Disordered systems and random graphs 2

Amin Coja-Oghlan
Goethe University

based on joint work with Dimitris Achlioptas, Oliver Gebhard, Max Hahn-Klimroth, Joon Lee, Philipp Loick, Noela Müller, Manuel Penschuck, Guangyan Zhou

Overview

This lecture: random 2-SAT

- Belief Propagation and density evolution
- the contraction method
- spatial mixing
- the Aizenman-Sims-Starr scheme
- the interpolation method

Random 2-SAT

The 2-SAT problem

- Boolean variables $x_1, \ldots, x_n$
- truth values $+1$ and -1
- four types of clauses:

$$x_i \vee x_j \qquad x_i \vee \neg x_j \qquad \neg x_i \vee x_j \qquad \neg x_i \vee \neg x_j$$

- a 2-SAT formula is a conjunction $\Phi = \bigwedge_{i=1}^{m} a_i$ of clauses
- $S(\Phi) =$ set of satisfying assignments
- $Z(\Phi) = |S(\Phi)|$

Random 2-SAT

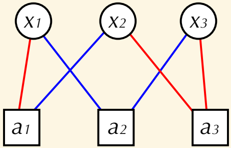

Random 2-SAT

- for a fixed $0 < d < \infty$ let $\boldsymbol{m} = \text{Po}(dn/2)$
- $\boldsymbol{\Phi}$ = conjunction of $\boldsymbol{m}$ independent random clauses
- variable degrees have distribution $\text{Po}(d)$
- *Key questions:* is $Z(\boldsymbol{\Phi}) > 0$ and if so, what is

$$\lim_{n\to\infty} \frac{1}{n} \log Z(\boldsymbol{\Phi}) \quad ?$$

The cavity method

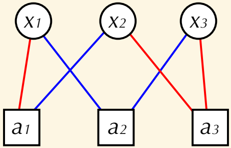

The factor graph

- vertices $x_1, \ldots, x_n$ represent variables
- vertices $a_1, \ldots, a_m$ represent clauses
- the graph $G(\mathbf{\Phi})$ contains few short cycles
- locally $G(\mathbf{\Phi})$ resembles a Galton-Watson branching process

The cavity method

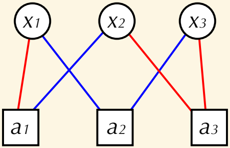

The Boltzmann distribution

- assuming $S(\mathbf{\Phi}) \neq \emptyset$ define

$$\mu_{\mathbf{\Phi}}(\sigma) = \frac{\mathbf{1}\{\sigma \in S(\mathbf{\Phi})\}}{Z(\mathbf{\Phi})} \qquad (\sigma \in \{\pm 1\}^{\{x_1,\ldots,x_n\}})$$

- let $\boldsymbol{\sigma} = \boldsymbol{\sigma}_{\mathbf{\Phi}}$ be a sample from $\mu_{\mathbf{\Phi}}$

The cavity method

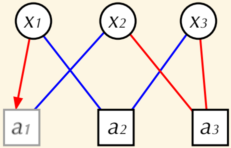

Belief Propagation

- define the variable–to–clause messages by

$$\mu_{\mathbf{\Phi},x\to a}(\sigma) = \mu_{\mathbf{\Phi}-a}(\boldsymbol{\sigma}_x = \sigma) \qquad (\sigma = \pm 1)$$

- “marginal of x upon removal of a”

The cavity method

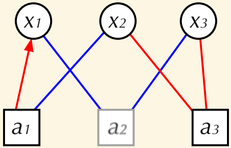

Belief Propagation

- define the clause–to–variable messages by

$$\mu_{\mathbf{\Phi},a\to x}(\sigma) = \mu_{\mathbf{\Phi}-(\partial x\setminus a)}(\boldsymbol{\sigma}_x = \sigma) \qquad (\sigma = \pm 1)$$

- “marginal of x upon removal of all neighbours $b \in \partial x$, $b \neq a$”

The cavity method

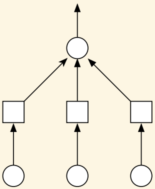

The replica symmetric ansatz

The messages (approximately) satisfy

$$\mu_{\mathbf{\Phi},x\to a}(\sigma) \propto \prod_{b\in\partial x\setminus a} \mu_{\mathbf{\Phi},b\to x}(\sigma)$$

$$\mu_{\mathbf{\Phi},a\to x}(\sigma) \propto 1 - \mathbf{1}\left\{\sigma \neq \mathrm{sign}(x,a)\right\} \mu_{\mathbf{\Phi},\partial a\setminus x}(-\mathrm{sign}(\partial a \setminus x))$$

The cavity method

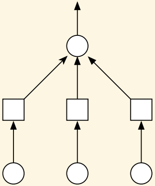

The Bethe free entropy

- we expect that

$$\log Z(\mathbf{\Phi}) \sim \sum_{i=1}^{n} \log \sum_{\sigma=\pm 1} \prod_{a \in \partial x_i} \mu_{\mathbf{\Phi}, a \to x}(\sigma)$$
$$+ \sum_{i=1}^{\boldsymbol{m}} \log \left(1 - \prod_{x \in \partial a_i} \boldsymbol{\mu}_{\mathbf{\Phi}, x \to a_i}(-\mathrm{sign}(x, a_i)) \right)$$
$$- \sum_{i=1}^{n} \sum_{a \in \partial x_i} \log \sum_{\sigma=\pm 1} \mu_{\mathbf{\Phi}, x \to a_i}(\sigma) \mu_{\mathbf{\Phi}, a_i \to x}(\sigma)$$

The cavity method

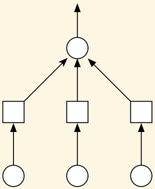

Density evolution

- consider the empirical distribution of the messages:

$$\pi_{\Phi} = \frac{1}{2\boldsymbol{m}} \sum_{i=1}^{n} \sum_{a \in \partial x_i} \delta_{\mu_{\Phi, x \to a}(+1)}$$

- $\boldsymbol{d}^+, \boldsymbol{d}^- \sim \mathrm{Po}(d/2)$, $\boldsymbol{\mu}_0, \boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \ldots$ samples from π_{Φ}

$$\boldsymbol{\mu}_0 \stackrel{\mathrm{d}}{=} \frac{\prod_{i=1}^{\boldsymbol{d}_+} \boldsymbol{\mu}_i}{\prod_{i=1}^{\boldsymbol{d}_+} \boldsymbol{\mu}_i + \prod_{i=1}^{\boldsymbol{d}_-} \boldsymbol{\mu}_{i+\boldsymbol{d}^+}}$$

The cavity method

Summary: the replica symmetric prediction [MZ96]

For $d < 2$ there is a unique distribution π_d on $(0,1)$ s.t.

$$\boldsymbol{\mu}_0 \overset{\mathrm{d}}{=} \frac{\prod_{i=1}^{\boldsymbol{d}_+} \boldsymbol{\mu}_i}{\prod_{i=1}^{\boldsymbol{d}_+} \boldsymbol{\mu}_i + \prod_{i=1}^{\boldsymbol{d}_-} \boldsymbol{\mu}_{i+\boldsymbol{d}^+}}$$

and $\lim_{n\to\infty} n^{-1} \log Z(\mathbf{\Phi}) = \mathscr{B}_d$ where

$$\mathscr{B}_d = \mathrm{E}\left[\log\left(\prod_{i=1}^{\boldsymbol{d}_+} \boldsymbol{\mu}_i + \prod_{i=1}^{\boldsymbol{d}_-} \boldsymbol{\mu}_{i+\boldsymbol{d}_+}\right) - \frac{d}{2}\log\left(1-\boldsymbol{\mu}_1\boldsymbol{\mu}_2\right)\right]$$

The cavity method

Theorem [ACOHKLMPZ20]

For $d < 2$ there is a unique distribution π_d on $(0,1)$ s.t.

$$\boldsymbol{\mu}_0 \overset{\mathrm{d}}{=} \frac{\prod_{i=1}^{\boldsymbol{d}_+} \boldsymbol{\mu}_i}{\prod_{i=1}^{\boldsymbol{d}_+} \boldsymbol{\mu}_i + \prod_{i=1}^{\boldsymbol{d}_-} \boldsymbol{\mu}_{i+\boldsymbol{d}^+}}$$

and $\lim_{n\to\infty} n^{-1} \log Z(\boldsymbol{\Phi}) = \mathscr{B}_d$ where

$$\mathscr{B}_d = \mathrm{E}\left[\log\left(\prod_{i=1}^{\boldsymbol{d}_+} \boldsymbol{\mu}_i + \prod_{i=1}^{\boldsymbol{d}_-} \boldsymbol{\mu}_{i+\boldsymbol{d}_+}\right) - \frac{d}{2}\log\left(1-\boldsymbol{\mu}_1\boldsymbol{\mu}_2\right)\right]$$

The cavity method

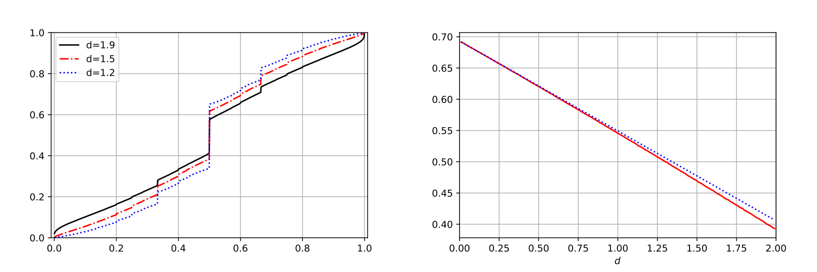

The proof strategy

Outline

1. *Contraction method:* unique solution to density evolution
2. *Spatial mixing:* the empirical distribution $\pi_{\mathbf{\Phi}}$
3. *Aizenman-Sims-Starr:* derivation of the Bethe formula
4. *Interpolation method:* concentration of $\log Z(\mathbf{\Phi})$

The proof strategy

Outline

1. *Contraction method:* unique solution to density evolution
2. *Spatial mixing:* the empirical distribution $\pi_{\mathbf{\Phi}}$
3. *Aizenman-Sims-Starr:* derivation of the Bethe formula
4. *Interpolation method:* concentration of $\log Z(\mathbf{\Phi})$

Comparison with prior work [DM10,DMS13,MS07,P14,T01]

- *zero temperature:* hard constraints
- *spatial mixing:* delicate construction of extremal boundaries
- *Aizenman-Sims-Starr* instead of varying temperature β

Step 1: the contraction method

Proposition

For $d < 2$ there is a unique distribution π_d on $(0,1)$ s.t.

$$\boldsymbol{\mu}_0 \overset{\mathrm{d}}{=} \frac{\prod_{i=1}^{\boldsymbol{d}_+} \boldsymbol{\mu}_i}{\prod_{i=1}^{\boldsymbol{d}_+} \boldsymbol{\mu}_i + \prod_{i=1}^{\boldsymbol{d}_-} \boldsymbol{\mu}_{i+\boldsymbol{d}^+}}$$

Step 1: the contraction method

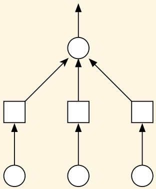

Log-likelihood ratios

- let $\boldsymbol{d} = \mathrm{Po}(d)$
- let $\boldsymbol{s}_1, \boldsymbol{s}'_1, \boldsymbol{s}_2, \boldsymbol{s}'_2, \ldots \in \{\pm 1\}$ be unifom and independent
- introducing

$$\boldsymbol{\eta}_i = \log \frac{\boldsymbol{\mu}_i}{1 - \boldsymbol{\mu}_i} \in \mathbb{R}$$

we obtain

$$\boldsymbol{\eta}_0 \stackrel{\mathrm{d}}{=} \sum_{i=1}^{\boldsymbol{d}} \boldsymbol{s}_i \log \frac{1 + \boldsymbol{s}'_i \tanh(\boldsymbol{\eta}_i / 2)}{2}$$

Step 1: the contraction method

The Wasserstein space

- $\mathcal{W}_2(\mathbb{R}) = \{\text{probability measures with finite 2nd moment}\}$
- for $\varrho, \varrho' \in \mathcal{W}_2(\mathbb{R})$ define

$$\Delta_2(\varrho, \varrho') = \inf_{\boldsymbol{X} \sim \varrho, \boldsymbol{X}' \sim \varrho'} \sqrt{\mathrm{E}[(\boldsymbol{X} - \boldsymbol{X}')^2]}$$

- this metric turns $\mathcal{W}_2(\mathbb{R})$ into a complete separable space

Step 1: the contraction method

The Banach fixed point theorem

- a map $F : \mathscr{W}_2(\mathbb{R}) \to \mathscr{W}_2(\mathbb{R})$ is a *contraction* if

$$\Delta_2(F(\varrho), F(\varrho')) \leq (1-\epsilon)\Delta_2(\varrho, \varrho') \qquad (\varrho, \varrho' \in \mathscr{W}_2(\mathbb{R}))$$

- a contraction has a unique fixed point

Step 1: the contraction method

Lemma

The map $F : \mathscr{W}_2(\mathbb{R}) \to \mathscr{W}_2(\mathbb{R})$ that maps ϱ to the distribution of

$$\sum_{i=1}^{\boldsymbol{d}} \boldsymbol{s}_i \log \frac{1 + \boldsymbol{s}_i' \tanh(\boldsymbol{\eta}_i/2)}{2}$$

is a contraction.

Step 1: the contraction method

Proof

With $(\boldsymbol{\eta}_1, \boldsymbol{\eta}_1'), (\boldsymbol{\eta}_2, \boldsymbol{\eta}_2'), \ldots$ be pairs drawn from ϱ, ϱ',

$$\begin{aligned}
\Delta_2(F(\varrho), F(\varrho'))^2 &\le \mathrm{E}\left[\left(\sum_{i=1}^{\boldsymbol{d}} \boldsymbol{s}_i \log \frac{1+\boldsymbol{s}_i' \tanh(\boldsymbol{\eta}_i/2)}{1+\boldsymbol{s}_i' \tanh(\boldsymbol{\eta}_i'/2)}\right)^2\right] \\
&= \mathrm{E}\left[\sum_{i=1}^{\boldsymbol{d}} \log^2 \frac{1+\boldsymbol{s}_i' \tanh(\boldsymbol{\eta}_i/2)}{1+\boldsymbol{s}_i' \tanh(\boldsymbol{\eta}_i'/2)}\right] \\
&= d \cdot \mathrm{E}\left[\log^2 \frac{1+\boldsymbol{s}_1' \tanh(\boldsymbol{\eta}_1/2)}{1+\boldsymbol{s}_1' \tanh(\boldsymbol{\eta}_1'/2)}\right] \\
&= \frac{d}{2} \sum_{s=\pm 1} \mathrm{E}\left[|\boldsymbol{\eta}_1 - \boldsymbol{\eta}_1'| \int_{\boldsymbol{\eta}_1 \wedge \boldsymbol{\eta}_1'}^{\boldsymbol{\eta}_1 \vee \boldsymbol{\eta}_1'} \left(\frac{1+s\tanh(z/2)}{2}\right)^2 \mathrm{d}z\right] \\
&\le \frac{d}{2}\mathrm{E}[(\boldsymbol{\eta}_1 - \boldsymbol{\eta}_1')^2] = \frac{d}{2}\Delta_2(\varrho, \varrho')^2
\end{aligned}$$

Step 2: spatial mixing

Proposition

For $d < 2$ the empirical distribution of marginals

$$\pi_{\mathbf{\Phi}} = \frac{1}{n} \sum_{i=1}^{n} \delta_{\boldsymbol{\mu}_{\mathbf{\Phi}}(\boldsymbol{\sigma}_{x_i}=1)}$$

converges to the density evolution fixed point π_d.

Step 2: spatial mixing

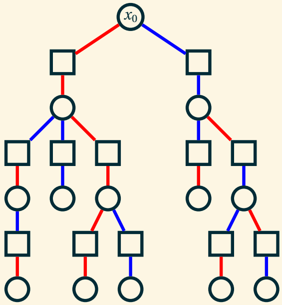

The Galton-Watson tree

- a random tree $\boldsymbol{T}$ comprising variable and clause nodes
- the root x_0 is a variable
- each variable node spawns $\mathrm{Po}(d)$ clause nodes
- each clause node has one variable node child

Step 2: spatial mixing

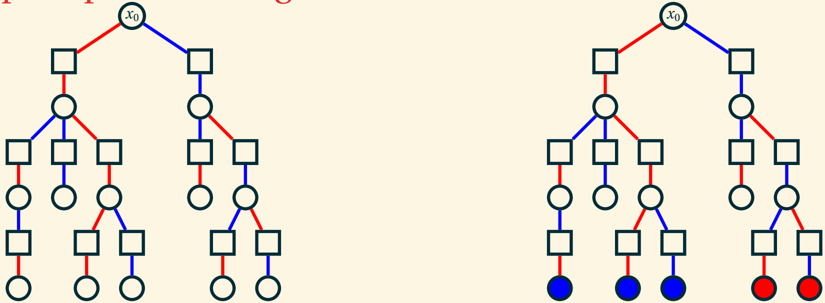

The Gibbs uniqueness property

- $\boldsymbol{T}^{(2\ell)}$ = top 2ℓ levels of $\boldsymbol{T}$
- we are going to show that

$$\lim_{\ell\to\infty} \mathrm{E}\left[\max_{\sigma\in S(\boldsymbol{T}^{(2\ell)})}\left|\mu_{\boldsymbol{T}^{(2\ell)}}(\boldsymbol{\sigma}_{x_0}=1)-\mu_{\boldsymbol{T}^{(2\ell)}}(\boldsymbol{\sigma}_{x_0}=1\mid\boldsymbol{\sigma}_{\partial^{2\ell}x_0}=\sigma_{\partial^{2\ell}x_0})\right|\right]=0$$

Step 2: spatial mixing

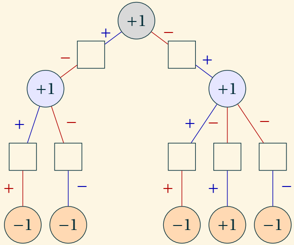

The extremal boundary condition

- given $\boldsymbol{T}^{(2\ell)}$ we construct $\sigma^+ \in S(\boldsymbol{T}^{(2\ell)})$ that maximises

$$\mu_{\boldsymbol{T}^{(2\ell)}}(\boldsymbol{\sigma}_{x_0} = 1 \mid \boldsymbol{\sigma}_{\partial^{2\ell} x_0} = \sigma^+_{\partial^{2\ell} x_0}) = 0$$

- we start by setting $\sigma^+_{x_0} = 1$ and proceed inductively
- given σ^+_x the spins σ^+_y nudge x towards σ^+_x

Step 2: spatial mixing

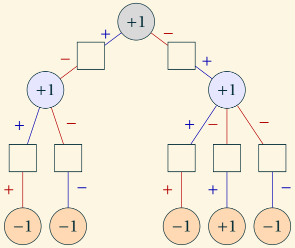

Extremal density evolution

- the process leads to a modified density evolution equation

$$\boldsymbol{\eta}_0 \overset{\mathrm{d}}{=} \sum_{i=1}^{\boldsymbol{d}} \boldsymbol{s}_i \log \frac{1 + \boldsymbol{s}_i \tanh(\boldsymbol{\eta}_i/2)}{2}$$

- the contraction method applies
- we re-discover the solution π_d to the original density evolution
- consequently, $\pi_{\boldsymbol{\Phi}}$ converges to π_d

Step 2: spatial mixing

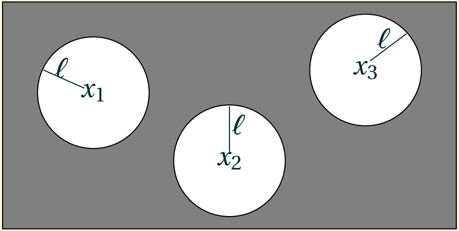

Corollary

For any fixed $k \geq 2$ we have

$$\lim_{n\to\infty} \sum_{\sigma\in\{\pm1\}^k} \mathrm{E}\left|\boldsymbol{\mu}_{\boldsymbol{\Phi}}(\boldsymbol{\sigma}_{x_1}=\sigma_1,\ldots,\boldsymbol{\sigma}_{x_k}=\sigma_k)-\prod_{i=1}^{k}\boldsymbol{\mu}_{\boldsymbol{\Phi}}(\boldsymbol{\sigma}_{x_i}=\sigma_i)\right| = 0$$

Step 3: Aizenman-Sims-Starr

Proposition

We have $\lim_{n\to\infty} \mathrm{E}\left[\log(1 \vee Z(\mathbf{\Phi}_{n+1}))\right] - \mathrm{E}\left[\log(1 \vee Z(\mathbf{\Phi}_n))\right] = \mathscr{B}_d$

Step 3: Aizenman-Sims-Starr

Proposition

We have $\lim_{n\to\infty} \mathrm{E}\left[\log(1 \vee Z(\boldsymbol{\Phi}_{n+1}))\right] - \mathrm{E}\left[\log(1 \vee Z(\boldsymbol{\Phi}_{n}))\right] = \mathscr{B}_d$

Corollary

We have $\lim_{n\to\infty} \frac{1}{n} \mathrm{E}\left[\log(1 \vee Z(\boldsymbol{\Phi}_{n}))\right] = \mathscr{B}_d$

Proof

Just write a telescoping sum

$$\mathrm{E}\left[\log(1 \vee Z(\boldsymbol{\Phi}_{n}))\right] = \sum_{N=1}^{n-1} \mathrm{E}\left[\log(1 \vee Z(\boldsymbol{\Phi}_{N+1}))\right] - \mathrm{E}\left[\log(1 \vee Z(\boldsymbol{\Phi}_{N}))\right]$$

Step 3: Aizenman-Sims-Starr

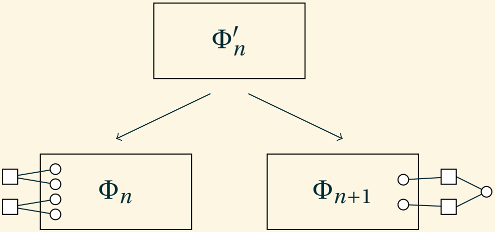

A coupling argument

- let $\mathbf{\Phi}'_n$ comprise $\boldsymbol{m}' \sim \mathrm{Po}(d(n-1)/2)$ random clauses
- obtain $\mathbf{\Phi}_n$ by adding $\mathbf{\Delta}'' \sim \mathrm{Po}(d/2)$ clauses
- to obtain $\mathbf{\Phi}_{n+1}$ add x_{n+1} and $\mathbf{\Delta}''' \sim \mathrm{Po}(d)$ clauses

$$\mathrm{E}\log\frac{Z(\mathbf{\Phi}_n)\vee 1}{Z(\mathbf{\Phi}'_n)\vee 1} \qquad \mathrm{E}\log\frac{Z(\mathbf{\Phi}_{n+1})\vee 1}{Z(\mathbf{\Phi}'_n)\vee 1}$$

Step 3: Aizenman-Sims-Starr

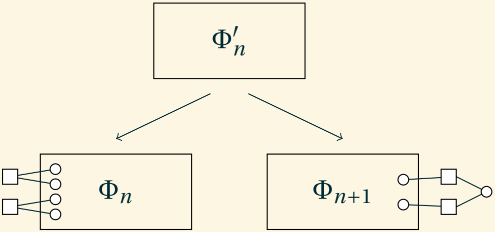

$$\begin{aligned}
\mathrm{E}\log\frac{Z(\boldsymbol{\Phi}_n)\vee 1}{Z(\boldsymbol{\Phi}'_n)\vee 1} &= \mathrm{E}\log\left\langle \prod_{i=1}^{\boldsymbol{\Delta}''} \mathbf{1}\left\{\boldsymbol{\sigma}\models b''_i\right\}, \mu_{\boldsymbol{\Phi}'_n} \right\rangle \\
&= \mathrm{E}\log\left\langle \prod_{i=1}^{\boldsymbol{\Delta}''} 1-\mathbf{1}\{\boldsymbol{\sigma}_{x_{2i-1}}=-\boldsymbol{s}_i, \boldsymbol{\sigma}_{x_{2i}}=-\boldsymbol{s}'_i\}, \mu_{\boldsymbol{\Phi}'_n} \right\rangle \\
&= \frac{d}{2}\mathrm{E}\log\left(1-\boldsymbol{\mu}_{\boldsymbol{\Phi}'_n}(\boldsymbol{\sigma}_{x_1}=1)\boldsymbol{\mu}_{\boldsymbol{\Phi}'_n}(\boldsymbol{\sigma}_{x_2}=1)\right) \\
&\sim \frac{d}{2}\mathrm{E}\log\left(1-\boldsymbol{\mu}_1\boldsymbol{\mu}_2\right)
\end{aligned}$$

Step 3: Aizenman-Sims-Starr

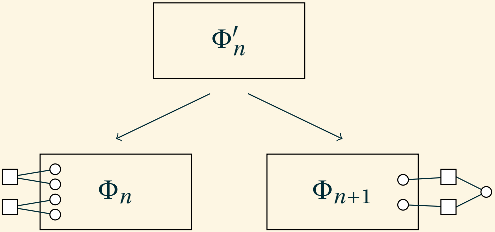

$$\mathrm{E}\log\frac{Z(\boldsymbol{\Phi}_{n+1})\vee 1}{Z(\boldsymbol{\Phi}'_n)\vee 1}\sim \mathrm{E}\left[\log\left(\prod_{i=1}^{\boldsymbol{d}_+}\boldsymbol{\mu}_i+\prod_{i=1}^{\boldsymbol{d}_-}\boldsymbol{\mu}_{i+\boldsymbol{d}_+}\right)\right]$$

Step 4: the interpolation method

Proposition

We have

$$\lim_{n\to\infty} \frac{\log(1 \vee Z(\mathbf{\Phi}))}{n} = \mathscr{B}_d \qquad \text{in probability.}$$

Step 4: the interpolation method

Proof strategy

- show that w.h.p.

$$\frac{1}{n}\log(1 \vee Z(\mathbf{\Phi})) \leq \mathscr{B}_d + o(1)$$

- then the assertion follows because

$$\frac{1}{n}\mathrm{E}[\log(1 \vee Z(\mathbf{\Phi}))] \sim \mathscr{B}_d$$

Step 4: the interpolation method

Soft constraints

- for $0 < \beta < \infty$ introduce

$$Z_\beta(\mathbf{\Phi}) = \sum_{\sigma \in \{\pm 1\}^n} \prod_{i=1}^{\boldsymbol{m}} \exp(-\beta \mathbf{1}\{\sigma \not\models a_i\})$$

- then $Z(\mathbf{\Phi}) \le Z_\beta(\mathbf{\Phi})$ for all $\beta > 0$
- Azuma–Hoeffding $\Rightarrow \log Z_\beta(\mathbf{\Phi}) = \mathrm{E}[\log Z_\beta(\mathbf{\Phi})] + o(n)$ w.h.p.
- hence, it suffices to prove

$$\lim_{\beta \to \infty} \lim_{n \to \infty} \mathrm{E}[\log Z_\beta(\mathbf{\Phi})] \le \mathscr{B}_d$$

Step 4: the interpolation method

Lemma

For any $\beta > 0$ we have

$$\lim_{n\to\infty} \frac{1}{n} \mathrm{E}[\log Z_\beta(\boldsymbol{\Phi})] \leq \mathscr{B}_{d,\beta}$$

where

$$\mathscr{B}_{d,\beta} = \mathrm{E}\left[\log\left(\sum_{s=\pm 1} \prod_{i=1}^{\boldsymbol{d}} 1 - \mathbf{1}\{s \neq \boldsymbol{s}_i\}(1-\mathrm{e}^{-\beta})\boldsymbol{\mu}_i\right)\right] - \frac{d}{2}\mathrm{E}\left[\log\left(1-\left(1-\mathrm{e}^{-\beta}\right)\boldsymbol{\mu}_1\boldsymbol{\mu}_2\right)\right]$$

Step 4: the interpolation method

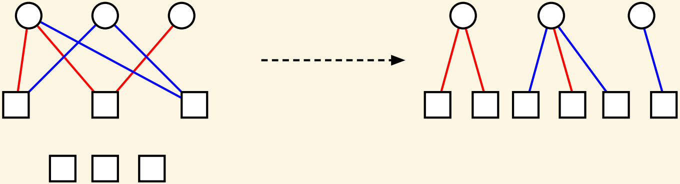

Proof via interpolation [G03,FL03,PT04]

Given $0 \le t \le 1$ introduce a random formula $\mathbf{\Phi}_t$ comprising

- $\boldsymbol{m}_t = \mathrm{Po}((1-t)dn/2)$ random 2-clauses $a_1, \ldots, a_{\boldsymbol{m}_t}$
- $\boldsymbol{m}'_t = \mathrm{Po}(tdn)$ random external fields $b_1, \ldots, b_{\boldsymbol{m}'_t}$
- $\boldsymbol{m}''_t = \mathrm{Po}((1-t)dn/2)$ random constant factors $c_1, \ldots, c_{\boldsymbol{m}''_t}$

Step 4: the interpolation method

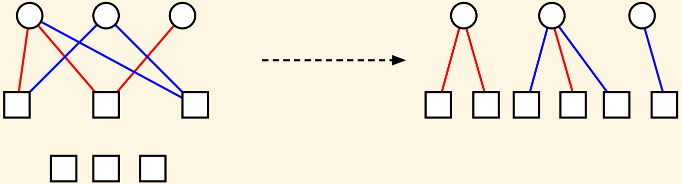

Proof via interpolation [G03,FL03,PT04]

$$Z_\beta(\boldsymbol{\Phi}_t) = \sum_{\sigma \in \{\pm 1\}^n} \prod_{i=1}^{\boldsymbol{m}_t} \exp(-\beta \mathbf{1}\{\sigma \not\models a_i\})$$

$$\times \prod_{i=1}^{\boldsymbol{m}'_t} 1 - (1 - \mathrm{e}^{-\beta}) \mathbf{1}\left\{\boldsymbol{\sigma}_{\partial b_i} = \boldsymbol{s}_i\right\} \boldsymbol{\mu}_i$$

$$\times \prod_{i=1}^{\boldsymbol{m}''_t} 1 - (1 - \mathrm{e}^{-\beta}) \boldsymbol{\mu}'_i \boldsymbol{\mu}''_i$$

Step 4: the interpolation method

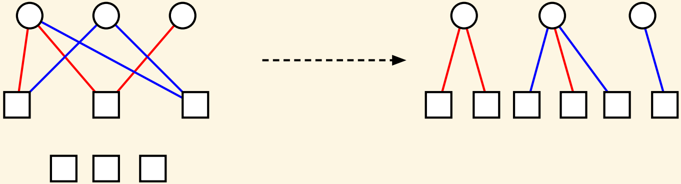

Proof via interpolation [G03,FL03,PT04]

$$\frac{\partial}{\partial t}\mathrm{E}[\log Z_\beta(\mathbf{\Phi}_t)] = \text{sum of squares} \geq 0$$

Outlook

Cavity method predictions

- the diluted mean-field spin glass [MP01,MP03]
- Belief/Survey Propagation [MPZ02]
- RSB and the condensation phase transition [KMRTSZ07]

Outlook

Rigorous results: replica symmetry

- ferromagnetic Ising/Potts [DM10,DMS13]
- random linear equations [DM02,PS16,ACOGM18,COEGHR20]
- condensation [GT04,BCOHRV16,COKPZ18,COEJKK18]

Rigorous results: replica symmetry breaking

- satisfiability phase transitions [COP12,COP16,DSS15]
- Bethe states/variational free energy [P13,DSS16,COP19]

References

- D. Achlioptas, A. Coja-Oghlan, M. Hahn-Klimroth, J. Lee, N. Müller, M. Penschuck, G. Zhou: The random 2-SAT partition function. arXiv:2002.03690
- V. Chvátal, B. Reed: Mick gets some (the odds are on his side). FOCS 1992
- J. Ding, A. Sly, N. Sun: Proof of the satisfiability conjecture for large k. STOC 2015
- A. Goerdt: A threshold for unsatisfiability. JCSS **53** (1996)
- M. Mézard, G. Parisi, R. Zecchina: Analytic and algorithmic solution of random satisfiability problems. Science **297** (2002)
- R. Monasson, R. Zecchina: The entropy of the k-satisfiability problem. Phys Rev Lett **76** (1996)
- D. Panchenko: On the replica symmetric solution of the K-sat model. EJP **19** (2014)
- D. Panchenko, M. Talagrand: Bounds for diluted mean-fields spin glass models. PTRF **130** (2004)