Potts and independent set models on d-regular graphs

Amir Dembo   Andrea Montanari   Allan Sly   Nike Sun

Stanford University   UC Berkeley

Bangalore   January 4 2013

Outline

1. The Potts and independent set models
2. Locally tree-like graphs and the Bethe prediction
3. Previous work and results
4. Verifying the Bethe prediction: proof ideas

Outline

Graphical models

$G = (V, E)$ finite undirected graph

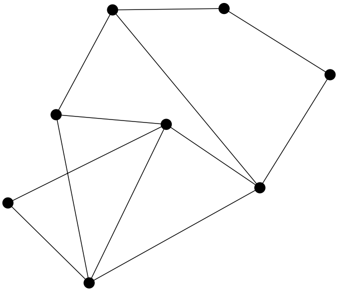

$G = (V, E)$ finite undirected graph

Spin configuration $\underline{\sigma} \in \mathscr{X}^V$
($\mathscr{X}$ finite alphabet)

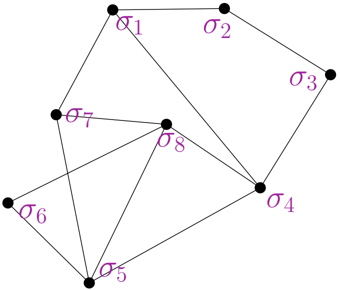

$G = (V, E)$ finite undirected graph

Spin configuration $\underline{\sigma} \in \mathscr{X}^V$
($\mathscr{X}$ finite alphabet)

Graphical model:

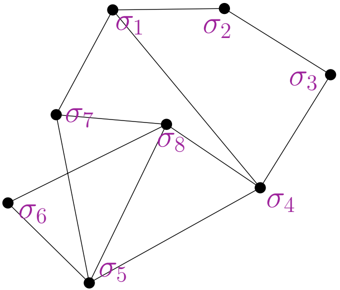

Graphical models

$G = (V, E)$ finite undirected graph

Spin configuration $\underline{\sigma} \in \mathscr{X}^V$
($\mathscr{X}$ finite alphabet)

Graphical model:
Model of random spin configuration

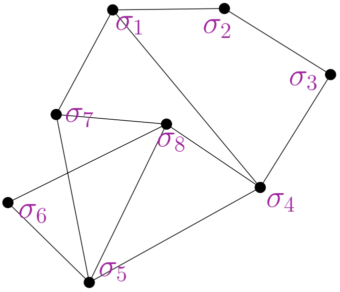

$G = (V, E)$ finite undirected graph

Spin configuration $\underline{\sigma} \in \mathscr{X}^V$
($\mathscr{X}$ finite alphabet)

Graphical model:
Model of random spin configuration defined by **local** interactions

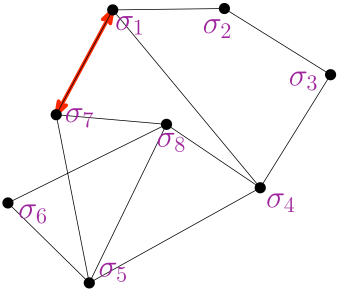

Factor models

Factor models

Factor model on $G = (V, E)$:

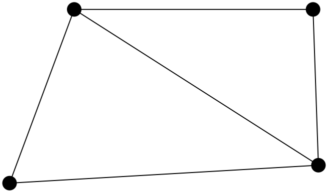

Factor models

Factor model on $G = (V, E)$:

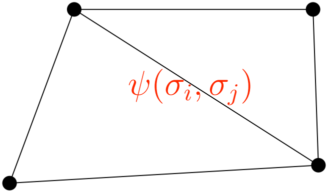

Edge interaction $\psi(\sigma_i, \sigma_j)$,

Factor models

Factor model on $G = (V, E)$:

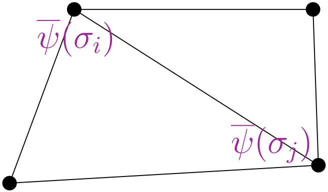

Edge interaction $\psi(\sigma_i, \sigma_j)$, vertex factor $\bar{\psi}(\sigma_i)$ (external field)

Factor model on $G = (V, E)$:

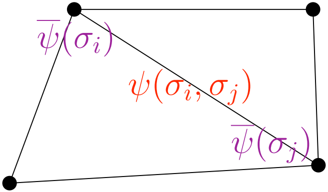

Edge interaction $\psi(\sigma_i, \sigma_j)$, vertex factor $\bar{\psi}(\sigma_i)$ (external field)
Taking product over all edges, vertices gives **factor model**

Factor model on $G = (V, E)$:

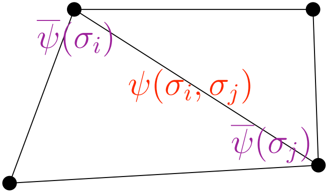

Edge interaction $\psi(\sigma_i, \sigma_j)$, vertex factor $\bar{\psi}(\sigma_i)$ (external field)
Taking product over all edges, vertices gives **factor model**

$$\nu_G(\underline{\sigma}) = \frac{1}{Z} \prod_{(ij)\in E} \psi(\sigma_i, \sigma_j) \prod_{i \in V} \bar{\psi}(\sigma_i)$$

Factor models

Factor model on $G = (V, E)$:

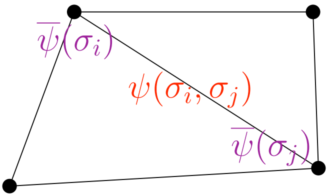

Edge interaction $\psi(\sigma_i, \sigma_j)$, vertex factor $\bar{\psi}(\sigma_i)$ (external field)
Taking product over all edges, vertices gives **factor model**

$$\nu_G(\underline{\sigma}) = \frac{1}{Z} \prod_{(ij)\in E} \psi(\sigma_i, \sigma_j) \prod_{i\in V} \bar{\psi}(\sigma_i)$$

$Z =$ normalizing constant or **partition funtion**

Potts model

Potts model

$$\mathscr{X} = [q] \equiv \{1, \ldots, q\}$$

Potts model

$\mathscr{X} = [q] \equiv \{1, \ldots, q\}$

q-**Potts model** on $G = (V, E)$:

Potts model

$\mathscr{X} = [q] \equiv \{1, \ldots, q\}$

q-**Potts model** on $G = (V, E)$:

$$\nu_G^{\beta,B}(\underline{\sigma}) = \frac{1}{Z_G(\beta, B)} \prod_{(ij)\in E} e^{\beta \mathbf{1}\{\sigma_i = \sigma_j\}} \prod_{i \in V} e^{B \mathbf{1}\{\sigma_i = 1\}}$$

Potts model

$\mathscr{X} = [q] \equiv \{1, \ldots, q\}$

q-**Potts model** on $G = (V, E)$:

$$\nu_G^{\beta,B}(\underline{\sigma}) = \frac{1}{Z_G(\beta, B)} \prod_{(ij)\in E} e^{\beta \mathbf{1}\{\sigma_i=\sigma_j\}} \prod_{i\in V} e^{B\mathbf{1}\{\sigma_i=1\}}$$

- β^{-1} = temperature

$\mathscr{X} = [q] \equiv \{1, \dots, q\}$

q-**Potts model** on $G = (V, E)$:

$$\nu_G^{\beta,B}(\underline{\sigma}) = \frac{1}{Z_G(\beta, B)} \prod_{(ij)\in E} e^{\beta \mathbf{1}\{\sigma_i = \sigma_j\}} \prod_{i\in V} e^{B \mathbf{1}\{\sigma_i = 1\}}$$

- β^{-1} = temperature
- B = external field, in direction of distinguished spin 1

Potts model

$\mathscr{X} = [q] \equiv \{1, \ldots, q\}$

q-**Potts model** on $G = (V, E)$:

$$\nu_G^{\beta,B}(\underline{\sigma}) = \frac{1}{Z_G(\beta, B)} \prod_{(ij)\in E} e^{\beta \mathbf{1}\{\sigma_i = \sigma_j\}} \prod_{i \in V} e^{B \mathbf{1}\{\sigma_i = 1\}}$$

- β^{-1} = temperature
- B = external field, in direction of distinguished spin 1
- $Z_G(\beta, B)$ = partition function

Potts model

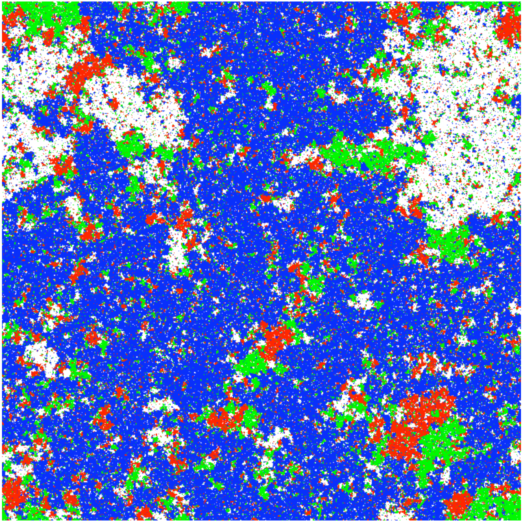

Figure: David Wilson

Potts model

$$\nu_G^{\beta,B}(\underline{\sigma}) = \frac{1}{Z_G(\beta,B)} \prod_{(ij)\in E} e^{\beta \mathbf{1}\{\sigma_i=\sigma_j\}} \prod_{i\in V} e^{B\mathbf{1}\{\sigma_i=1\}}$$

Potts model

$$\nu_G^{\beta,B}(\underline{\sigma}) = \frac{1}{Z_G(\beta,B)} \prod_{(ij)\in E} e^{\beta \mathbf{1}\{\sigma_i=\sigma_j\}} \prod_{i\in V} e^{B\mathbf{1}\{\sigma_i=1\}}$$

- $\beta > 0$ is ferromagnetic; $\beta < 0$ is anti-ferromagnetic (**AF**)

Potts model

$$\nu_G^{\beta,B}(\underline{\sigma}) = \frac{1}{Z_G(\beta,B)} \prod_{(ij)\in E} e^{\beta \mathbf{1}\{\sigma_i=\sigma_j\}} \prod_{i\in V} e^{B\mathbf{1}\{\sigma_i=1\}}$$

- $\beta > 0$ is ferromagnetic; $\beta < 0$ is anti-ferromagnetic (**AF**)
- $q = 2$: **Ising model**

Potts model

$$\nu_G^{\beta,B}(\underline{\sigma}) = \frac{1}{Z_G(\beta,B)} \prod_{(ij)\in E} e^{\beta \mathbf{1}\{\sigma_i=\sigma_j\}} \prod_{i\in V} e^{B\mathbf{1}\{\sigma_i=1\}}$$

- $\beta > 0$ is ferromagnetic; $\beta < 0$ is anti-ferromagnetic (**AF**)
- $q = 2$: **Ising model**
- $\beta = -\infty$: random proper q-colorings

Independent set (hard-core) model

Independent set (hard-core) model

$\mathscr{X} = \{0 = \text{unoccupied}, 1 = \text{occupied}\}$

Independent set (hard-core) model

$\mathscr{X} = \{0 = \text{unoccupied}, 1 = \text{occupied}\}$

The **independent set** (**IS**) or **hard-core model** on $G = (V, E)$:

Independent set (hard-core) model

$\mathscr{X} = \{0 = \text{unoccupied}, 1 = \text{occupied}\}$

The **independent set** (**IS**) or **hard-core model** on $G = (V, E)$:

$$\nu_G^\lambda(\underline{\sigma}) = \frac{1}{Z_G(\lambda)} \prod_{(ij)\in E} \mathbf{1}\{\sigma_i\sigma_j \neq 1\} \prod_{i\in V} \lambda^{\sigma_i}$$

Independent set (hard-core) model

$\mathscr{X} = \{0 = \text{unoccupied}, 1 = \text{occupied}\}$

The **independent set** (**IS**) or **hard-core model** on $G = (V, E)$:

$$\nu_G^\lambda(\underline{\sigma}) = \frac{1}{Z_G(\lambda)} \prod_{(ij)\in E} \mathbf{1}\{\sigma_i\sigma_j \neq 1\} \prod_{i\in V} \lambda^{\sigma_i}$$

- $\mathbf{1}\{\sigma_i\sigma_j \neq 1\}$: hard constraints; repulsive interactions

Independent set (hard-core) model

$\mathscr{X} = \{0 = \text{unoccupied}, 1 = \text{occupied}\}$

The **independent set** (**IS**) or **hard-core model** on $G = (V, E)$:

$$\nu_G^\lambda(\underline{\sigma}) = \frac{1}{Z_G(\lambda)} \prod_{(ij)\in E} \mathbf{1}\{\sigma_i \sigma_j \neq 1\} \prod_{i\in V} \lambda^{\sigma_i}$$

- $\mathbf{1}\{\sigma_i \sigma_j \neq 1\}$: hard constraints; repulsive interactions
- $\lambda =$ fugacity or activity

Independent set (hard-core) model

$\mathscr{X} = \{0 = \text{unoccupied}, 1 = \text{occupied}\}$

The **independent set** (**IS**) or **hard-core model** on $G = (V, E)$:

$$\nu_G^\lambda(\underline{\sigma}) = \frac{1}{Z_G(\lambda)} \prod_{(ij)\in E} \mathbf{1}\{\sigma_i\sigma_j \neq 1\} \prod_{i\in V} \lambda^{\sigma_i}$$

- $\mathbf{1}\{\sigma_i\sigma_j \neq 1\}$: hard constraints; repulsive interactions
- λ = fugacity or activity
- $Z_G(\lambda)$ = partition function,
  with $Z_G(1)$ = number of independent sets

Independent set (hard-core) model

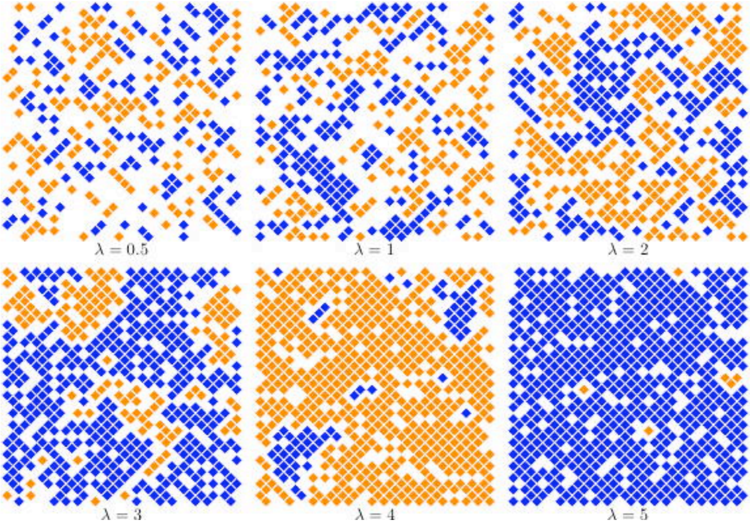

Figure: David Wilson

Free energy density

Free energy density

Consider a sequence of (random) graphs G_n (n vertices) in the thermodynamic limit $n \to \infty$

Free energy density

Consider a sequence of (random) graphs G_n (n vertices) in the thermodynamic limit $n \to \infty$

Asymptotics of partition function $Z_n \equiv Z_{G_n}$?

Free energy density

Consider a sequence of (random) graphs G_n (n vertices) in the thermodynamic limit $n \to \infty$

Asymptotics of partition function $Z_n \equiv Z_{G_n}$?

Free energy $\phi_n \equiv n^{-1}\mathbb{E}_n[\log Z_n]$

Free energy density

Consider a sequence of (random) graphs G_n (n vertices) in the thermodynamic limit $n \to \infty$

Asymptotics of partition function $Z_n \equiv Z_{G_n}$?

Free energy $\phi_n \equiv n^{-1}\mathbb{E}_n[\log Z_n]$

Free energy density $\phi \equiv \lim_{n\to\infty} \phi_n$

Free energy density

Consider a sequence of (random) graphs G_n (n vertices) in the thermodynamic limit $n \to \infty$

Asymptotics of partition function $Z_n \equiv Z_{G_n}$?

Free energy $\phi_n \equiv n^{-1}\mathbb{E}_n[\log Z_n]$

Free energy density $\phi \equiv \lim_{n\to\infty} \phi_n$
($\approx$ exponential growth rate of Z_n)

Consider a sequence of (random) graphs G_n (n vertices) in the thermodynamic limit $n \to \infty$

Asymptotics of partition function $Z_n \equiv Z_{G_n}$?

Free energy $\phi_n \equiv n^{-1}\mathbb{E}_n[\log Z_n]$

Free energy density $\phi \equiv \lim_{n\to\infty} \phi_n$
($\approx$ exponential growth rate of Z_n)

Does ϕ exist?

Free energy density

Consider a sequence of (random) graphs G_n (n vertices) in the thermodynamic limit $n \to \infty$

Asymptotics of partition function $Z_n \equiv Z_{G_n}$?

Free energy $\phi_n \equiv n^{-1}\mathbb{E}_n[\log Z_n]$

Free energy density $\phi \equiv \lim_{n\to\infty} \phi_n$
($\approx$ exponential growth rate of Z_n)

Does ϕ exist? Can its value be computed?

Consider a sequence of (random) graphs G_n (n vertices) in the thermodynamic limit $n \to \infty$

Asymptotics of partition function $Z_n \equiv Z_{G_n}$?

Free energy $\phi_n \equiv n^{-1}\mathbb{E}_n[\log Z_n]$

Free energy density $\phi \equiv \lim_{n\to\infty} \phi_n$
($\approx$ exponential growth rate of Z_n)

Does ϕ exist? Can its value be computed?

The purpose of this work is to give an answer in the setting of locally tree-like graphs

Outline

1. The Potts and independent set models
2. Locally tree-like graphs and the Bethe prediction
3. Previous work and results
4. Verifying the Bethe prediction: proof ideas

Outline

In what sense is the random 3-regular graph

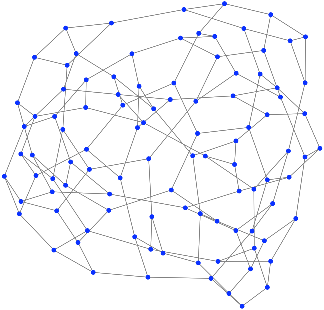

random 3-regular graph

Locally tree-like graphs

In what sense is the random 3-regular graph locally like T_3?

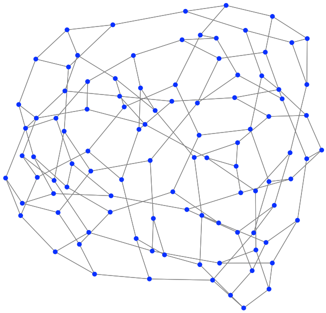

random 3-regular graph

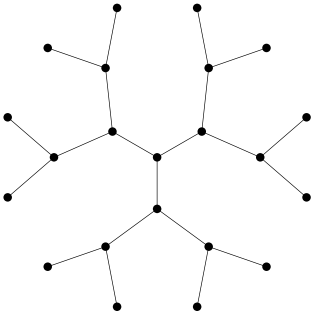

first few levels of T_3

Locally tree-like graphs

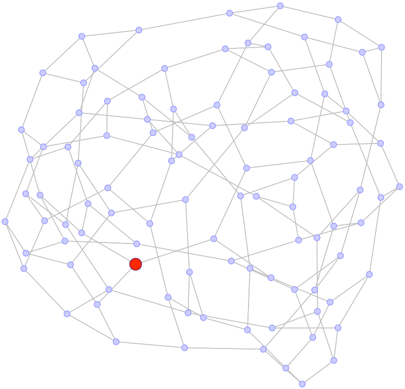

$I_n \in V_n$ unif. random

$B_t(I_n)$ radius t ball about I_n

Isomorphic to T_d^t
(first t levels of T_d)?

Locally tree-like graphs

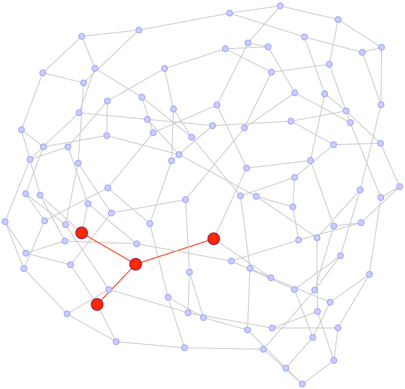

$I_n \in V_n$ unif. random

$B_t(I_n)$ radius t ball about I_n

Isomorphic to T_d^t
(first t levels of T_d)?

Locally tree-like graphs

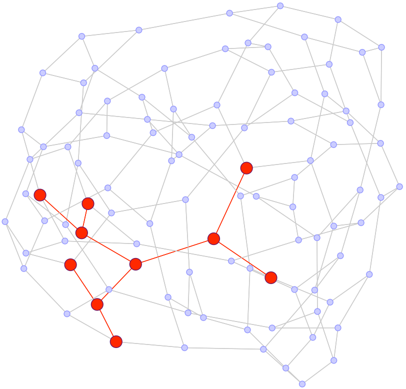

$I_n \in V_n$ unif. random

$B_t(I_n)$ radius t ball about I_n

Isomorphic to T_d^t
(first t levels of T_d)?

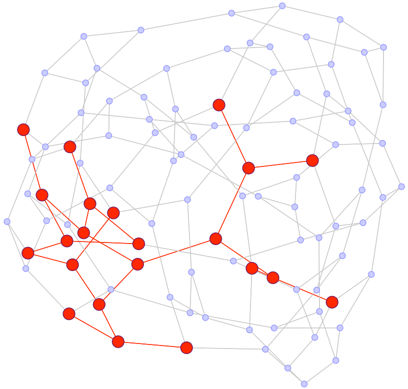

$I_n \in V_n$ unif. random

$B_t(I_n)$ radius t ball about I_n

Isomorphic to T_d^t
(first t levels of T_d)?

Locally tree-like graphs

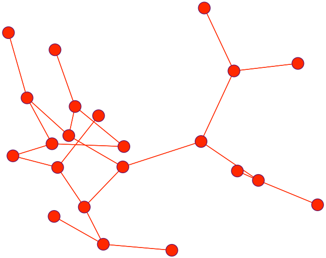

$I_n \in V_n$ unif. random

$B_t(I_n)$ radius t ball about I_n

Isomorphic to T_d^t
(first t levels of T_d)?

Locally tree-like graphs: definition

$G_n = (V_n, E_n)$ random graph sequence

Locally tree-like graphs: definition

$G_n = (V_n, E_n)$ random graph sequence
$I_n \in V_n$ uniformly random vertex

Locally tree-like graphs: definition

$G_n = (V_n, E_n)$ random graph sequence
$I_n \in V_n$ uniformly random vertex

Definition.

$G_n = (V_n, E_n)$ random graph sequence
$I_n \in V_n$ uniformly random vertex

Definition.

> G_n **converges locally** to T_d if for all $t \geq 0$,

Locally tree-like graphs: definition

$G_n = (V_n, E_n)$ random graph sequence
$I_n \in V_n$ uniformly random vertex

Definition.

> G_n **converges locally** to T_d if for all $t \geq 0$,
> $B_t(I_n)$ converges in probability to T_d^t

$G_n = (V_n, E_n)$ random graph sequence
$I_n \in V_n$ uniformly random vertex

Definition.

> G_n **converges locally** to T_d if for all $t \geq 0$,
> $B_t(I_n)$ converges in probability to T_d^t

Notation: $G_n \rightarrow_{loc} T_d$

Locally tree-like graphs: definition

$G_n = (V_n, E_n)$ random graph sequence
$I_n \in V_n$ uniformly random vertex

Definition.

> G_n **converges locally** to T_d if for all $t \geq 0$,
> $B_t(I_n)$ converges in probability to T_d^t

Notation: $G_n \to_{loc} T_d$

[Can also make definition
with general (random) limiting tree]

Examples.

Examples.

The random d-regular graph converges locally to T_d

Locally tree-like graphs: examples

Examples.

The random d-regular graph converges locally to T_d
More generally, so does the random k-partite d-regular graph

Examples.

The random d-regular graph converges locally to T_d
More generally, so does the random k-partite d-regular graph
The Erdős-Rényi graph $G(n, \gamma/n)$ converges locally to
the $\mathrm{Pois}(\gamma)$ Galton–Watson tree

Examples.

The random d-regular graph converges locally to T_d
More generally, so does the random k-partite d-regular graph
The Erdős-Rényi graph $G(n, \gamma/n)$ converges locally to
the $\mathrm{Pois}(\gamma)$ Galton–Watson tree
T_d^t does not converge locally to T_d,

Examples.

The random d-regular graph converges locally to T_d
More generally, so does the random k-partite d-regular graph
The Erdős-Rényi graph $G(n, \gamma/n)$ converges locally to
the $\mathrm{Pois}(\gamma)$ Galton–Watson tree
T_d^t does not converge locally to T_d,
but rather to the random d-canopy tree

Locally tree-like graphs: examples

Examples.

The random d-regular graph converges locally to T_d
More generally, so does the random k-partite d-regular graph
The Erdős-Rényi graph $G(n, \gamma/n)$ converges locally to
the $\mathrm{Pois}(\gamma)$ Galton–Watson tree
T_d^t does not converge locally to T_d,
but rather to the random d-canopy tree

Local weak limits are unimodular measures
on the space of rooted graphs.

The Bethe prediction

For factor models on graph sequence $G_n \to_{loc} T$,

The Bethe prediction

For factor models on graph sequence $G_n \to_{loc} T$,

non-rigorous methods of statistical physics give an explicit prediction for free energy density $\phi \equiv \lim_n n^{-1}\mathbb{E}_n[\log Z_n]$:

The Bethe prediction

For factor models on graph sequence $G_n \to_{loc} T$,

non-rigorous methods of statistical physics give an explicit prediction for free energy density $\phi \equiv \lim_n n^{-1}\mathbb{E}_n[\log Z_n]$:

the **Bethe prediction** (or replica symmetric solution)

The Bethe prediction

For factor models on graph sequence $G_n \to_{loc} T$,

non-rigorous methods of statistical physics give an explicit prediction for free energy density $\phi \equiv \lim_n n^{-1}\mathbb{E}_n[\log Z_n]$:

the **Bethe prediction** (or replica symmetric solution)

Bethe prediction is defined only in terms of limiting tree — not the finite graphs G_n

The Bethe prediction: definition

The Bethe prediction: definition

Bethe prediction for factor models on $G_n \to_{loc} T_d$:

The Bethe prediction: definition

Bethe prediction for factor models on $G_n \to_{loc} T_d$:
$\phi \equiv \lim_n n^{-1}\mathbb{E}_n[\log Z_n]$ exists

The Bethe prediction: definition

Bethe prediction for factor models on $G_n \to_{loc} T_d$:
$\phi \equiv \lim_n n^{-1}\mathbb{E}_n[\log Z_n]$ exists and equals the Bethe free energy

The Bethe prediction: definition

Bethe prediction for factor models on $G_n \to_{loc} T_d$:
$\phi \equiv \lim_n n^{-1}\mathbb{E}_n[\log Z_n]$ exists and equals the Bethe free energy

$$\Phi \equiv \Phi(h)$$

The Bethe prediction: definition

Bethe prediction for factor models on $G_n \to_{loc} T_d$:
$\phi \equiv \lim_n n^{-1}\mathbb{E}_n[\log Z_n]$ exists and equals the Bethe free energy

$$\Phi \equiv \Phi(h)$$

for $h \in \Delta$ ($\mathscr{X}$-simplex) a distinguished fixed point of the **Bethe** or **belief propagation** (BP) **recursion**:

The Bethe prediction: definition

Bethe prediction for factor models on $G_n \to_{loc} T_d$:
$\phi \equiv \lim_n n^{-1}\mathbb{E}_n[\log Z_n]$ exists and equals the Bethe free energy

$$\Phi \equiv \Phi(h)$$

for $h \in \Delta$ ($\mathscr{X}$-simplex) a distinguished fixed point of the
Bethe or **belief propagation** (BP) **recursion**:

$$h(\sigma) \cong \bar{\psi}(\sigma) \left(\sum_{\sigma'} \psi(\sigma, \sigma') h(\sigma') \right)^{d-1}$$

The Bethe prediction: functional form

The Bethe prediction: functional form

Functional form:

The Bethe prediction: functional form

Functional form: $\Phi \equiv \Phi^{\mathsf{vx}} - \Phi^{\mathsf{e}}$ where

The Bethe prediction: functional form

Functional form: $\Phi \equiv \Phi^{\mathsf{vx}} - \Phi^{\mathsf{e}}$ where

$$\Phi^{\mathsf{vx}} \equiv \log \left\{ \sum_{\sigma} \bar{\psi}(\sigma) \left(\sum_{\sigma'} \psi(\sigma, \sigma') h(\sigma') \right)^{d} \right\}$$

The Bethe prediction: functional form

Functional form: $\Phi \equiv \Phi^{\mathsf{vx}} - \Phi^{\mathsf{e}}$ where

$$\Phi^{\mathsf{vx}} \equiv \log\left\{\sum_{\sigma} \bar{\psi}(\sigma)\left(\sum_{\sigma'} \psi(\sigma,\sigma')h(\sigma')\right)^{d}\right\}$$

$$\Phi^{\mathsf{e}} \equiv \frac{d}{2}\log\left\{\sum_{\sigma,\sigma'} \psi(\sigma,\sigma')h(\sigma)h(\sigma')\right\}$$

Interpretation of the BP fixed point:

Interpretation of the BP fixed point:

Suppose factor model ν_n on G_n has local weak limit ν —

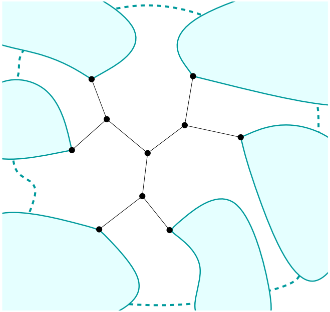

Interpretation of the BP fixed point:
Suppose factor model ν_n on G_n has local weak limit ν — trans.-inv. Gibbs measure for factor model on T_d

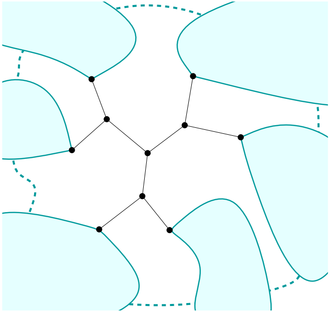

Interpretation of the BP fixed point:
Suppose factor model ν_n on G_n has local weak limit ν — trans.-inv. Gibbs measure for factor model on T_d

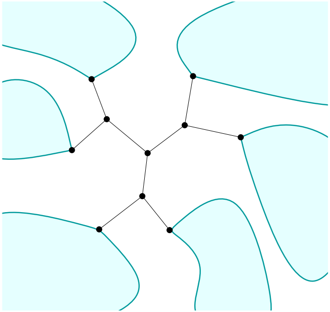

Ignore long cycles

Interpretation of the BP fixed point:

Suppose factor model ν_n on G_n has local weak limit ν — trans.-inv. Gibbs measure for factor model on T_d

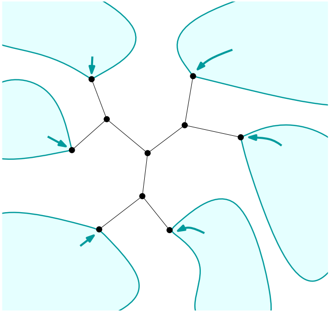

Ignore long cycles

Boundary data $\approx$ i.i.d.

Interpretation of the BP fixed point:
Suppose factor model ν_n on G_n has local weak limit ν — trans.-inv. Gibbs measure for factor model on T_d

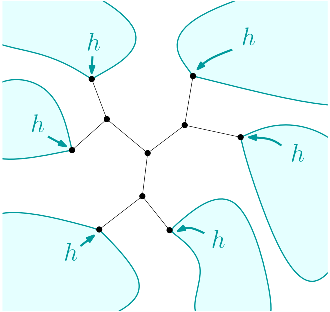

Ignore long cycles

Boundary data ≈ i.i.d.

Marginal of ν on $U \approx \nu(\underline{\sigma}_U | \underline{\sigma}_{\partial U})$

Interpretation of the BP fixed point:
Suppose factor model ν_n on G_n has local weak limit ν — trans.-inv. Gibbs measure for factor model on T_d

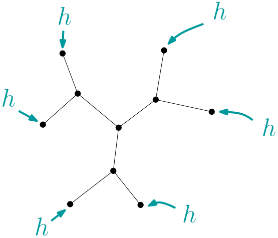

Ignore long cycles

Boundary data $\approx$ i.i.d.

Marginal of ν on $U \approx$ $\nu(\underline{\sigma}_U | \underline{\sigma}_{\partial U}) \times \prod_{v \in \partial U} h(\sigma)$

Interpretation of the BP fixed point:
Suppose factor model ν_n on G_n has local weak limit ν — trans.-inv. Gibbs measure for factor model on T_d

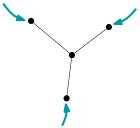

Ignore long cycles

Boundary data $\approx$ i.i.d.

Marginal of ν on U $\approx$ $\nu(\underline{\sigma}_U | \underline{\sigma}_{\partial U}) \times \prod_{v \in \partial U} h(\sigma)$

Interpretation of the BP fixed point:
Suppose factor model ν_n on G_n has local weak limit ν — trans.-inv. Gibbs measure for factor model on T_d

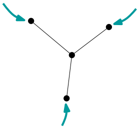

Ignore long cycles

Boundary data $\approx$ i.i.d.

Marginal of ν on $U \approx$ $\nu(\underline{\sigma}_U | \underline{\sigma}_{\partial U}) \times \prod_{v \in \partial U} h(\sigma)$

Consistent family of marginals

The Bethe prediction: interpretation of BP recursion

Interpretation of the BP fixed point:
Suppose factor model ν_n on G_n has local weak limit ν — trans.-inv. Gibbs measure for factor model on T_d

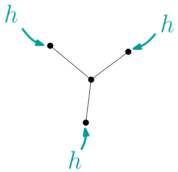

Ignore long cycles

Boundary data $\approx$ i.i.d.

Marginal of ν on $U \approx$
$\nu(\underline{\sigma}_U | \underline{\sigma}_{\partial U}) \times \prod_{v \in \partial U} h(\sigma)$

Consistent family of marginals precisely when h is a BP fixed point

Interpretation of the BP fixed point:
Suppose factor model ν_n on G_n has local weak limit ν — trans.-inv. Gibbs measure for factor model on T_d

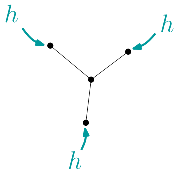

Ignore long cycles

Boundary data ≈ i.i.d.

Marginal of ν on U ≈ $\nu(\underline{\sigma}_U|\underline{\sigma}_{\partial U})\times\prod_{v\in\partial U} h(\sigma)$

Consistent family of marginals precisely when h is a BP fixed point

BP fixed point h

Interpretation of the BP fixed point:
Suppose factor model ν_n on G_n has local weak limit ν — trans.-inv. Gibbs measure for factor model on T_d

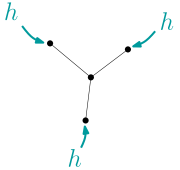

Ignore long cycles

Boundary data ≈ i.i.d.

Marginal of ν on $U \approx$ $\nu(\underline{\sigma}_U | \underline{\sigma}_{\partial U}) \times \prod_{v \in \partial U} h(\sigma)$

Consistent family of marginals precisely when h is a BP fixed point

BP fixed point h
$\longleftrightarrow$ $\nu \equiv \nu_h$ **candidate local weak limit of** ν_n

The Bethe prediction: interpretation of function Φ

BP fixed point $h \longleftrightarrow \nu \equiv \nu_h$ candidate local weak limit of ν_n

BP fixed point $h \longleftrightarrow \nu \equiv \nu_h$ candidate local weak limit of ν_n
Heuristic for $\Phi \equiv \Phi^{\mathsf{vx}} - \Phi^{\mathsf{e}}$:

The Bethe prediction: interpretation of function Φ

BP fixed point $h \longleftrightarrow \nu \equiv \nu_h$ candidate local weak limit of ν_n
Heuristic for $\Phi \equiv \Phi^{\mathsf{vx}} - \Phi^{\mathsf{e}}$:

The Bethe prediction: interpretation of function Φ

BP fixed point $h \longleftrightarrow \nu \equiv \nu_h$ candidate local weak limit of ν_n
Heuristic for $\Phi \equiv \Phi^{\mathsf{vx}} - \Phi^{\mathsf{e}}$:

$\Phi^{\mathsf{vx}}(h) =$ log-partition of marginal of ν_h on star graph T_d^1

BP fixed point $h \longleftrightarrow \nu \equiv \nu_h$ candidate local weak limit of ν_n
Heuristic for $\Phi \equiv \Phi^{\mathsf{vx}} - \Phi^{\mathsf{e}}$:

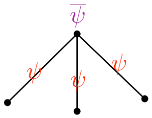

$\Phi^{\mathsf{vx}}(h) =$ log-partition of marginal of ν_h on star graph T_d^1

BP fixed point $h \longleftrightarrow \nu \equiv \nu_h$ candidate local weak limit of ν_n
Heuristic for $\Phi \equiv \Phi^{\mathsf{vx}} - \Phi^{\mathsf{e}}$:

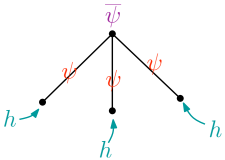

$\Phi^{\mathsf{vx}}(h) =$ log-partition of marginal of ν_h on star graph T_d^1

The Bethe prediction: interpretation of function Φ

BP fixed point $h \longleftrightarrow \nu \equiv \nu_h$ candidate local weak limit of ν_n
Heuristic for $\Phi \equiv \Phi^{\mathsf{vx}} - \Phi^{\mathsf{e}}$:

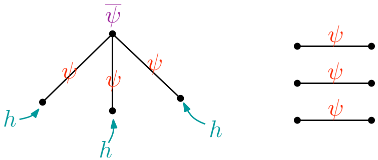

$\Phi^{\mathsf{vx}}(h) =$ log-partition of marginal of ν_h on star graph T_d^1
$\Phi^{\mathsf{e}}(h) = 1/2$ log-partition on d disjoint edges

The Bethe prediction: interpretation of function Φ

BP fixed point $h \longleftrightarrow \nu \equiv \nu_h$ candidate local weak limit of ν_n
Heuristic for $\Phi \equiv \Phi^{\mathsf{vx}} - \Phi^{\mathsf{e}}$:

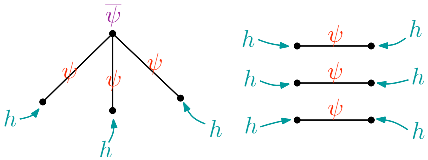

$\Phi^{\mathsf{vx}}(h) =$ log-partition of marginal of ν_h on star graph T_d^1
$\Phi^{\mathsf{e}}(h) = 1/2$ log-partition on d disjoint edges

The Bethe prediction: interpretation of function Φ

BP fixed point $h \longleftrightarrow \nu \equiv \nu_h$ candidate local weak limit of ν_n
Heuristic for $\Phi \equiv \Phi^{\mathsf{vx}} - \Phi^{\mathsf{e}}$:

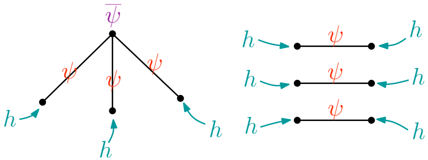

$\Phi^{\mathsf{vx}}(h) =$ log-partition of marginal of ν_h on star graph T_d^1
$\Phi^{\mathsf{e}}(h) = 1/2$ log-partition on d disjoint edges

> If G_n is finite tree, can compute $\log Z_n$ by cutting edges recursively

The Bethe prediction: interpretation of function Φ

BP fixed point $h \longleftrightarrow \nu \equiv \nu_h$ candidate local weak limit of ν_n
Heuristic for $\Phi \equiv \Phi^{\mathsf{vx}} - \Phi^{\mathsf{e}}$:

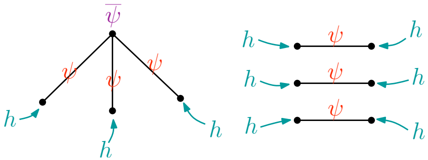

$\Phi^{\mathsf{vx}}(h) =$ log-partition of marginal of ν_h on star graph T_d^1
$\Phi^{\mathsf{e}}(h) = 1/2$ log-partition on d disjoint edges

> If G_n is finite tree, can compute $\log Z_n$ by cutting edges recursively
> n stars, contribution $\approx n \cdot \Phi^{\mathsf{vx}}$

The Bethe prediction: interpretation of function Φ

BP fixed point $h \longleftrightarrow \nu \equiv \nu_h$ candidate local weak limit of ν_n
Heuristic for $\Phi \equiv \Phi^{\text{vx}} - \Phi^{\text{e}}$:

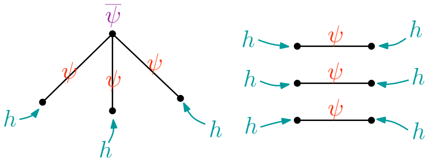

$\Phi^{\text{vx}}(h) =$ log-partition of marginal of ν_h on star graph T_d^1
$\Phi^{\text{e}}(h) = 1/2$ log-partition on d disjoint edges

> If G_n is finite tree, can compute $\log Z_n$ by cutting edges recursively
> n stars, contribution $\approx n \cdot \Phi^{\text{vx}}$
> $nd/2$ edges, each participating in two stars:
> to correct for overcounting, subtract $n \cdot \Phi^{\text{e}}$

The Bethe prediction: interpretation of function Φ

BP fixed point $h \longleftrightarrow \nu \equiv \nu_h$ candidate local weak limit of ν_n
Heuristic for $\Phi \equiv \Phi^{\mathsf{vx}} - \Phi^{\mathsf{e}}$:

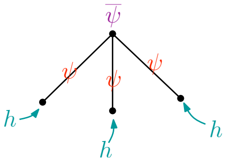

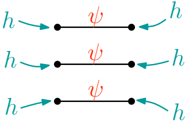

$\Phi^{\mathsf{vx}}(h) =$ log-partition of marginal of ν_h on star graph T_d^1
$\Phi^{\mathsf{e}}(h) = 1/2$ log-partition on d disjoint edges

> If G_n is finite tree, can compute $\log Z_n$ by cutting edges recursively
> - n stars, contribution $\approx n \cdot \Phi^{\mathsf{vx}}$
> - $nd/2$ edges, each participating in two stars:
>   - to correct for overcounting, subtract $n \cdot \Phi^{\mathsf{e}}$

Only a heuristic: G_n are typically not trees!

The Bethe prediction: multiple fixed points

The Bethe prediction: multiple fixed points

$\Phi(h) \equiv \Phi(\nu_h)$ **is (heuristic) formula for** ϕ **assuming** $\nu_n \to_{\mathsf{loc}} \nu_h$

The Bethe prediction: multiple fixed points

$\Phi(h) \equiv \Phi(\nu_h)$ **is (heuristic) formula for** ϕ **assuming** $\nu_n \to_{loc} \nu_h$

If BP fixed point h is **non-unique**, assume $\nu_n \to_{loc}$ mixture(ν_h)

The Bethe prediction: multiple fixed points

$\Phi(h) \equiv \Phi(\nu_h)$ **is (heuristic) formula for** ϕ **assuming** $\nu_n \to_{loc} \nu_h$

If BP fixed point h is **non-unique**, assume $\nu_n \to_{loc}$ mixture(ν_h)

Bethe prediction becomes supremum of $\Phi(h)$ over fixed points h

The Bethe prediction: some remarks

The Bethe prediction: some remarks

Bethe prediction computation **only** involves infinite tree T_d,
not specific graph sequence G_n

The Bethe prediction: some remarks

Bethe prediction computation **only** involves infinite tree T_d,
not specific graph sequence G_n

Can make prediction for general (random) limiting trees

The Bethe prediction: some remarks

Bethe prediction computation **only** involves infinite tree T_d,
not specific graph sequence G_n

Can make prediction for general (random) limiting trees

Bethe prediction is "replica symmetric" in the sense that there is a fixed Gibbs measure ν in definition of $\Phi(\nu)$

The Bethe prediction: some remarks

Bethe prediction computation **only** involves infinite tree T_d,
not specific graph sequence G_n
Can make prediction for general (random) limiting trees

Bethe prediction is "replica symmetric" in the sense that there is a fixed Gibbs measure ν in definition of $\Phi(\nu)$
— equivalently, take *same* h at each boundary vertex

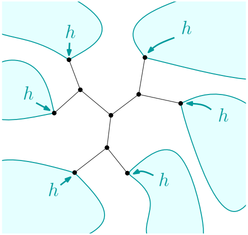

Bethe prediction for ferromagnetic Potts

Bethe prediction specialized to ferromagnetic Potts:

Bethe prediction for ferromagnetic Potts

Bethe prediction specialized to ferromagnetic Potts:

Translation-invariant Gibbs measures

Bethe prediction specialized to ferromagnetic Potts:

Translation-invariant Gibbs measures
ν^{f} (free) and ν^1 (maximally 1-biased)

Bethe prediction for ferromagnetic Potts

Bethe prediction specialized to ferromagnetic Potts:

Translation-invariant Gibbs measures
ν^{f} (free) and ν^{1} (maximally 1-biased)

Bethe prediction is

Bethe prediction for ferromagnetic Potts

Bethe prediction specialized to ferromagnetic Potts:

Translation-invariant Gibbs measures
ν^{f} (free) and ν^1 (maximally 1-biased)

Bethe prediction is $\Phi(\nu^{\mathsf{f}}) \vee \Phi(\nu^1)$

Bethe prediction for AF two-spin systems

Bethe prediction specialized to IS and AF Ising:

Bethe prediction for AF two-spin systems

Bethe prediction specialized to IS and AF Ising:

For G_n bipartite, local weak limits of ν_n
need only be semi-trans.-inv.

Bethe prediction for AF two-spin systems

Bethe prediction specialized to IS and AF Ising:

For G_n bipartite, local weak limits of ν_n
need only be semi-trans.-inv.
Extremal semi-trans.-inv. Gibbs measures

Bethe prediction specialized to IS and AF Ising:

For G_n bipartite, local weak limits of ν_n
need only be semi-trans.-inv.
Extremal semi-trans.-inv. Gibbs measures
ν^0, ν^1 which disagree in **non-uniqueness regimes**

Bethe prediction specialized to IS and AF Ising:

For G_n bipartite, local weak limits of ν_n
need only be semi-trans.-inv.
Extremal semi-trans.-inv. Gibbs measures
ν^0, ν^1 which disagree in **non-uniqueness regimes**
Bethe prediction is $\Phi = \Phi(\nu^0) = \Phi(\nu^1)$

Bethe prediction specialized to IS and AF Ising:

For G_n bipartite, local weak limits of ν_n
need only be semi-trans.-inv.
Extremal semi-trans.-inv. Gibbs measures
ν^0, ν^1 which disagree in **non-uniqueness regimes**
Bethe prediction is $\Phi = \Phi(\nu^0) = \Phi(\nu^1)$

For G_n non-bipartite, same prediction believed to hold
in uniqueness regimes only

Outline

1. The Potts and independent set models
2. Locally tree-like graphs and the Bethe prediction
3. Previous work and results
4. Verifying the Bethe prediction: proof ideas

Outline

Ferromagnetic Ising:

Previous work: ferromagnetic Ising

Ferromagnetic Ising:

[Dembo–Montanari AAP '10] verified Bethe prediction for all $\beta \geq 0$, $B \in \mathbb{R}$, for graphs converging locally to Galton-Watson trees

Ferromagnetic Ising:

[Dembo–Montanari AAP '10] verified Bethe prediction for all $\beta \geq 0$, $B \in \mathbb{R}$, for graphs converging locally to Galton-Watson trees

Moment condition on root vertex degree later removed [Dommers–Giardinà–van der Hofstad JSP '10]

Previous work: ferromagnetic Ising

Ferromagnetic Ising:

[Dembo–Montanari AAP '10] verified Bethe prediction for all $\beta \geq 0$, $B \in \mathbb{R}$, for graphs converging locally to Galton-Watson trees

Moment condition on root vertex degree later removed [Dommers–Giardinà–van der Hofstad JSP '10]

Proofs use an interpolation scheme, comparing $\partial_\beta \phi_n$ with $\partial_\beta \Phi$

Results: Ferro. Potts on general limiting tree

We developed a generalized interpolation scheme to show:

We developed a generalized interpolation scheme to show:

Theorem (Dembo, Montanari, Sun '11).

We developed a generalized interpolation scheme to show:

Theorem (Dembo, Montanari, Sun '11).

> The Bethe prediction $\phi = \Phi$ holds on locally tree-like graphs with general limiting tree for:

Results: Ferro. Potts on general limiting tree

We developed a generalized interpolation scheme to show:

Theorem (Dembo, Montanari, Sun '11).

> The Bethe prediction $\phi = \Phi$ holds on locally tree-like graphs with general limiting tree for:
>
> Ferro. Ising at any $B \in \mathbb{R}$

Results: Ferro. Potts on general limiting tree

We developed a generalized interpolation scheme to show:

Theorem (Dembo, Montanari, Sun '11).

> The Bethe prediction $\phi = \Phi$ holds on locally tree-like graphs with general limiting tree for:
>
> Ferro. Ising at any $B \in \mathbb{R}$
>
> Ferro. Potts at $B \geq 0$ with β sufficiently low (high)

We developed a generalized interpolation scheme to show:

Theorem (Dembo, Montanari, Sun '11).

> The Bethe prediction $\phi = \Phi$ holds on locally tree-like graphs with general limiting tree for:
> Ferro. Ising at any $B \in \mathbb{R}$
> Ferro. Potts at $B \geq 0$ with β sufficiently low (high)

$(\nu^{\mathrm{f}} = \nu^{1})$

Results: Potts on T_d

Results: Potts on T_d

Can obtain sharper results when $G_n \to_{loc} T_d$:

Results: Potts on T_d

Can obtain sharper results when $G_n \to_{\mathit{loc}} T_d$:

Theorem (Dembo, Montanari, Sun '11).

Can obtain sharper results when $G_n \to_{loc} T_d$:

Theorem (Dembo, Montanari, Sun '11).

> For Potts model on $G_n \to_{loc} T_d$,

Can obtain sharper results when $G_n \to_{loc} T_d$:

Theorem (Dembo, Montanari, Sun '11).

For Potts model on $G_n \to_{loc} T_d$,
$\liminf_n \phi_n \geq \Phi$ for all $\beta, B \geq 0$.

Can obtain sharper results when $G_n \to_{loc} T_d$:

Theorem (Dembo, Montanari, Sun '11).

For Potts model on $G_n \to_{loc} T_d$,
$\liminf_n \phi_n \geq \Phi$ for all $\beta, B \geq 0$.

Theorem (Dembo, Montanari, Sly, Sun '12).

Can obtain sharper results when $G_n \to_{loc} T_d$:

Theorem (Dembo, Montanari, Sun '11).

> For Potts model on $G_n \to_{loc} T_d$,
> $\liminf_n \phi_n \geq \Phi$ for all $\beta, B \geq 0$.

Theorem (Dembo, Montanari, Sly, Sun '12).

> For Potts model on $G_n \to_{loc} T_d$ with d even,

Results: Potts on T_d

Can obtain sharper results when $G_n \to_{loc} T_d$:

Theorem (Dembo, Montanari, Sun '11).

For Potts model on $G_n \to_{loc} T_d$,
$\liminf_n \phi_n \geq \Phi$ for all $\beta, B \geq 0$.

Theorem (Dembo, Montanari, Sly, Sun '12).

For Potts model on $G_n \to_{loc} T_d$ with d even,
$\phi = \Phi$ for all $\beta, B \geq 0$.

Previous work: AF two-spin free energy density

IS, AF Ising:

Previous work: AF two-spin free energy density

IS, AF Ising:

Bethe prediction $\phi = \Phi$ holds for random regular graphs below uniqueness threshold
[Bandyopadhyay–Gamarnik SODA '06]

Previous work: AF two-spin free energy density

IS, AF Ising:

Bethe prediction $\phi = \Phi$ holds for random regular graphs below uniqueness threshold
[Bandyopadhyay–Gamarnik SODA '06]

Existence of ϕ for random regular graphs and Erdős-Rényi graphs
[Bayati–Gamarnik–Tetali STOC '10]

Results: AF two-spin free energy density

Results: AF two-spin free energy density

Theorem (Dembo, Montanari, Sun '11).

Results: AF two-spin free energy density

Theorem (Dembo, Montanari, Sun '11).

> The Bethe prediction $\phi = \Phi$ holds on locally tree-like graphs for the IS model at sufficiently low λ.

Theorem (Dembo, Montanari, Sun '11).

> The Bethe prediction $\phi = \Phi$ holds on locally tree-like graphs for the IS model at sufficiently low λ.

($\nu^0 = \nu^1$ & reg. conds)

Theorem (Dembo, Montanari, Sun '11).

> The Bethe prediction $\phi = \Phi$ holds on locally tree-like graphs for the IS model at sufficiently low λ.

($\nu^0 = \nu^1$ & reg. conds)

Theorem (Sly, Sun '12).

Theorem (Dembo, Montanari, Sun '11).

The Bethe prediction $\phi = \Phi$ holds on locally tree-like graphs for the IS model at sufficiently low λ.

($\nu^0 = \nu^1$ & reg. conds)

Theorem (Sly, Sun '12).

For the Ising and IS models on $G_n \to_{loc} T_d$ with G_n bipartite, $\phi = \Phi$ for all parameter values.

Two-spin systems — algorithmic results:

Two-spin systems — algorithmic results:

Ferromagnetic:

Two-spin systems — algorithmic results:

Ferromagnetic:
FPRAS for ferro. Ising at all temperatures, arbitrary magnetic field [Jerrum–Sinclair ALP '90]

Two-spin systems — algorithmic results:

Ferromagnetic:

FPRAS for ferro. Ising at all temperatures, arbitrary magnetic field [Jerrum–Sinclair ALP '90]

Anti-ferromagnetic:

Two-spin systems — algorithmic results:

Ferromagnetic:
FPRAS for ferro. Ising at all temperatures, arbitrary magnetic field [Jerrum–Sinclair ALP '90]

Anti-ferromagnetic:
AF two-spin systems have uniqueness thresholds on T_d:

Two-spin systems — algorithmic results:

Ferromagnetic:

FPRAS for ferro. Ising at all temperatures, arbitrary magnetic field [Jerrum–Sinclair ALP '90]

Anti-ferromagnetic:

AF two-spin systems have uniqueness thresholds on T_d:
$\lambda_c(d)$ for IS, $\beta_c^{\mathrm{af}}(B,d) < 0$ for AF Ising

Two-spin systems — algorithmic results:

Ferromagnetic:
FPRAS for ferro. Ising at all temperatures, arbitrary magnetic field [Jerrum–Sinclair ALP '90]

Anti-ferromagnetic:
AF two-spin systems have uniqueness thresholds on T_d:
$\lambda_c(d)$ for IS, $\beta_c^{\mathrm{af}}(B,d) < 0$ for AF Ising

- FPTAS for IS partition function $Z_G(\lambda)$ on bdd. deg. graphs, $\lambda < \lambda_c(d)$ [Weitz STOC '06]

Two-spin systems — algorithmic results:

Ferromagnetic:

FPRAS for ferro. Ising at all temperatures, arbitrary magnetic field [Jerrum–Sinclair ALP '90]

Anti-ferromagnetic:

AF two-spin systems have uniqueness thresholds on T_d:
$\lambda_c(d)$ for IS, $\beta_c^{\mathrm{af}}(B,d) < 0$ for AF Ising

- FPTAS for IS partition function $Z_G(\lambda)$ on bdd. deg. graphs, $\lambda < \lambda_c(d)$ [Weitz STOC '06]
- FPTAS for AF Ising partition function $Z_G(\beta, B)$ on bdd. deg. graphs, $\beta_c^{\mathrm{af}}(B,d) < \beta < 0$ [Sinclair–Srivastava–Thurley '11]

Complexity of AF two-spin systems

Hardness results for IS:

Hardness results for IS:

$Z_G(\lambda)$ hard to approximate on d-regular graphs when
$\lambda > c/d$ [Luby–Vigoda STOC '97];

Complexity of AF two-spin systems

Hardness results for IS:

$Z_G(\lambda)$ hard to approximate on d-regular graphs when

$\lambda > c/d$ [Luby–Vigoda STOC '97];

$\lambda = 1$ and $d > 25$ [Dyer–Frieze–Jerrum FOCS '99]

Complexity of AF two-spin systems

Hardness results for IS:

$Z_G(\lambda)$ hard to approximate on d-regular graphs when
$\lambda > c/d$ [Luby–Vigoda STOC '97];
$\lambda = 1$ and $d > 25$ [Dyer–Frieze–Jerrum FOCS '99]

Phase transition at $\lambda_c(d)$:

Complexity of AF two-spin systems

Hardness results for IS:

$Z_G(\lambda)$ hard to approximate on d-regular graphs when
$\lambda > c/d$ [Luby–Vigoda STOC '97];
$\lambda = 1$ and $d > 25$ [Dyer–Frieze–Jerrum FOCS '99]

Phase transition at $\lambda_c(d)$:

- [Mossel–Weitz–Wormald PTRF '09] Local MCMC mixes slowly on random bipartite d-reg. graphs, $\lambda_c(d) < \lambda < \lambda_c(d) + \epsilon(d)$

Hardness results for IS:

$Z_G(\lambda)$ hard to approximate on d-regular graphs when
$\lambda > c/d$ [Luby–Vigoda STOC '97];
$\lambda = 1$ and $d > 25$ [Dyer–Frieze–Jerrum FOCS '99]

Phase transition at $\lambda_c(d)$:

- [Mossel–Weitz–Wormald PTRF '09] Local MCMC mixes slowly on random bipartite d-reg. graphs, $\lambda_c(d) < \lambda < \lambda_c(d) + \epsilon(d)$
- [Sly FOCS '10] $Z_G(\lambda)$ hard to approximate on d-regular graphs for $\lambda_c(d) < \lambda < \lambda_c(d) + \epsilon(d)$

Complexity of AF two-spin systems

Hardness results for IS:

$Z_G(\lambda)$ hard to approximate on d-regular graphs when
$\lambda > c/d$ [Luby–Vigoda STOC '97];
$\lambda = 1$ and $d > 25$ [Dyer–Frieze–Jerrum FOCS '99]

Phase transition at $\lambda_c(d)$:

- [Mossel–Weitz–Wormald PTRF '09] Local MCMC mixes slowly on random bipartite d-reg. graphs, $\lambda_c(d) < \lambda < \lambda_c(d) + \epsilon(d)$
- [Sly FOCS '10] $Z_G(\lambda)$ hard to approximate on d-regular graphs for $\lambda_c(d) < \lambda < \lambda_c(d) + \epsilon(d)$
  — first rigorous indication that
  computational transition for finite d-regular graphs

Complexity of AF two-spin systems

Hardness results for IS:

$Z_G(\lambda)$ hard to approximate on d-regular graphs when
$\lambda > c/d$ [Luby–Vigoda STOC '97];
$\lambda = 1$ and $d > 25$ [Dyer–Frieze–Jerrum FOCS '99]

Phase transition at $\lambda_c(d)$:

- [Mossel–Weitz–Wormald PTRF '09] Local MCMC mixes slowly on random bipartite d-reg. graphs, $\lambda_c(d) < \lambda < \lambda_c(d) + \epsilon(d)$
- [Sly FOCS '10] $Z_G(\lambda)$ hard to approximate on d-regular graphs for $\lambda_c(d) < \lambda < \lambda_c(d) + \epsilon(d)$
  — first rigorous indication that
  computational transition for finite d-regular graphs $\longleftrightarrow$
  statistical physics phase transition for the model on T_d

Hardness results for IS:

$Z_G(\lambda)$ hard to approximate on d-regular graphs when
$\lambda > c/d$ [Luby–Vigoda STOC '97];
$\lambda = 1$ and $d > 25$ [Dyer–Frieze–Jerrum FOCS '99]

Phase transition at $\lambda_c(d)$:

- [Mossel–Weitz–Wormald PTRF '09] Local MCMC mixes slowly on random bipartite d-reg. graphs, $\lambda_c(d) < \lambda < \lambda_c(d) + \epsilon(d)$
- [Sly FOCS '10] $Z_G(\lambda)$ hard to approximate on d-regular graphs for $\lambda_c(d) < \lambda < \lambda_c(d) + \epsilon(d)$
  — first rigorous indication that
  computational transition for finite d-regular graphs $\longleftrightarrow$ statistical physics phase transition for the model on T_d
- Subsequently improved to $\lambda > \lambda_c(d)$ for $d \neq 4, 5$
  [Galanis–Ge–Štefankovič–Vigoda–Yang '11]

Recent: complexity of AF two-spin systems

Recent: complexity of AF two-spin systems

Theorem (Sly, Sun '12).

Recent: complexity of AF two-spin systems

Theorem (Sly, Sun '12).

(a) For $d \geq 3, \lambda > \lambda_c(d)$ the IS partition function $Z_G(\lambda)$ is hard to approx. on the class of d-regular graphs.

Recent: complexity of AF two-spin systems

Theorem (Sly, Sun '12).

(a) For $d \geq 3, \lambda > \lambda_c(d)$ the IS partition function $Z_G(\lambda)$ is hard to approx. on the class of d-regular graphs.

(b) For $d \geq 3, \beta < \beta_c^{\mathrm{af}}(B, d)$, the Ising partition function $Z_G(\beta, B)$ is hard to approx. on the class of d-regular graphs.

Recent: complexity of AF two-spin systems

Theorem (Sly, Sun '12).

(a) For $d \geq 3, \lambda > \lambda_c(d)$ the IS partition function $Z_G(\lambda)$ is hard to approx. on the class of d-regular graphs.

(b) For $d \geq 3, \beta < \beta_c^{\mathrm{af}}(B, d)$, the Ising partition function $Z_G(\beta, B)$ is hard to approx. on the class of d-regular graphs.

Non-trivial two-spin systems on d-regular graphs reduce to IS/Ising

Recent: complexity of AF two-spin systems

Theorem (Sly, Sun '12).

(a) For $d \geq 3, \lambda > \lambda_c(d)$ the IS partition function $Z_G(\lambda)$ is hard to approx. on the class of d-regular graphs.

(b) For $d \geq 3, \beta < \beta_c^{\mathrm{af}}(B, d)$, the Ising partition function $Z_G(\beta, B)$ is hard to approx. on the class of d-regular graphs.

Non-trivial two-spin systems on d-regular graphs reduce to IS/Ising

(a) & [Weitz STOC '06]

complete classification of hard-core complexity except at $\lambda_c(d)$

Theorem (Sly, Sun '12).

(a) For $d \geq 3, \lambda > \lambda_c(d)$ the IS partition function $Z_G(\lambda)$ is hard to approx. on the class of d-regular graphs.

(b) For $d \geq 3, \beta < \beta_c^{\mathrm{af}}(B, d)$, the Ising partition function $Z_G(\beta, B)$ is hard to approx. on the class of d-regular graphs.

Non-trivial two-spin systems on d-regular graphs reduce to IS/Ising

(a) & [Weitz STOC '06]
complete classification of hard-core complexity except at $\lambda_c(d)$

(b) & [Jerrum–Sinclair ALP '90] & [Sinclair–Srivastava–Thurley '11]
complete classification of Ising complexity except at $\beta_c^{\mathrm{af}}(B, d)$

Recent: complexity of AF two-spin systems

Theorem (Sly, Sun '12).

(a) For $d \geq 3, \lambda > \lambda_c(d)$ the IS partition function $Z_G(\lambda)$ is hard to approx. on the class of d-regular graphs.

(b) For $d \geq 3, \beta < \beta_c^{\mathrm{af}}(B, d)$, the Ising partition function $Z_G(\beta, B)$ is hard to approx. on the class of d-regular graphs.

Non-trivial two-spin systems on d-regular graphs reduce to IS/Ising

(a) & [Weitz STOC '06]
complete classification of hard-core complexity except at $\lambda_c(d)$

(b) & [Jerrum–Sinclair ALP '90] & [Sinclair–Srivastava–Thurley '11]
complete classification of Ising complexity except at $\beta_c^{\mathrm{af}}(B, d)$

Interpolation & methods from [Montanari–Mossel–Sly PTRF '12]

Recent: complexity of AF two-spin systems

Theorem (Sly, Sun '12).

(a) For $d \geq 3, \lambda > \lambda_c(d)$ the IS partition function $Z_G(\lambda)$ is hard to approx. on the class of d-regular graphs.

(b) For $d \geq 3, \beta < \beta_c^{\mathrm{af}}(B, d)$, the Ising partition function $Z_G(\beta, B)$ is hard to approx. on the class of d-regular graphs.

Non-trivial two-spin systems on d-regular graphs reduce to IS/Ising

(a) & [Weitz STOC '06]
complete classification of hard-core complexity except at $\lambda_c(d)$

(b) & [Jerrum–Sinclair ALP '90] & [Sinclair–Srivastava–Thurley '11]
complete classification of Ising complexity except at $\beta_c^{\mathrm{af}}(B, d)$

Interpolation & methods from [Montanari–Mossel–Sly PTRF '12]
circumvent difficult second moment calculations

Recent: complexity of AF two-spin systems

Theorem (Sly, Sun '12).

(a) For $d \geq 3, \lambda > \lambda_c(d)$ the IS partition function $Z_G(\lambda)$ is hard to approx. on the class of d-regular graphs.

(b) For $d \geq 3, \beta < \beta_c^{\mathrm{af}}(B,d)$, the Ising partition function $Z_G(\beta, B)$ is hard to approx. on the class of d-regular graphs.

Non-trivial two-spin systems on d-regular graphs reduce to IS/Ising

(a) & [Weitz STOC '06]
complete classification of hard-core complexity except at $\lambda_c(d)$

(b) & [Jerrum–Sinclair ALP '90] & [Sinclair–Srivastava–Thurley '11]
complete classification of Ising complexity except at $\beta_c^{\mathrm{af}}(B,d)$

Interpolation & methods from [Montanari–Mossel–Sly PTRF '12] circumvent difficult second moment calculations

Independently, Galanis–Štefankovič–Vigoda '12 establish (a), and (b) with $B = 0$.

Outline

Outline

Proof ideas: interpolation scheme for factor models

Proof ideas: interpolation scheme for factor models

(Dembo, Montanari, Sun '11)

Proof ideas: interpolation scheme for factor models

(Dembo, Montanari, Sun '11)
Generalized interpolation scheme for abstract factor models

Proof ideas: interpolation scheme for factor models

(Dembo, Montanari, Sun '11)
Generalized interpolation scheme for abstract factor models

> Basic idea: if $\limsup_n \partial_B \phi_n \leq \partial_B \Phi$ ★

Proof ideas: interpolation scheme for factor models

(Dembo, Montanari, Sun '11)
Generalized interpolation scheme for abstract factor models

> Basic idea: if $\limsup_n \partial_B \phi_n \le \partial_B \Phi$ ★
> then $\limsup_n [\phi_n(B_1) - \phi_n(B_0)] \le \Phi(B_1) - \Phi(B_0)$.

Proof ideas: interpolation scheme for factor models

(Dembo, Montanari, Sun '11)
Generalized interpolation scheme for abstract factor models

> Basic idea: if $\limsup_n \partial_B \phi_n \leq \partial_B \Phi$ ★
> then $\limsup_n [\phi_n(B_1) - \phi_n(B_0)] \leq \Phi(B_1) - \Phi(B_0)$.

Recall $\phi_n = n^{-1} \log Z_n$:

Proof ideas: interpolation scheme for factor models

(Dembo, Montanari, Sun '11)
Generalized interpolation scheme for abstract factor models

> Basic idea: if $\limsup_n \partial_B \phi_n \le \partial_B \Phi$ ★
> then $\limsup_n [\phi_n(B_1) - \phi_n(B_0)] \le \Phi(B_1) - \Phi(B_0)$.

Recall $\phi_n = n^{-1} \log Z_n$:
$\Rightarrow \partial_B \phi_n =$ avg. of local observable w.r.t. ν_n

Proof ideas: interpolation scheme for factor models

(Dembo, Montanari, Sun '11)
Generalized interpolation scheme for abstract factor models

> Basic idea: if $\limsup_n \partial_B \phi_n \le \partial_B \Phi$ ★
> then $\limsup_n [\phi_n(B_1) - \phi_n(B_0)] \le \Phi(B_1) - \Phi(B_0)$.

Recall $\phi_n = n^{-1} \log Z_n$:
$\Rightarrow \partial_B \phi_n =$ avg. of local observable w.r.t. ν_n
Can show $\partial_B \Phi(\nu) =$ avg. of same observable at root of T_d
w.r.t. Gibbs measure ν

Proof ideas: interpolation scheme for factor models

(Dembo, Montanari, Sun '11)
Generalized interpolation scheme for abstract factor models

> Basic idea: if $\limsup_n \partial_B \phi_n \le \partial_B \Phi$ ★
> then $\limsup_n [\phi_n(B_1) - \phi_n(B_0)] \le \Phi(B_1) - \Phi(B_0)$.

Recall $\phi_n = n^{-1} \log Z_n$:
$\Rightarrow \partial_B \phi_n =$ avg. of local observable w.r.t. ν_n
Can show $\partial_B \Phi(\nu) =$ avg. of same observable at root of T_d
w.r.t. Gibbs measure ν

If Gibbs measure unique,

Proof ideas: interpolation scheme for factor models

(Dembo, Montanari, Sun '11)
Generalized interpolation scheme for abstract factor models

> Basic idea: if $\limsup_n \partial_B \phi_n \leq \partial_B \Phi$ ★
> then $\limsup_n [\phi_n(B_1) - \phi_n(B_0)] \leq \Phi(B_1) - \Phi(B_0)$.

Recall $\phi_n = n^{-1} \log Z_n$:
$\Rightarrow \partial_B \phi_n =$ avg. of local observable w.r.t. ν_n
Can show $\partial_B \Phi(\nu) =$ avg. of same observable at root of T_d
w.r.t. Gibbs measure ν

If Gibbs measure unique, observable averages on G_n
converge to averages on T_d by general theory

Proof ideas: interpolation scheme for factor models

(Dembo, Montanari, Sun '11)
Generalized interpolation scheme for abstract factor models

> Basic idea: if $\limsup_n \partial_B \phi_n \leq \partial_B \Phi$ ★
> then $\limsup_n [\phi_n(B_1) - \phi_n(B_0)] \leq \Phi(B_1) - \Phi(B_0)$.

Recall $\phi_n = n^{-1} \log Z_n$:
$\Rightarrow \partial_B \phi_n =$ avg. of local observable w.r.t. ν_n
Can show $\partial_B \Phi(\nu) =$ avg. of same observable at root of T_d
w.r.t. Gibbs measure ν

If Gibbs measure unique, observable averages on G_n
converge to averages on T_d by general theory $\Rightarrow$ ★

Proof ideas: interpolation scheme for factor models

(Dembo, Montanari, Sun '11)
Generalized interpolation scheme for abstract factor models

> Basic idea: if $\limsup_n \partial_B \phi_n \leq \partial_B \Phi$ ★
> then $\limsup_n [\phi_n(B_1) - \phi_n(B_0)] \leq \Phi(B_1) - \Phi(B_0)$.

Recall $\phi_n = n^{-1} \log Z_n$:
$\Rightarrow \partial_B \phi_n =$ avg. of local observable w.r.t. ν_n
Can show $\partial_B \Phi(\nu) =$ avg. of same observable at root of T_d
w.r.t. Gibbs measure ν

If Gibbs measure unique, observable averages on G_n
converge to averages on T_d by general theory $\Rightarrow$ ★

Can sometimes obtain ★ beyond uniqueness from
(model-specific) (anti-)monotonicity properties

Proof ideas: BP recursion on T_d

Proof ideas: BP recursion on T_d

BP recursion on general limiting trees is complicated, but

Proof ideas: BP recursion on T_d

BP recursion on general limiting trees is complicated, but
BP recursion on T_d is simply a map $\Delta \to \Delta$:

BP recursion on general limiting trees is complicated, but
BP recursion on T_d is simply a map $\Delta \to \Delta$:

$$h(\sigma) \cong \bar{\psi}(\sigma) \left(\sum_{\sigma'} \psi(\sigma, \sigma') h(\sigma') \right)^{d-1}$$

Proof ideas: BP recursion on T_d

BP recursion on general limiting trees is complicated, but
BP recursion on T_d is simply a map $\Delta \to \Delta$:

$$h(\sigma) \cong \bar{\psi}(\sigma) \left(\sum_{\sigma'} \psi(\sigma, \sigma') h(\sigma') \right)^{d-1}$$

By explicitly analyzing this mapping,
can obtain more exact results for T_d than are implied by interpolation scheme for general trees

Proof ideas: AF two-spin systems on bipartite graphs

For AF two-spin systems on bipartite graphs,
complete Bethe prediction can be verified by interpolation
with a good choice of the local observable

For AF two-spin systems on bipartite graphs,
complete Bethe prediction can be verified by interpolation
with a good choice of the local observable

$\partial_B \phi_n = \mathbb{E}_n[\sigma_{I_n}]$ (with $B \equiv \log \lambda$ for IS)

Proof ideas: AF two-spin systems on bipartite graphs

For AF two-spin systems on bipartite graphs,
complete Bethe prediction can be verified by interpolation with a good choice of the local observable

$\partial_B \phi_n = \mathbb{E}_n[\sigma_{I_n}]$ (with $B \equiv \log \lambda$ for IS)
With obvious observable $i \mapsto \sigma_i$, can show $\phi = \Phi$ for $\lambda \leq \lambda_c$

Proof ideas: AF two-spin systems on bipartite graphs

For AF two-spin systems on bipartite graphs,
complete Bethe prediction can be verified by interpolation with a good choice of the local observable

$\partial_B \phi_n = \mathbb{E}_n[\sigma_{I_n}]$ (with $B \equiv \log \lambda$ for IS)
With obvious observable $i \mapsto \sigma_i$, can show $\phi = \Phi$ for $\lambda \leq \lambda_c$
But by taking observable $i \mapsto (\sigma_i + d^{-1} \sum_{j \in \partial i} \sigma_j)/2$
can show $\phi = \Phi$ for all $\lambda > 0$

IS BP recursion

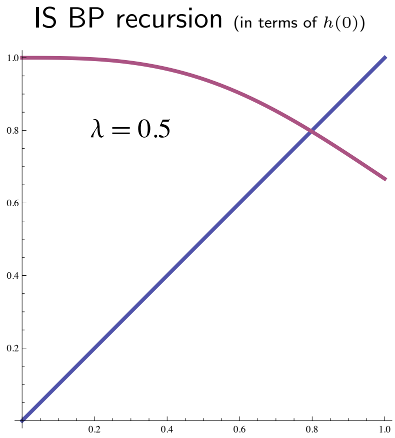

IS BP recursion

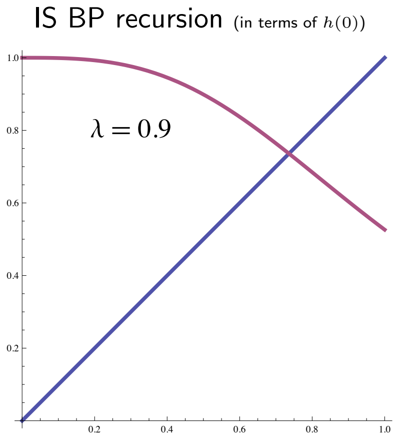

IS BP recursion

IS BP recursion (in terms of $h(0)$)

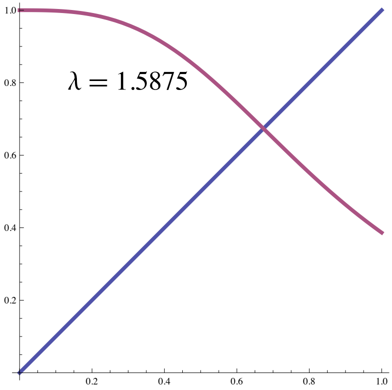

IS BP recursion

IS BP recursion (in terms of $h(0)$)

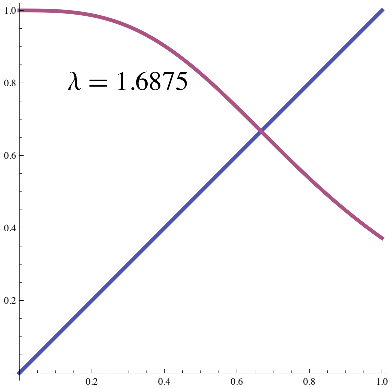

IS BP recursion

IS BP recursion (in terms of $h(0)$)

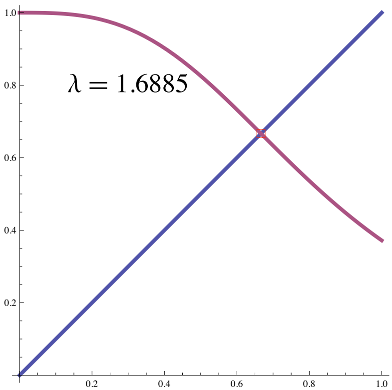

Semi-translation-invariant solutions arise above λ_c

IS BP recursion

IS BP recursion (in terms of $h(0)$)

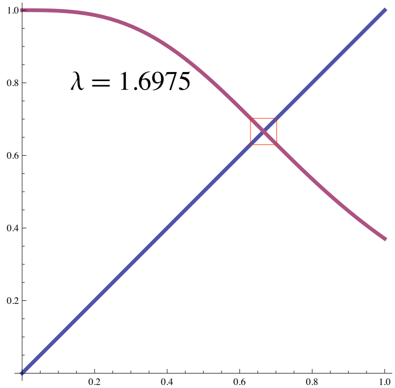

Semi-translation-invariant solutions arise above λ_c

IS BP recursion (in terms of $h(0)$)

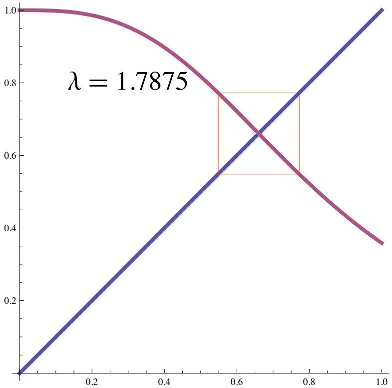

Semi-translation-invariant solutions arise above λ_c

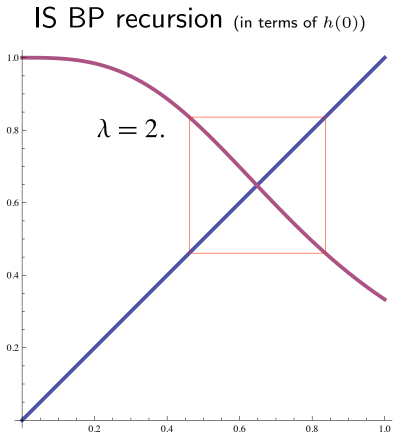

Semi-translation-invariant solutions arise above λ_c

IS BP recursion

IS BP recursion (in terms of $h(0)$)

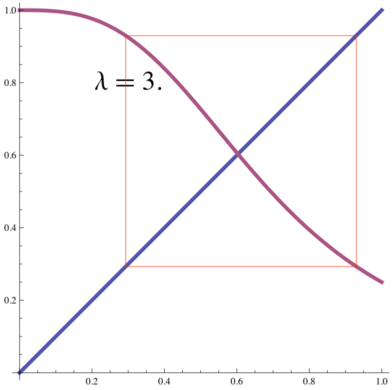

Semi-translation-invariant solutions arise above λ_c

IS BP recursion

IS BP recursion (in terms of $h(0)$)

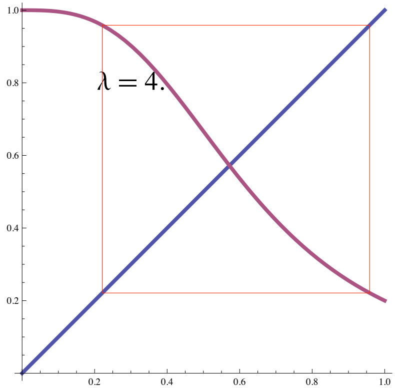

Semi-translation-invariant solutions arise above λ_c

IS BP recursion (in terms of $h(0)$)

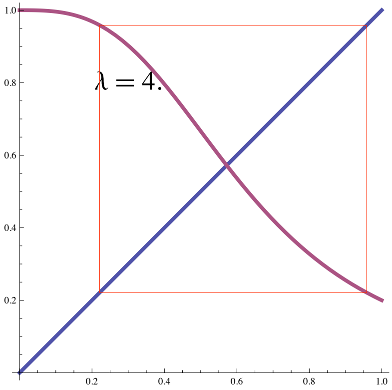

BP solutions as function of λ

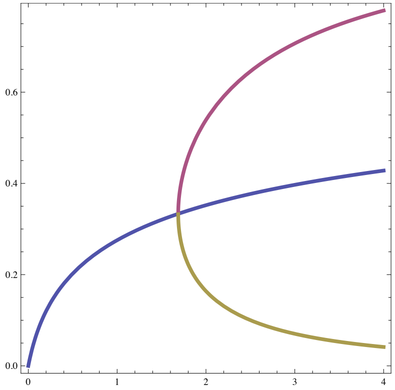

Semi-translation-invariant solutions arise above λ_c

IS free energy density

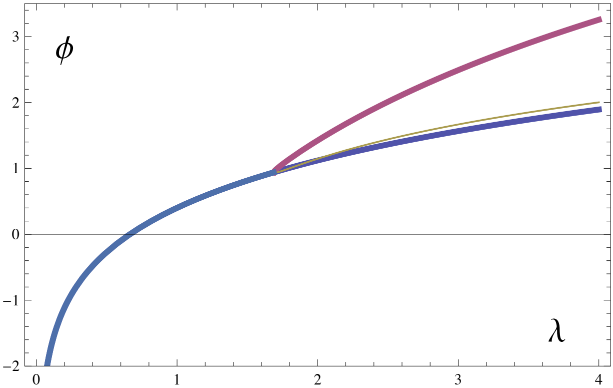

IS free energy density

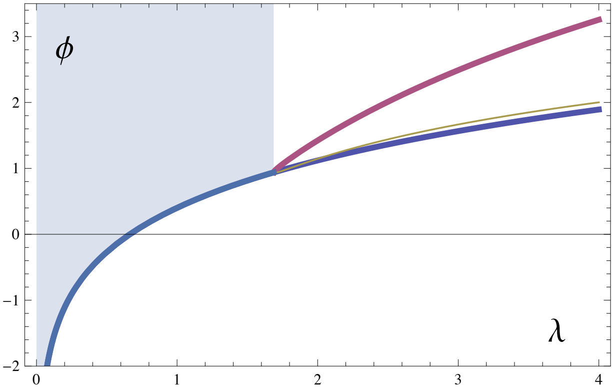

IS free energy density

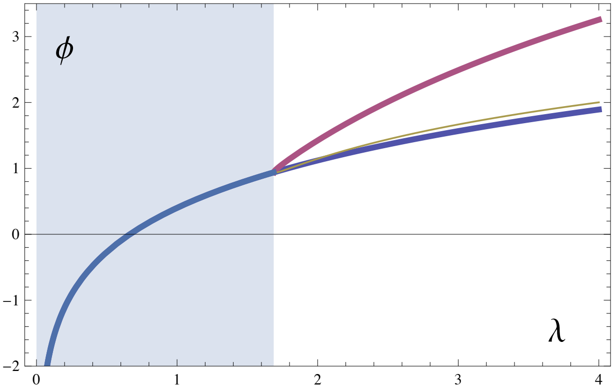

IS free energy density

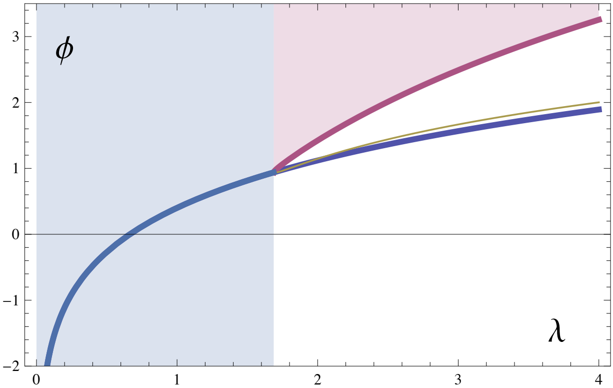

IS free energy density

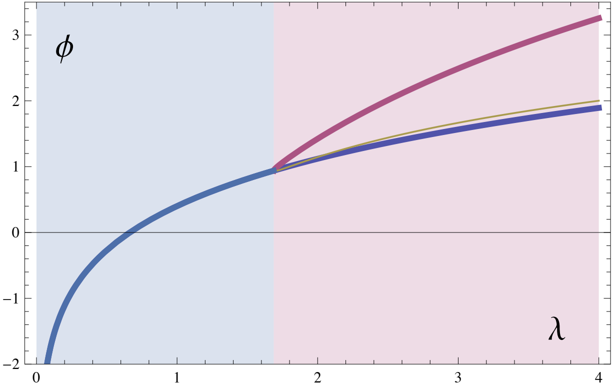

Use bipartite property to interpolate
semi-trans.-inv. fixed point from $\lambda = \infty$

Proof ideas: interpolation for Potts

In Potts model, $\partial_B \phi_n = \mathbb{E}_n \mathbb{E}_{\nu_n}[\delta_{\sigma_{I_n},1}]$,

Proof ideas: interpolation for Potts

In Potts model, $\partial_B \phi_n = \mathbb{E}_n \mathbb{E}_{\nu_n}[\delta_{\sigma_{I_n},1}]$,
so local observable is simply $v \mapsto \delta_{\sigma_v,1}$

Proof ideas: interpolation for Potts

In Potts model, $\partial_B \phi_n = \mathbb{E}_n \mathbb{E}_{\nu_n}[\delta_{\sigma_{I_n},1}]$,
so local observable is simply $v \mapsto \delta_{\sigma_v,1}$

Similarly $\partial_\beta \phi_n = \mathbb{E}_n \mathbb{E}_{\nu_n}[\sum_{j\in\partial I_n} \delta_{\sigma_{I_n},\sigma_j}]$

In Potts model, $\partial_B \phi_n = \mathbb{E}_n \mathbb{E}_{\nu_n}[\delta_{\sigma_{I_n},1}]$,
so local observable is simply $v \mapsto \delta_{\sigma_v,1}$

Similarly $\partial_\beta \phi_n = \mathbb{E}_n \mathbb{E}_{\nu_n}[\sum_{j \in \partial I_n} \delta_{\sigma_{I_n},\sigma_j}]$

In non-uniqueness regimes, can take advantage of

Proof ideas: interpolation for Potts

In Potts model, $\partial_B \phi_n = \mathbb{E}_n \mathbb{E}_{\nu_n}[\delta_{\sigma_{I_n},1}]$,
so local observable is simply $v \mapsto \delta_{\sigma_v,1}$

Similarly $\partial_\beta \phi_n = \mathbb{E}_n \mathbb{E}_{\nu_n}[\sum_{j \in \partial I_n} \delta_{\sigma_{I_n},\sigma_j}]$

In non-uniqueness regimes, can take advantage of **random-cluster** (FK) **representation** for Potts model

Proof ideas: interpolation for Potts

In Potts model, $\partial_B \phi_n = \mathbb{E}_n \mathbb{E}_{\nu_n}[\delta_{\sigma_{I_n},1}]$,
so local observable is simply $v \mapsto \delta_{\sigma_v,1}$

Similarly $\partial_\beta \phi_n = \mathbb{E}_n \mathbb{E}_{\nu_n}[\sum_{j \in \partial I_n} \delta_{\sigma_{I_n},\sigma_j}]$

In non-uniqueness regimes, can take advantage of **random-cluster** (FK) **representation** for Potts model to get monotonicity properties, thereby restricting range of admissible Gibbs measures

Ising vs. Potts

Ising BP (in terms of $\log[h(+)/h(-)]$)

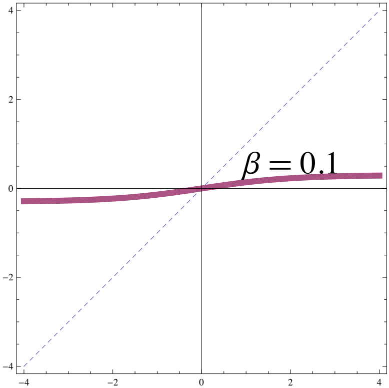

Ising vs. Potts

Ising BP (in terms of $\log[h(+)/h(-)]$)

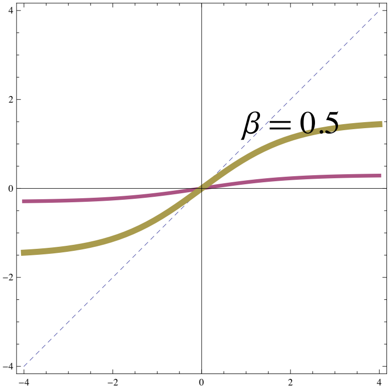

Ising BP (in terms of $\log[h(+)/h(-)]$)

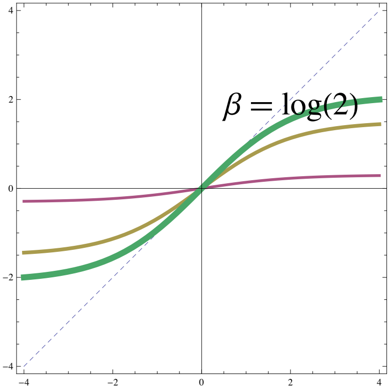

Ising vs. Potts

Ising BP (in terms of $\log[h(+)/h(-)]$)

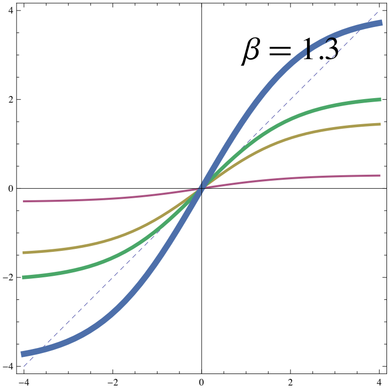

Ising vs. Potts

Ising BP (in terms of $\log[h(+)/h(-)]$)

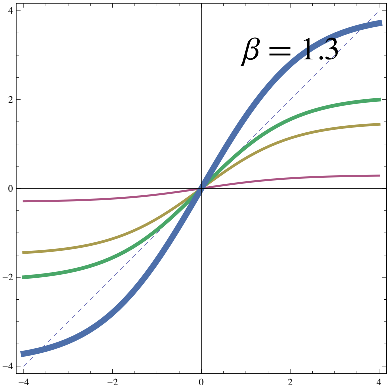

BP solutions as function of β

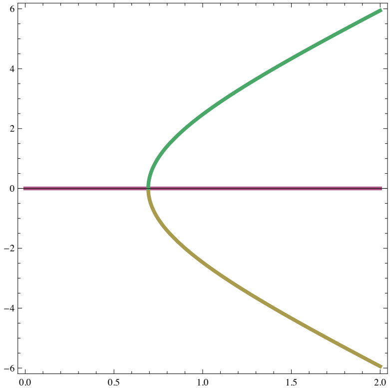

Ising vs. Potts

Ising BP (in terms of $\log[h(+)/h(-)]$)

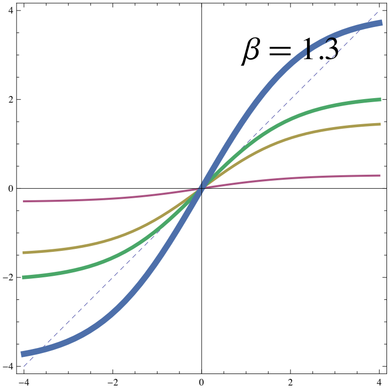

BP solutions as function of β

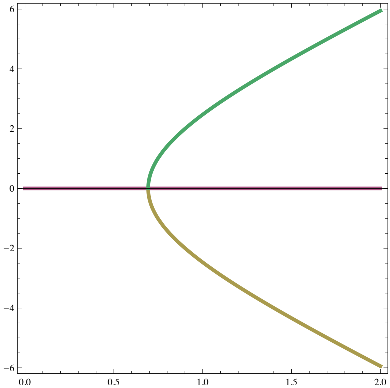

Adding small field $B > 0$ resolves non-uniqueness

Ising vs. Potts

Potts BP (in terms of $\log[h(1)/h(2)]$)

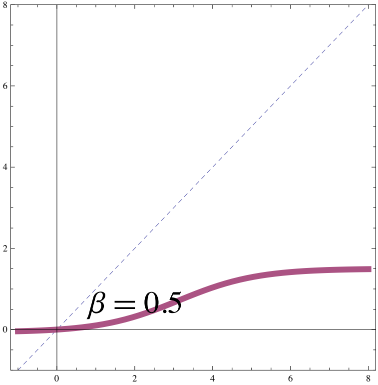

Ising vs. Potts

Potts BP (in terms of $\log[h(1)/h(2)]$)

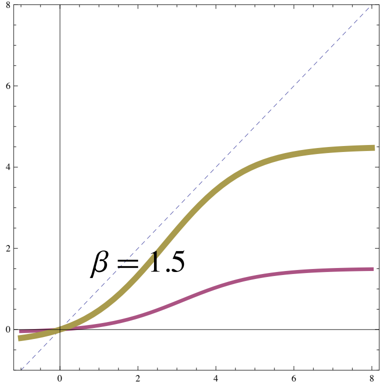

Ising vs. Potts

Potts BP (in terms of $\log[h(1)/h(2)]$)

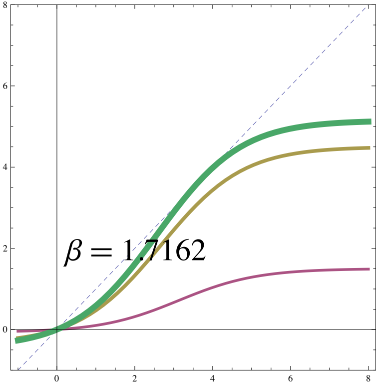

Potts BP (in terms of $\log[h(1)/h(2)]$)

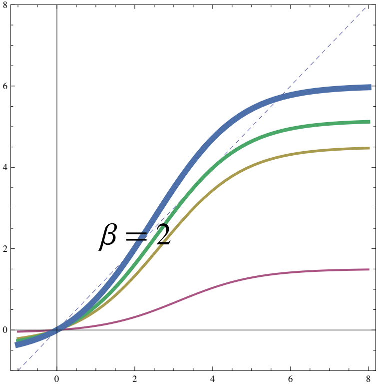

Ising vs. Potts

Potts BP (in terms of $\log[h(1)/h(2)]$)

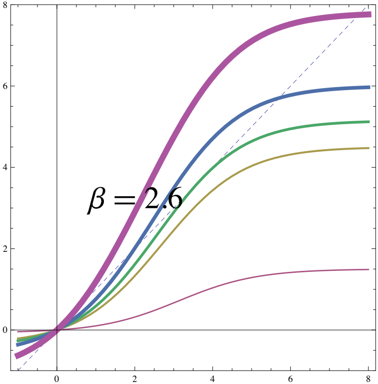

Ising vs. Potts

Potts BP (in terms of $\log[h(1)/h(2)]$)

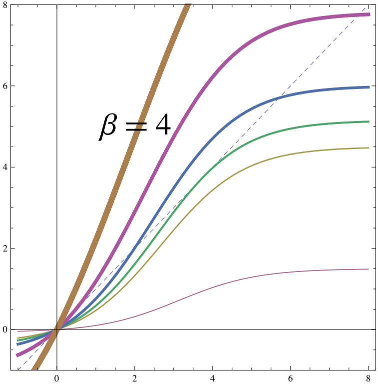

Ising vs. Potts

Potts BP (in terms of $\log[h(1)/h(2)]$)

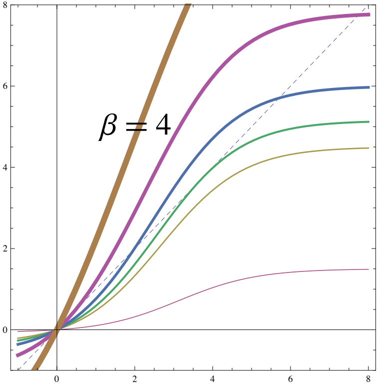

BP solutions as function of β

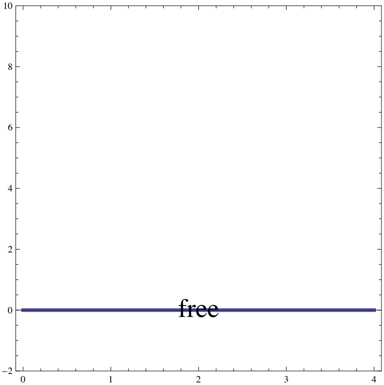

Ising vs. Potts

Potts BP (in terms of $\log[h(1)/h(2)]$)

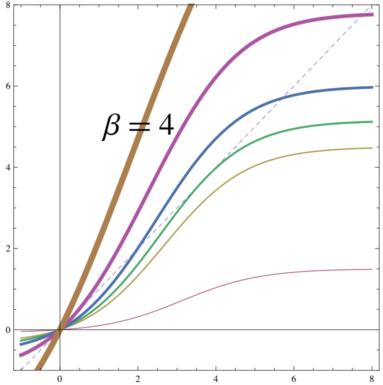

BP solutions as function of β

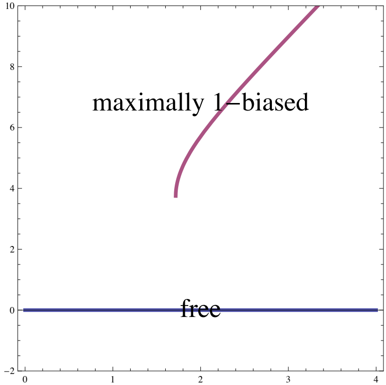

Ising vs. Potts

Potts BP (in terms of $\log[h(1)/h(2)]$)

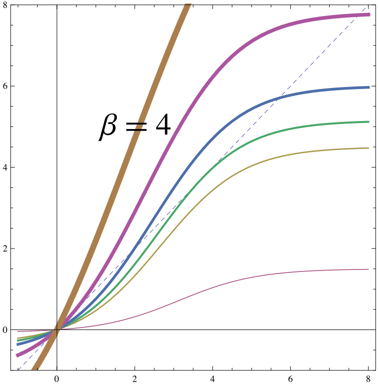

BP solutions as function of β

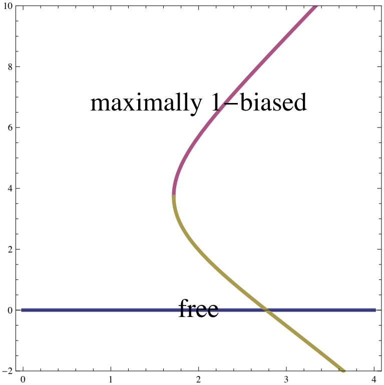

Ising vs. Potts

Potts BP (in terms of $\log[h(1)/h(2)]$)

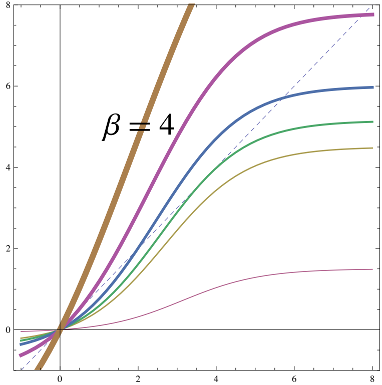

BP solutions as function of β

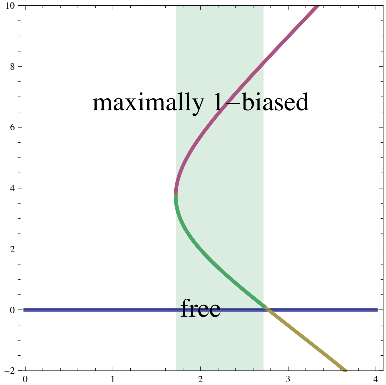

Adding $B > 0$ **not** enough to resolve non-uniqueness

Potts: $\phi \geq \Phi$ by interpolation

Potts: $\phi \geq \Phi$ by interpolation

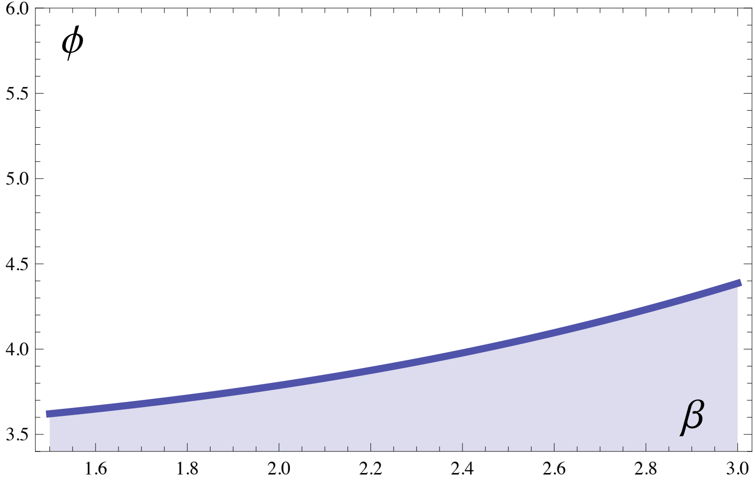

Potts: $\phi \geq \Phi$ by interpolation

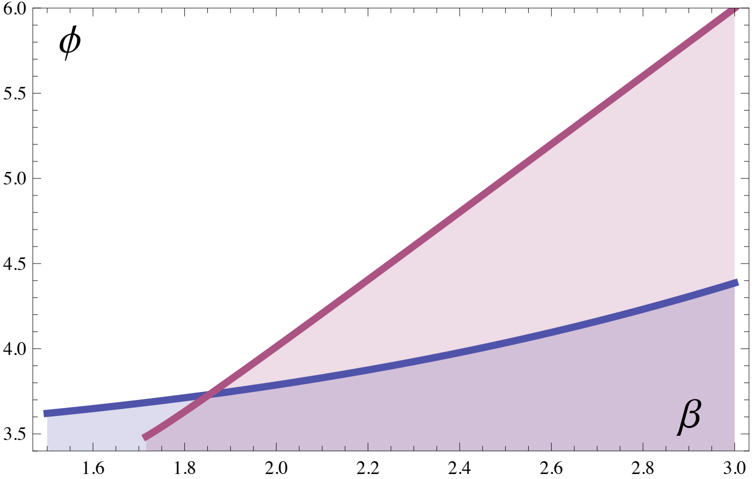

Potts: $\phi \geq \Phi$ by interpolation

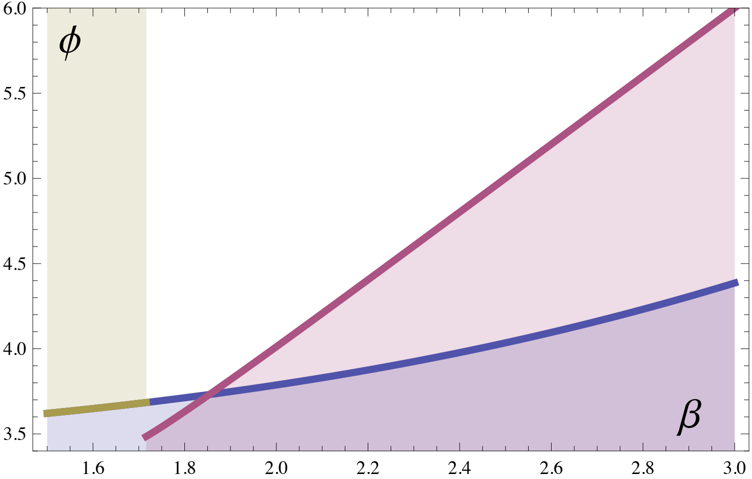

Potts: $\phi \geq \Phi$ by interpolation

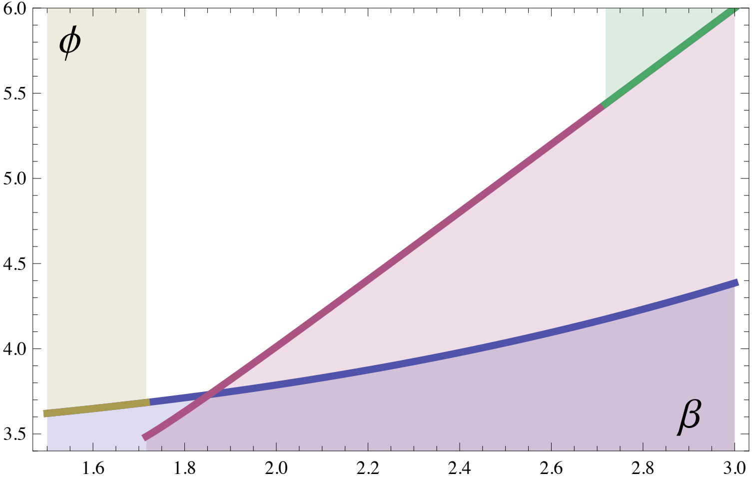

Potts: $\phi \geq \Phi$ by interpolation

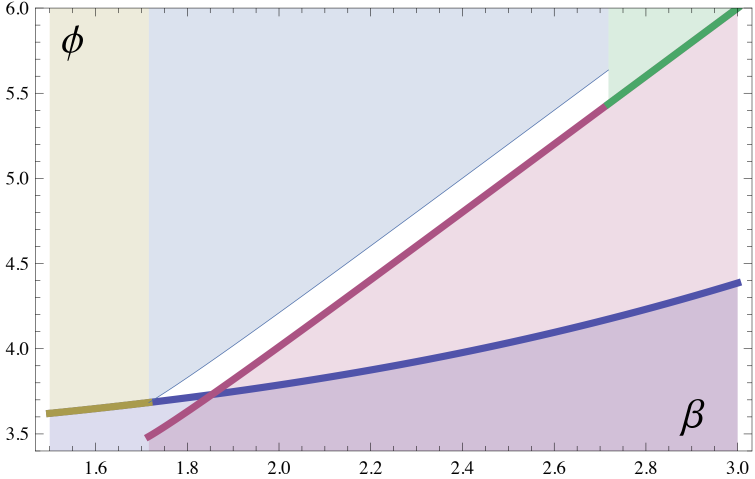

Potts: $\phi \geq \Phi$ by interpolation

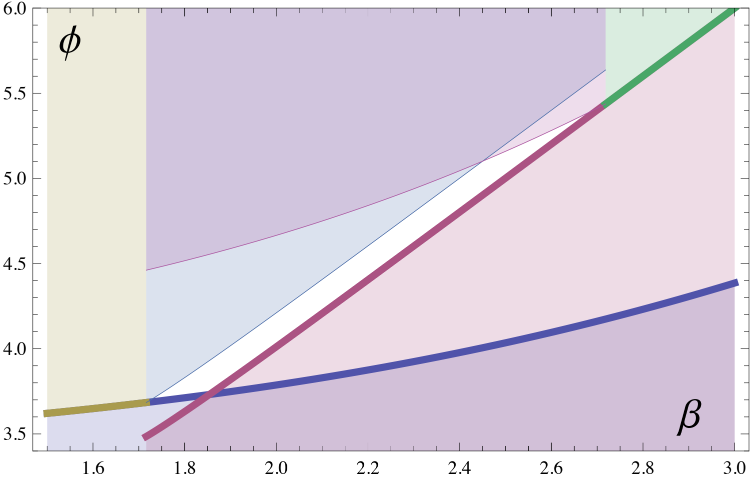

Interpolation gives $\phi \geq \Phi$,

Potts: $\phi \geq \Phi$ by interpolation

Interpolation gives $\phi \geq \Phi$, with equality for $(\beta, B) \notin \mathcal{R}_{\neq}$ (shaded)

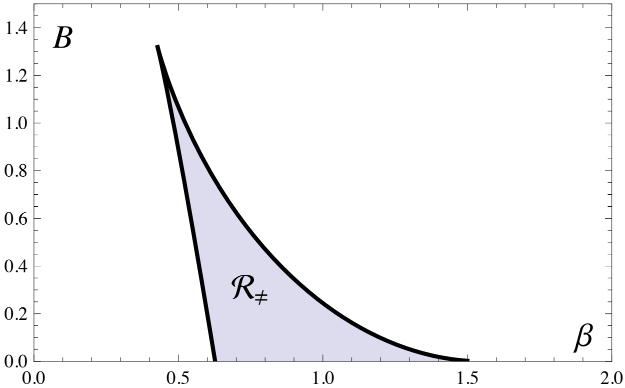

Potts: $\phi \geq \Phi$ by interpolation

Interpolation gives $\phi \geq \Phi$, with equality for $(\beta, B) \notin \mathcal{R}_{\neq}$ (shaded)

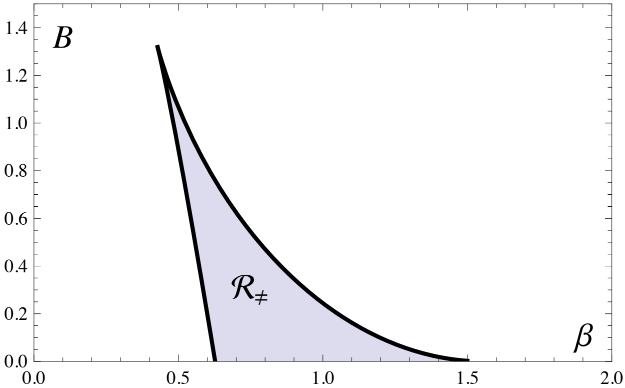

Different approach needed to obtain equality inside $\mathcal{R}_{\neq}$

Potts: $\phi \leq \Phi$ by graph deconstruction

Potts: $\phi \le \Phi$ by graph deconstruction

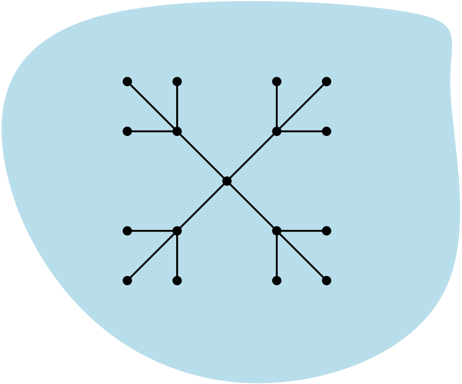

Potts: $\phi \leq \Phi$ by graph deconstruction

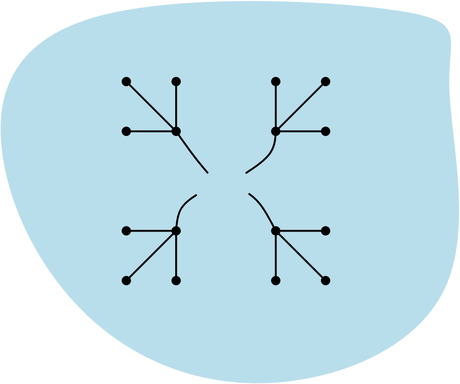

Delete a vertex

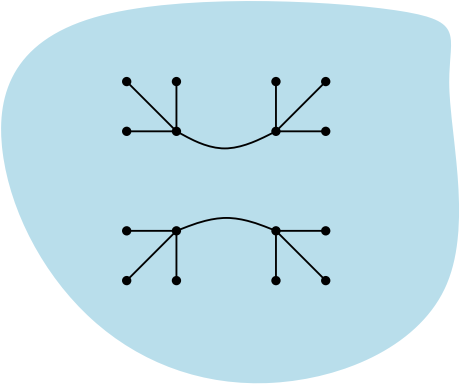

Delete a vertex

Match up half edges

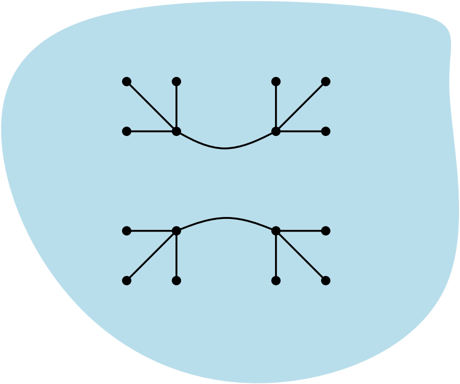

Delete a vertex

Match up half edges

Show decrease in $\log Z$ at each step is $\leq \Phi$ ★

Potts: $\phi \leq \Phi$ by graph deconstruction

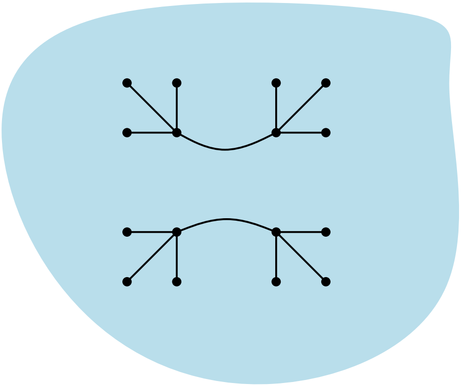

Delete a vertex

Match up half edges

Show decrease in $\log Z$ at each step is $\leq \Phi$ ★

Matching **not** done u.a.r. but to guarantee ★

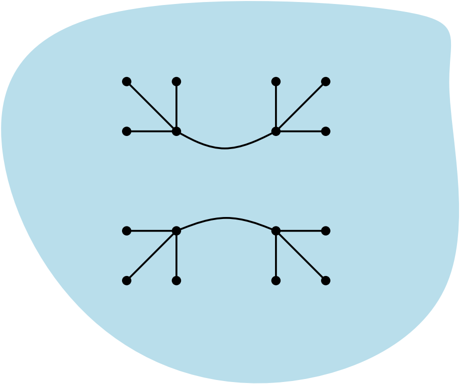

Delete a vertex

Match up half edges

Show decrease in $\log Z$ at each step is $\le \Phi$ ★

Matching **not** done u.a.r. but to guarantee ★

Argue graphs remain uniformly locally tree-like

Potts: $\phi \leq \Phi$ by graph deconstruction

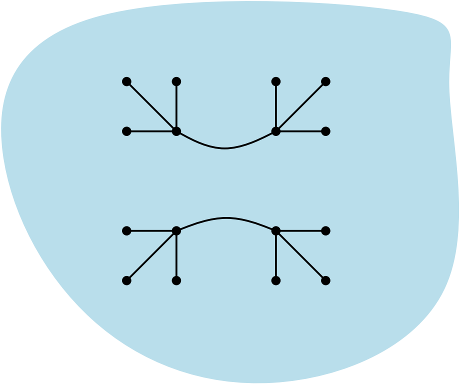

Delete a vertex

Match up half edges

Show decrease in $\log Z$ at each step is $\leq \Phi$ ★

Matching **not** done u.a.r. but to guarantee ★

Argue graphs remain uniformly locally tree-like

This procedure reduces the upper bound to showing ★, which is a difficult (but tractable) calculus problem

Two questions

Two questions

- We make crucial use of the fact that the limiting tree is T_d. Can these methods be extended to more general graph ensembles, e.g. Erdős-Rényi?

Two questions

- We make crucial use of the fact that the limiting tree is T_d. Can these methods be extended to more general graph ensembles, e.g. Erdős-Rényi?

- The Bethe prediction is believed to be false for IS at high fugacity on typical non-bipartite graphs converging to T_d. Can one describe what happens in this case?