Critical Exponents for Marked Random Connection Models

Matthew Dickson

Universität Augsburg

Aug 1, 2023

Joint work with Alejandro Caicedo and Markus Heydenreich

The Marked Random Connection Model

- We produce a random graph on the space $\mathbb{X} = \mathbb{R}^d \times \mathcal{E}$, where $\mathcal{E}$ is the **Mark Space**.
- The **Vetex Set** $\eta \subset \mathbb{X}$ is a Poisson Point Process with intensity measure $\lambda\nu$. $\lambda > 0$ and $\nu = \mathrm{Leb} \otimes \mathcal{P}$, where $\mathcal{P}$ is a probability measure on $\mathcal{E}$.
- Edges form independently according to a given symmetric **Adjacency function:**

$$\varphi(x, y) = \mathbb{P}(x \sim y)$$

where $x = (\overline{x}, a)$ and $y = (\overline{y}, b)$. In order to have spatial translation invariance, we require $\varphi(x, y) = \varphi(\overline{x} - \overline{y}; a, b)$.

Example Model: Boolean Hyper-Sphere Model

- First assign vertices to $\mathbb{R}^d$ according to a PPP with intensity $\lambda\mathrm{Leb}$, with $\lambda > 0$.

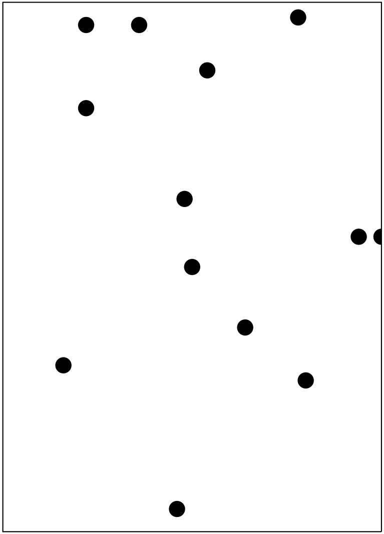

Example Model: Boolean Hyper-Sphere Model

- First assign vertices to $\mathbb{R}^d$ according to a PPP with intensity $\lambda\mathrm{Leb}$, with $\lambda > 0$.
- Independently assign each vertex a disc with random radius in $\mathcal{E} = \mathbb{R}_+$ with distribution $\mathcal{P}$.

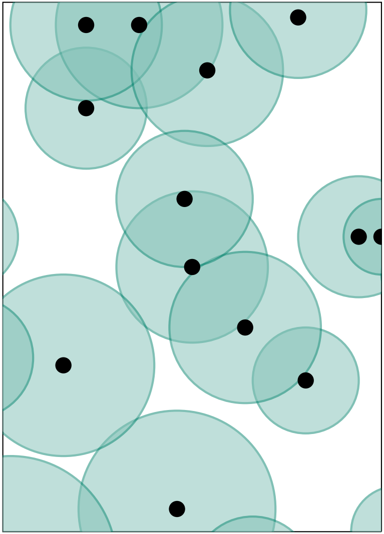

Example Model: Boolean Hyper-Sphere Model

- First assign vertices to $\mathbb{R}^d$ according to a PPP with intensity $\lambda \mathrm{Leb}$, with $\lambda > 0$.
- Independently assign each vertex a disc with random radius in $\mathcal{E} = \mathbb{R}_+$ with distribution $\mathcal{P}$.
- Edges form between vertices when their discs overlap. This corresponds to

$$\varphi\left(\overline{x}; a, b\right) = \mathbb{1}\left\{\left|\overline{x}\right| < a + b\right\}.$$

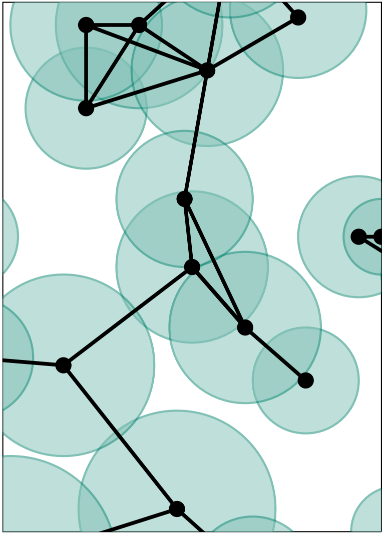

Other Examples

- **Factorisable Models:** Let $\psi\colon \mathbb{R}^d \to [0,1]$ and $K\colon \mathcal{E}\times\mathcal{E} \to [0,1]$ be symmetric and measurable. Then let

$$\varphi\left(\overline{x}; a, b\right) = \psi\left(\overline{x}\right) K\left(a, b\right).$$

- **Gaussian Model:** Let $\Sigma\colon \mathcal{E}^2 \to \mathbb{R}^{d\times d}$ be a measurable map where for every $a, b \in \mathcal{E}$, $\Sigma\left(a, b\right)$ is itself a symmetric positive definite covariance matrix. Then let

$$\varphi(\overline{x}; a, b) = (2\pi)^{-d/2}\left(\det \Sigma(a,b)\right)^{-1/2} \exp\left(-\frac{1}{2}\overline{x}^{\intercal}\Sigma(a,b)^{-1}\overline{x}\right).$$

- **Weight-Dependent Models:** Let $\rho\colon \mathbb{R}_+ \to [0,1]$ be non-increasing and $g\colon\ (0,1)\times(0,1) \to \mathbb{R}_+$ be non-decreasing in both arguments. Then let

$$\varphi(\overline{x}; a, b) = \rho\left(g\left(a, b\right)\left|\overline{x}\right|^d\right).$$

Other Examples

- **Weight-Dependent Models:** Let $\rho\colon \mathbb{R}_+ \to [0,1]$ be non-increasing and $g\colon (0,1)\times(0,1) \to \mathbb{R}_+$ be non-decreasing in both arguments. Then let
$$\varphi(\overline{x}; a, b) = \rho\left(g\left(a,b\right)\left|\overline{x}\right|^d\right).$$

Different choices of ρ and g produce various models in the literature:

- Boolean model, Gilbert disk model [Gilbert '61, Hall '85]
- (Soft) random geometric graph [Penrose '93]
- Ultra-small scale-free geometric networks [Yukich '03]
- Scale-free Gilbert graph [Hirsch '17]
- Continuum scale-free percolation [Deprez-Wüthrich '19]
- Geometric inhomogeneous random graphs [Bringmann-Keusch-Lengler '19]
- Age-dependent random connection model [Gracar-Mönch-Mörters '19].

Critical Intensities

- The *cluster* of x is given by

$$\mathscr{C}(x) = \{y \in \eta^x : x \longleftrightarrow y \text{ in } \xi^x\}.$$

- Define *susceptibility*:

$$\chi_\lambda \colon \mathcal{E} \to [0,\infty], \qquad \chi_\lambda(a) = \mathbb{E}_\lambda\left[\left|\mathscr{C}\left(\overline{0},a\right)\right|\right],$$

and the *percolation probability*:

$$\theta_\lambda \colon \mathcal{E} \to [0,1], \qquad \theta_\lambda(a) = \mathbb{P}_\lambda\left(\left|\mathscr{C}\left(\overline{0},a\right)\right| = \infty\right).$$

- For $p \in [1,\infty]$ there are the associated critical intensities

$$\lambda_T^{(p)} := \inf\left\{\lambda > 0 \colon \left\|\chi_\lambda\right\|_p = \infty\right\}, \quad \lambda_c^{(p)} := \inf\left\{\lambda > 0 \colon \left\|\theta_\lambda\right\|_p > 0\right\}.$$

Critical Intensities

- Recall

$$\lambda_T^{(p)} := \inf\left\{\lambda > 0\colon \|\chi_\lambda\|_p = \infty\right\}, \quad \lambda_c^{(p)} := \inf\left\{\lambda > 0\colon \|\theta_\lambda\|_p > 0\right\}.$$

- Observe
  - For all $p \in [1,\infty]$, $\lambda_c^{(p)} = \lambda_c$.
  - For all $p \in [1,\infty]$, $\lambda_T^{(p)} \leq \lambda_c$.
  - For all $1 \leq p_1 \leq p_2 \leq \infty$, $\lambda_T^{(p_1)} \geq \lambda_T^{(p_2)}$.

Lemma (D., Heydenreich '22+ & Caicedo, D. '23+)

$$\operatorname*{ess\,sup}_{a,b\in\mathcal{E}} \int_{\mathbb{R}^d} \varphi(\overline{x}; a, b)\,\mathrm{d}\overline{x} < \infty \implies \lambda_T^{(\infty)} = \lambda_T^{(1)} = \lambda_T.$$

If furthermore $\varphi > 0$ on a $\mathrm{Leb}\times\mathcal{P}^2$-positive set and $d \geq 2$, then

$$0 < \lambda_T \leq \lambda_c < \infty.$$

For comparison, in the single-mark case we have $\lambda_T = \lambda_c$ [Meester '95], and $\lambda_c \in (0,\infty)$ iff $\int_{\mathbb{R}^d} \varphi(\overline{x})\,\mathrm{d}\overline{x} \in (0,\infty)$ (for $d \geq 2$) [Penrose '91].

Critical Exponents

Recall

$$\chi_\lambda(a) = \mathbb{E}_\lambda\left[\left|\mathscr{C}\left(\overline{0}, a\right)\right|\right], \qquad \theta_\lambda(a) = \mathbb{P}_\lambda\left(\left|\mathscr{C}\left(\overline{0}, a\right)\right| = \infty\right).$$

- How do χ_λ and θ_λ behave near λ_T and λ_c? Do there exist γ, β such that (in a bounded ratio sense)

$$\left\|\chi_\lambda\right\|_p \asymp \frac{1}{\left(\lambda_T - \lambda\right)^\gamma}, \qquad \text{as } \lambda \nearrow \lambda_T,$$

$$\left\|\theta_\lambda\right\|_p \asymp \left(\lambda - \lambda_c\right)^\beta, \qquad \text{as } \lambda \searrow \lambda_c?$$

- If you were to consider a spatial branching process with offspring kernel $\lambda\varphi$, the analogous quantities would have exponents $\gamma = 1$ and $\beta = 1$. These are called the *mean-field* exponents.

Some Assumptions ...

Given $a, b \in \mathcal{E}$ and $n \geq 1$, let us define

$$D(a,b) := \int_{\mathbb{R}^d} \varphi\left(\overline{x}; a, b\right) \mathrm{d}\overline{x},$$

$$D^{(n)}(a,b) := \int_{\mathcal{E}^{n-1}} \left(\prod_{j=1}^{n} D(c_{j-1}, c_j)\right) \mathcal{P}^{\otimes(n-1)}\left(\mathrm{d}\vec{c}_{[1,\dots,n-1]}\right),$$

where $c_0 = a$ and $c_k = b$. $D^{(n)}$ is the "matrix product" of n copies of D.

Assumptions (D)

(D.1) *"Every mark has bounded expected degree with every other mark"*

$$\operatorname*{ess\,sup}_{a,b\in\mathcal{E}} D(a,b) < \infty,$$

(D.2) *"Some mark can be connected to every other mark in exactly k steps for some k"*

$$\operatorname*{ess\,sup}_{a\in\mathcal{E}} \sup_{k\geq 1} \operatorname*{ess\,inf}_{b\in\mathcal{E}} D^{(k)}(a,b) > 0.$$

Mean Field Bounds

Theorem (Caicedo, D. '23+)

If Assumptions **(D.1)** *and* **(D.2)** *hold, then there exist* $\varepsilon > 0$ *and* $C > 0$ *such that*

$$\|\chi_\lambda\|_p \geq C\left(\lambda_T - \lambda\right)^{-1} \qquad \text{for } \lambda < \lambda_T,$$
$$\|\theta_\lambda\|_p \geq C\left(\lambda - \lambda_T\right)_+ \qquad \text{for } \lambda < \lambda_T + \varepsilon,$$

for all $p \in [1, \infty]$.

Mean Field Bounds

Theorem (Caicedo, D. '23+)

If Assumptions **(D.1)** *and* **(D.2)** *hold, then there exist* $\varepsilon > 0$ *and* $C > 0$ *such that*

$$\left\|\chi_\lambda\right\|_p \geq C\left(\lambda_T - \lambda\right)^{-1} \qquad \text{for } \lambda < \lambda_T,$$
$$\left\|\theta_\lambda\right\|_p \geq C\left(\lambda - \lambda_T\right)_+ \qquad \text{for } \lambda < \lambda_T + \varepsilon,$$

for all $p \in [1, \infty]$.

Oh look!

Mean Field Bounds

Theorem (Caicedo, D. '23+)

If Assumptions **(D.1)** *and* **(D.2)** *hold, then there exist* $\varepsilon > 0$ *and* $C > 0$ *such that*

$$\left\lVert \chi_\lambda \right\rVert_p \geq C \left(\lambda_T - \lambda\right)^{-1} \qquad \text{for } \lambda < \lambda_T,$$
$$\left\lVert \theta_\lambda \right\rVert_p \geq C \left(\lambda - \lambda_T\right)_+ \qquad \text{for } \lambda < \lambda_T + \varepsilon,$$

for all $p \in [1, \infty]$.

Oh look!

$$\lambda_T = \lambda_c$$

This theorem had already been proven for the Boolean Hyper-Sphere model [Dembin, Tassion '22] if the radius distribution has finite d-moments. This is a weaker radius condition than **(D.1)**, but our result can be applied to a wider class of models.

Mean-Field Behaviour

Theorem (Caicedo, D. '23+)

If Assumptions **(D.1)**, **(D.2)**, *and* **(T)** *hold, then there exist* $\varepsilon > 0$ *and* $C' > 0$ *such that*

$$\|\chi_\lambda\|_p \leq C' (\lambda_T - \lambda)^{-1} \qquad \text{for } \lambda < \lambda_T,$$
$$\|\theta_\lambda\|_p \leq C' (\lambda - \lambda_T)_+ \qquad \text{for } \lambda < \lambda_T + \varepsilon,$$

for all $p \in [1, \infty]$.
That is, $\gamma = 1$ *and* $\beta = 1$ *(they take their mean-field values).*

What is **(T)**?

What is **(T)**?

Define the *two-point function*:

$$\tau_\lambda(x,y) := \mathbb{P}_\lambda\left(x \longleftrightarrow y \text{ in } \xi^{x,y}\right).$$

For $\lambda \geq 0$, the "triangle diagram" is defined as

$$\triangle_\lambda := \lambda^2 \operatorname*{ess\,sup}_{x,y\in\mathbb{X}} \int \tau_\lambda(x,u)\tau_\lambda(u,v)\tau_\lambda(v,y)\nu^{\otimes 2}\left(\mathrm{d}u, \mathrm{d}v\right).$$

Triangle Condition

(T) We have

$$\triangle_{\lambda_T} < C_\triangle,$$

where $C_\triangle > 0$ is a specific constant.

When does **(T)** hold?

For the single-mark RCM, (Heydenreich, van der Hofstad, Last, Matzke '19) gave conditions under which there exists d^* such that the triangle condition holds for $d > d^*$. (It is expected that $d^* = 6$. This is not proven.)

When does **(T)** hold?

Let $\{\varphi_d\}_{d\geq 1}$ be a sequence of adjacency functions, each on $\left(\mathbb{R}^d \times \mathcal{E}\right)^2$.

Theorem (D., Heydenreich '22+)

Given conditions on $\{\varphi_d\}_{d\geq 1}$ (to be seen shortly), there exists a critical dimension $d^* \in \mathbb{N}$, a constant $C > 0$, and $\alpha = \alpha(d)$ such that for $d > d^*$ and all $\lambda \in [0, \lambda_T]$,

$$\triangle_\lambda \leq C\alpha.$$

The 'lace expansion' proof relies on deriving a linear operator equation (an Ornstein-Zernike Equation). We therefore need to introduce our operator formalism (relevant for the conditions).

Operator Notation

- For all $k \in \mathbb{R}^d$, let $\widehat{\varphi}(k; a, b)$ is the Fourier transform of $\varphi(\overline{x}; a, b)$ and define

$$\widehat{\Phi}(k) : L^2(\mathcal{E}) \to L^2(\mathcal{E}), \quad \left(\widehat{\Phi}(k) f\right)(a) = \int \widehat{\varphi}(k; a, b) f(b) \mathcal{P}(\mathrm{d}b).$$

- The important properties of these operators are encoded in their spectra. We have the spectral radius:

$$\rho\left(\widehat{\Phi}(k)\right) = \sup\left\{|z| : z \in \sigma\left(\widehat{\Phi}(k)\right)\right\}.$$

Since $\widehat{\Phi}(k)$ is self-adjoint, we can also define the *spectral supremum*:

$$\mathbb{S}\left(\widehat{\Phi}(k)\right) = \sup\left\{z : z \in \sigma\left(\widehat{\Phi}(k)\right) \subset \mathbb{R}\right\}.$$

- Sometimes easier to work with

$$\left\|\widehat{\Phi}(k)\right\|_{\infty,\infty} = \operatorname*{ess\,sup}_{a,b} |\widehat{\varphi}(k; a, b)|.$$

Marked RCM Assumptions

Assumption 1

For all dimensions d, there exists a d-independent constant $C > 0$ such that:

$$\mathbb{S}\left(\widehat{\Phi}(0)\right) < \infty, \qquad \left\|\widehat{\Phi}(0)\right\|_{\infty,\infty} \leq C\mathbb{S}\left(\widehat{\Phi}(0)\right),$$
$$\left\|\widehat{\Phi}(0) - \widehat{\Phi}(k)\right\|_{\infty,\infty} \leq C\left(\mathbb{S}\left(\widehat{\Phi}(0)\right) - \mathbb{S}\left(\widehat{\Phi}(k)\right)\right).$$

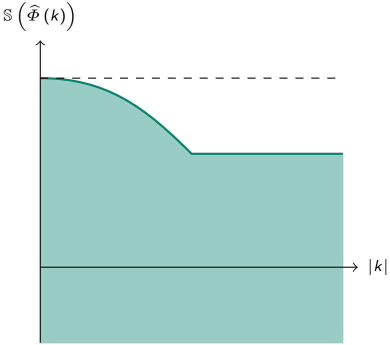

Assumption 2

There exist d-independent constants $C_1 \in (0,1)$ and $C_2 > 0$ such that

$$\mathbb{S}\left(\widehat{\Phi}(k)\right) \leq \left[C_1 \vee \left(1 - C_2|k|^2\right)\right]\mathbb{S}\left(\widehat{\Phi}(0)\right).$$

Marked RCM Assumptions

By a spatial scaling argument, w.l.o.g. $\mathbb{S}\left(\widehat{\Phi}(0)\right) = 1$.

Assumption 3

There exists a function $g\colon \mathbb{N} \to \mathbb{R}_{\geq 0}$ with the following three properties:

1. that $g(d) \to 0$ as $d \to \infty$,
2. that
$$\operatorname*{ess\,sup}_{\overline{x}\in\mathbb{R}^d, a_1,\ldots,a_6\in\mathcal{E}} \left(\varphi\left(\cdot; a_1, a_2\right) \star \varphi\left(\cdot; a_3, a_4\right) \star \varphi\left(\cdot; a_5, a_6\right)\right)\left(\overline{x}\right) \leq g(d),$$
3. and that the family of sets $\left\{B\left(x\right)\right\}_{x\in\mathbb{X}}$ given by
$$B\left(x\right) := \left\{y \in \mathbb{R}^d \times \mathcal{E} : \int \varphi\left(x, u\right)\varphi\left(u, y\right)\nu\left(\mathrm{d}u\right) > g\left(d\right)\right\}$$
satisfy $\operatorname{ess\,sup}_{x\in\mathbb{R}^d\times\mathcal{E}} \nu\left(B\left(x\right)\right) \leq g\left(d\right)$.

Proof Outline

- Recall the *two-point function*:

$$\tau_\lambda(x,y) := \mathbb{P}_\lambda\left(x \longleftrightarrow y \text{ in } \xi^{x,y}\right),$$

and define the associated $\widehat{\tau}_\lambda(k;a,b)$ and $\widehat{\mathcal{T}}_\lambda(k)$.

- **Lace Expansion:** For $\lambda \in [0,\lambda_T)$,

$$\widehat{\mathcal{T}}_\lambda(k) = \widehat{\Phi}(k) + \widehat{\Pi}_{\lambda,n}(k) + \lambda\widehat{\mathcal{T}}_\lambda(k)\left(\widehat{\Phi}(k) + \widehat{\Pi}_{\lambda,n}(k)\right) + \widehat{R}_{\lambda,n}(k).$$

These $\widehat{\Pi}_{\lambda,n}(k)$ are constructed by counting configurations with $\leq n$ pivotal points.

- **Bounding $\widehat{\Pi}_{\lambda,n}(k)$ with triangles:** We can count these configurations using thinnings, Mecke's formula, and BK inequality to get bounds in terms of integrals. By supremum bounds, we can extract triangles and other similar shapes.

Proof Outline

- **Compare to Random walk:** Define

$$\widehat{G}_{\mu_\lambda}(k) := \frac{1}{1-\mu_\lambda \mathbb{S}\left(\widehat{\Phi}(k)\right)}, \qquad f(\lambda) := \operatorname*{ess\,sup}_{k\in\mathbb{R}^d} \frac{\rho\left(\widehat{\mathcal{T}}_\lambda(k)\right)}{\widehat{G}_{\mu_\lambda}(k)}.$$

  In the single-mark model, $\widehat{G}_{\mu_\lambda}(k)$ is the Fourier transform of the Green's function of a random walk with jump density φ. We can use $f(\lambda)$ to replace factors of the 'unknown' τ_λ in the triangles with the 'known' φ.

- **Sub-critical Convergence:** For d sufficiently large, we can show that the $\widehat{G}_{\mu_\lambda}$-triangle (and the original τ_λ-triangle) are small, and therefore the expansion converges for $\lambda < \lambda_T$.

$$\widehat{\mathcal{T}}_\lambda(k) = \widehat{\Phi}(k) + \widehat{\Pi}_\lambda(k) + \lambda\widehat{\mathcal{T}}_\lambda(k)\left(\widehat{\Phi}(k) + \widehat{\Pi}_\lambda(k)\right). \qquad \text{(OZE)}$$

- **Uniform Sub-critical Convergence:** Using this expansion, we show that $f(\lambda)$ is uniformly bounded on the entire sub-critical regime (by a forbidden range argument). This also gives a uniform bound $\triangle_\lambda \leq C\alpha$ for $\lambda < \lambda_T$.
- **Extend (OZE) to Criticality:** By Monotone and Dominated convergence arguments, (OZE) holds at $\lambda = \lambda_T$ and $\triangle_{\lambda_T} \leq C\alpha$.

References

- Alejandro Caicedo and Matthew Dickson. *Critical exponents for marked random connection models*. Preprint arXiv:2305.07398 [math.PR].
- Matthew Dickson and Markus Heydenreich. *The triangle condition for the marked random connection model*. Preprint arXiv:2210.07727 [math.PR].
- Markus Heydenreich, Remco van der Hofstad, Günter Last, and Kilian Matzke. *Lace expansion and mean-field behavior for the random connection model*, Preprint arXiv:1908.11356 [math.PR].