c$ and $1 - \epsilon$ on the rectangles. If $1/3 < \epsilon < 2/3$, then the value of $g(x, y)$ on the small square $x, y < c$ is approximately $2 - 3\epsilon$; the value is close to 1 if $\epsilon < 1/3$ and close to zero if $\epsilon > 2/3$.
3. Conversely, as τ moves away from ϵ^3 , c grows quickly, to first order in $\delta\tau$. When $\epsilon < 1/2$, the small- c region is quite narrow, as one sees in the second row of Fig. 8, and is easy to miss with numerical experiments. Although rapid, the transition from small- c to large- c appears to be continuous and smooth, with no sign of any phase transitions.
4. As τ approaches $\epsilon^{3/2}$ (from below, of course), the value of $g(x, y)$ approaches 1 on a square of size $\sqrt{\epsilon}$, approaches zero relatively slowly on the rectangles, and approaches zero more quickly on the other square. In a graph represented by such a graphon, there is a cluster representing a fraction $\sqrt{\epsilon}$ of the vertices. The probability of two vertices within the cluster being connected is close to 1, the probability of a vertex inside the cluster being connected to a vertex outside the cluster is small, and the probability of two vertices outside the cluster being connected is very small.

3.3 Phases II and III: $\tau < \epsilon^3$

When $\tau < \epsilon^3$, we want to minimize the μ^2 term in $\delta\tau$ and make the ν^3 term as negative as possible. This is done by simply taking $\mu = 0$ and $\nu = (\delta\tau)^{1/3}$. This gives a graphon of the form

$$g(x, y) = \begin{cases} \epsilon + \nu \frac{1-c}{c} & x, y < c \\ \epsilon - \nu & x < c < y \text{ or } y < c < x \\ \epsilon + \nu \frac{c}{1-c} & x, y > c \end{cases} \quad (35)$$

which is displayed in the left of Fig. 12.

Note that ν is negative, so the value of the graphon is less than ϵ on the squares and greater than ϵ on the rectangles. The two vertex clusters of fraction c and $1 - c$ prefer to connect with each other than with themselves.

When $c \neq 1/2$, this gives a graphon that takes on 3 different values. When $c = 1/2$, the graphon takes on only two values, but these values do not sum to 1. Only when $\epsilon = c = 1/2$ can we simultaneously solve both optimization problems.

To understand which problem “wins” the competition, we look at the situation when $\epsilon < 1/2$ and τ is slightly less than ϵ^3 . In this region, the entropy cost for having $\mu \neq 0$ is lower order in $\epsilon^3 - \tau$ than the entropy cost for violating the second problem. This shows that




Figure 12: The graphons for the expressions in (35) (left), (36) (middle) and (37) (right) respectively.

the optimal bipodal graphon must have c close to $1/2$, validating the use of perturbation theory around $c = 1/2$.

In subsection 3.4, we compute the second variation of the rate function with respect to c at $c = 1/2$, and see that it is positive for all values of $\tau < \epsilon^3$ when $\epsilon < 1/2$, and for some values of τ when $1/2 < \epsilon < 0.629497839$. In this regime, the best we can do is to pick $c = 1/2$ and μ exactly zero. Our graphon then takes the form

$$g(x, y) = \begin{cases} \epsilon - (\epsilon^3 - \tau)^{1/3} & x, y < 1/2 \text{ or } x, y > 1/2 \\ \epsilon + (\epsilon^3 - \tau)^{1/3} & x < \frac{1}{2} < y \text{ or } y < \frac{1}{2} < x \end{cases} \quad (36)$$

These symmetric bipodal graphons, displayed in the middle of Fig. 12, are the optimizers for phase II.

The symmetric bipodal phase II has a natural boundary. If $\epsilon > 1/2$, then the smallest possible value of τ is when $\nu = \epsilon - 1$, in which case $\tau = 2\epsilon^3 - 3\epsilon^2 + 3\epsilon - 1$. It is possible to get lower values of τ with $c \neq 1/2$, and still lower values if the graphon is not bipodal. Thus there must be a phase transition between phase II and the asymmetric bipodal phase III.

Phase III has its own natural boundary. Among bipodal graphons, the minimum triangle density for a given edge density (with $\epsilon > 1/2$) is given by a graphon of the form

$$g(x, y) = \begin{cases} 0 & x, y < c \\ 1 & x < c < y \text{ or } y < c < x \\ \alpha & x, y > c \end{cases} \quad (37)$$

displayed in the right of Fig. 12, where $c, \alpha \in [0, 1]$ are free parameters.

The edge and triangle densities are then

$$\epsilon = 2c(1 - c) + (1 - c)^2\alpha; \quad \tau = (1 - c)^3\alpha^3 + 3c(1 - c)^2\alpha. \quad (38)$$

Minimizing τ for fixed ϵ yields a sixth order polynomial equation in c , which we can solve numerically.

The main qualitative features of phases II and III are

1. The minimizing graphon consists of two squares, of side c and $1 - c$, and two $c \times (1 - c)$ rectangular regions. The value of the graphon is largest in the rectangular regions and smaller in the squares. In phase II we have $c = 1/2$, and the value of the graphon is the same in both squares. In phase III we have $c < 1/2$, and the value of the graphon is smallest in the $c \times c$ square.
2. All points with $\epsilon < 1/2$ and $\tau < \epsilon^3$ lie in phase II. As $\tau \rightarrow 0$, the graphon describes graphs that are close to being bipartite with two clusters of equal size such that edges *within* a cluster appear with probability close to zero, and edges *between* clusters appear with probability close to 2ϵ .
3. When $\epsilon > 1/2$, phase II does **not** extend to its natural boundary. There comes a point where one can reduce the rate function by making c less than $1/2$. For instance, when $\epsilon = 0.6$ one can construct symmetric bipodal graphons with any value of τ between $2\epsilon^3 - 3\epsilon^2 + 3\epsilon - 1 = 0.152$ and $\epsilon^3 = 0.216$. However, phase II only extends between $\tau = 0.152704753$ and $\tau = 0.20651775$, with the remaining intervals being phase III. The actual boundary of phase II is computed in the next subsection.
4. Phase III, by contrast, appears to extend all the way to its natural boundary, shown in Fig. 2, to the current numerical resolution in the τ variable. The union of phases I, II, and III appears to be all values of (ϵ, τ) that can be achieved with a bipodal graphon.
5. When we cross the natural boundary of phase III, the larger of the two clusters breaks into two equal-sized pieces, with slightly different probabilities of having edges within and between those sub-clusters. The resulting tripodal graphons have $g(x, y)$ strictly between 0 and 1, and (for $\epsilon < 2/3$) extend continuously down to the $k = 3$ scallop.
6. Throughout phases I, II, and III, there appears to be a unique graphon that minimizes the rate function. We conjecture that this is true for all (ϵ, τ) , and that the minimizing graphon for each (ϵ, τ) is k -podal with k as small as possible.

3.4 The boundary between phase II and phase III

In both phase II and phase III, we can express our graphons in the form (27). For each value of c we can vary μ and ν to minimize the rate function, while preserving the constraint on τ . Let $\mu(c)$ and $\nu(c)$ be the values of μ and ν that achieve this minimum for fixed c . Note that changing c to $1 - c$ and changing μ to $-\mu$ results in the same graphon, up to reparametrization of $[0, 1]$. By this symmetry, $\mu(c)$ must be an odd function of $c - 1/2$ while $\nu(c)$ must be an even function.

Since $\tau = \epsilon^3 + \nu^3 + 3(\epsilon + \nu)\mu^2/2$ is independent of c , we obtain constraints on derivatives of $\mu(c)$ and $\nu(c)$ by computing

$$0 = \tau'(c) = 3\nu^2\nu' + \frac{3\mu^2\nu'}{2} + 3(\epsilon + \nu)\mu\mu'$$

$$0 = \tau''(c) = 6\nu(\nu')^2 + 3\nu^2\nu'' + \frac{3\mu^2\nu''}{2} + 6\mu\mu'\nu' + 3(\epsilon + \nu)(\mu')^2 + 3(\epsilon + \nu)\mu\mu'', \quad (39)$$

where ' denotes d/dc . Since $\mu(1/2) = 0$, the first equation implies that $\nu'(1/2) = 0$, while the second implies that

$$\nu'' = -\frac{\epsilon + \nu}{\nu^2}(\mu')^2 \quad (40)$$

at $c = 1/2$.

We next compute the rate function. We expand the formula (27) as

$$g(x, y) = \begin{cases} \epsilon + \nu\frac{1-c}{c} + \mu\sqrt{2}\sqrt{\frac{1-c}{c}} & x, y < c \\ \epsilon - \nu + \frac{\mu}{\sqrt{2}} \left(\sqrt{\frac{1-c}{c}} - \sqrt{\frac{c}{1-c}} \right) & x < c < y \text{ or } y < c < x \\ \epsilon + \nu\frac{c}{1-c} - \mu\sqrt{2}\sqrt{\frac{c}{1-c}} & x, y > c. \end{cases} \quad (41)$$

This makes the rate function

$$\begin{aligned} I(g) &= c^2 I_0 \left(\epsilon + \nu\frac{1-c}{c} + \sqrt{2}\mu\sqrt{\frac{1-c}{c}} \right) \\ &\quad + 2c(1-c) I_0 \left(\epsilon - \nu + \frac{\mu}{\sqrt{2}} \left(\sqrt{\frac{1-c}{c}} - \sqrt{\frac{c}{1-c}} \right) \right) \\ &\quad + (1-c)^2 I_0 \left(\epsilon + \nu\frac{c}{1-c} - \mu\sqrt{2}\sqrt{\frac{c}{1-c}} \right) \end{aligned} \quad (42)$$

Taking derivatives at $c = 1/2$ and plugging in (40) yields

$$\left. \frac{d^2 I}{dc^2} \right|_{c=1/2} = A(\mu')^2 + B\mu' + C, \quad (43)$$

where

$$\begin{aligned} A &= I_0''(\epsilon + \nu) + \frac{(\epsilon + \nu)[I_0'(\epsilon - \nu) - I_0'(\epsilon + \nu)]}{2\nu^2} \\ B &= 2\sqrt{2} [I_0'(\epsilon + \nu) - I_0'(\epsilon - \nu) - 2\nu I_0''(\epsilon + \nu)] \\ C &= 4[I_0(\epsilon + \nu) - 2\nu I_0'(\epsilon + \nu) + 2\nu^2 I_0''(\epsilon + \nu) - I_0(\epsilon - \nu)], \end{aligned} \quad (44)$$

and $\nu = -(\epsilon^3 - \tau)^{1/3}$. The actual value of μ' is the one that minimizes this quadratic expression, namely $\mu' = -B/(2A)$. At this value of μ' , $d^2 I/dc^2$ equals $(4AC - B^2)/(4A)$. As long as the discriminant $B^2 - 4AC$ is negative, $d^2 I/dc^2 > 0$ and the phase II graphon is stable against changes in c . When $B^2 - 4AC$ goes positive, the phase II graphon becomes unstable and the minimizing graphon has c different from $1/2$.

Since the function I_0 is transcendental, it is presumably impossible to express the solution to $B^2 - 4AC = 0$ in closed form. However, it is easy to compute this boundary numerically, as shown in Fig. 2.

4 Exponential Random Graph Models

There is a large, rapidly growing literature on the analysis of large graphs, much of it devoted to the analysis of some *specific* large graph of particular interest; see [N] and the references therein. There is also a significant fraction of the literature which deals with the asymptotics of graphs g as the vertex number diverges (see [Lov]), and in these works a popular tool is ‘exponential random graph models’ (‘ERGMS’) in which a few densities $t_j(g)$, $j = 1, 2, \dots$, such as edge and triangle densities, are selected, conjugate parameters β_j are introduced, and the family of relative probability densities:

$$\rho_{\beta_1, \beta_2, \dots} = \exp[n^2(\beta_1 t_1(g) + \beta_2 t_2(g) + \dots)] \quad (45)$$

is implemented on the space $G(n)$ of all simple graphs g on n vertices. Densities are assumed normalized to have value 1 in the complete graph, so the analysis is relevant to ‘dense graphs’, which have order n^2 edges.

$\rho_{\beta_1, \beta_2, \dots}$ has the obvious form of a grandcanonical ensemble, with free energy density

$$\psi_{\beta_1, \beta_2, \dots} = \frac{1}{n^2} \ln \left[\sum_{g \in G(n)} \exp[n^2(\beta_1 t_1(g) + \beta_2 t_2(g) + \dots)] \right] \quad (46)$$

and suggests the introduction of other ensembles. This paper is the third in a sequence in which we have studied dense networks in the microcanonical ensemble, in this paper specializing to edge and triangle densities.

One goal in the study of the asymptotics of graphs is *emergent phenomena*. The prototypical emergent phenomena are the thermodynamic phases (and phase transitions) which ‘emerge’ as volume diverges in the statistical mechanics of particles interacting through short range forces. The results motivated by these – the mathematics of phase transitions in large dense graphs [CD, RY, AR, LZ, RS1, Y, RS2, YRF] – are a significant instance of emergence within mathematics. We conclude with some observations on the relationship between the phase transitions in dense networks as they appear in different ensembles.

In the statistical mechanics of particles with short range forces there is a well-developed theory of the equivalence of ensembles, and in particular any phase transition seen in one ensemble automatically appears in any other ensemble because corresponding free energies are Legendre transforms of one another and transitions are represented by singularities in free energies [Ru]. For particles with long range forces (mean-field models) the equivalence of ensembles can break down [TET]. For dense networks this is known to be the case [RS1]. So the precise relationship between the phase decomposition in the various ensembles of dense networks is an open problem of importance. For instance in ERGMS one of the key results is a transition within a single phase, structurally similar to a gas/liquid transition, whose optimal graphons are all Erdős-Rényi [CD, RY]. This phenomenon does not seem to have an obvious image in the microcanonical ensemble of this paper, but this should be explored.

Another important open problem is the phase structure associated with the remaining scallops.

Acknowledgment

The authors gratefully acknowledge useful discussions with Professor Rick Kenyon (Brown University). The computational codes involved in this research were developed and debugged on the computational cluster of the Mathematics Department of UT Austin. The main computational results were obtained on the computational facilities in the Texas Super Computing Center (TACC). We gratefully acknowledge these computational supports. This work was partially supported by NSF grants DMS-1208941, DMS-1321018 and DMS-1101326.

References

- [AR] D. Aristoff and C. Radin, Emergent structures in large networks, *J. Appl. Probab.* 50 (2013) 883-888.
- [BCL] C. Borgs, J. Chayes and L. Lovász, Moments of two-variable functions and the uniqueness of graph limits, *Geom. Funct. Anal.* 19 (2010) 1597-1619.
- [BCLSV] C. Borgs, J. Chayes, L. Lovász, V.T. Sós and K. Vesztergombi, Convergent graph sequences I: subgraph frequencies, metric properties, and testing, *Adv. Math.* 219 (2008) 1801-1851.
- [CD] S. Chatterjee and P. Diaconis, Estimating and understanding exponential random graph models, *Ann. Statist.* 41 (2013) 2428-2461.
- [CV] S. Chatterjee and S.R.S. Varadhan, The large deviation principle for the Erdős-Rényi random graph, *Eur. J. Comb.* 32 (2011) 1000-1017.
- [GMSW] P.E. Gill, W. Murray, M.A. Saunders and M.H. Wright, User's guide for NPSOL 5.0: A Fortran package for nonlinear programming, Technical Report SOL 86-6, System Optimization Laboratory, Stanford University, 2001.
- [Lov] L. Lovász, Large networks and graph limits, American Mathematical Society, Providence, 2012.
- [LS1] L. Lovász and B. Szegedy, Limits of dense graph sequences, *J. Combin. Theory Ser. B* 98 (2006) 933-957.
- [LS2] L. Lovász and B. Szegedy, Szemerédi's lemma for the analyst, *GAFA* 17 (2007) 252-270.
- [LS3] L. Lovász and B. Szegedy, Finitely forcible graphons, *J. Combin. Theory Ser. B* 101 (2011) 269-301.
- [LZ] E. Lubetzky and Y. Zhao, On replica symmetry of large deviations in random graphs, *Random Structures and Algorithms* (to appear), arXiv:1210.7013 (2012).

- [N] M.E.J. Newman, *Networks: an Introduction*, Oxford University Press, 2010.
- [NW] J. Nocedal and S.J. Wright, *Numerical Optimization*, Springer-Verlag, New York, 1999.
- [PR] O. Pikhurko and A. Razborov, Asymptotic structure of graphs with the minimum number of triangles, arXiv:1203.4393 (2012).
- [Ra] A. Razborov, On the minimal density of triangles in graphs, *Combin. Probab. Comput.* 17 (2008) 603-618.
- [RS1] C. Radin and L. Sadun, Phase transitions in a complex network, *J. Phys. A: Math. Theor.* 46 (2013) 305002.
- [RS2] C. Radin and L. Sadun, Singularities in the entropy of asymptotically large simple graphs, arXiv:1302.3531 (2013).
- [Ru] D. Ruelle *Statistical Mechanics; Rigorous Results*, Benjamin, New York, 1969.
- [RY] C. Radin and M. Yin, Phase transitions in exponential random graphs, *Ann. Appl. Probab.* 23 (2013) 2458-2471.
- [TET] H. Touchette, R.S. Ellis and B. Turkington, *Physica A* 340 (2004) 138-146.
- [Y] M. Yin, Critical phenomena in exponential random graphs, *J. Stat. Phys.* 153 (2013) 1008-1021.
- [YRF] M. Yin, A. Rinaldo and S. Fadnavis, Asymptotic quantization of exponential random graphs, arXiv:1311.1738 (2013).