Bethe Approximations for Matrix Permanents and Contingency Table Counting

Yuwen Huang

Public Seminar
Data Science & Analytics Thrust, HKUST (Guangzhou)
January 2026

Fast message passing $\Rightarrow$ certified counting

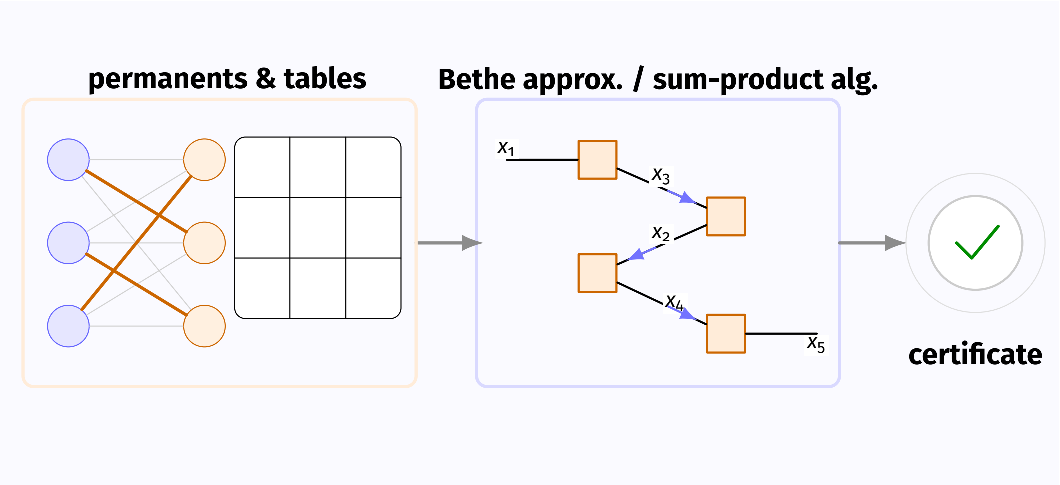

Takeaway: fast inference $\rightarrow$ decision-ready numbers (with guarantees).

One inference engine, many pipelines

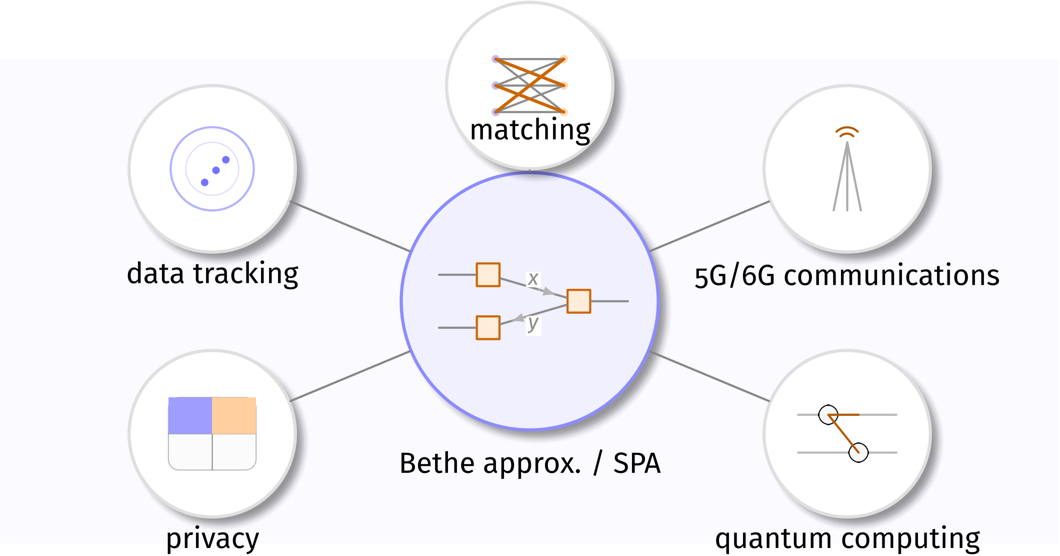

Theme: scalable inference + certificates → trustworthy analytics across domains.

Talk roadmap

Intro and Overview

Exemplar I (matrix permanent)

Exemplar II (contingency tables)

Future directions in data science

Intro and Overview

Introduce **permanent** and our **contribution**.

Introduce **contingency table counting** and our **contribution**.

Introduce the **SPA** and **research thesis**.

Data association (multi-target tracking) $\to$ permanent

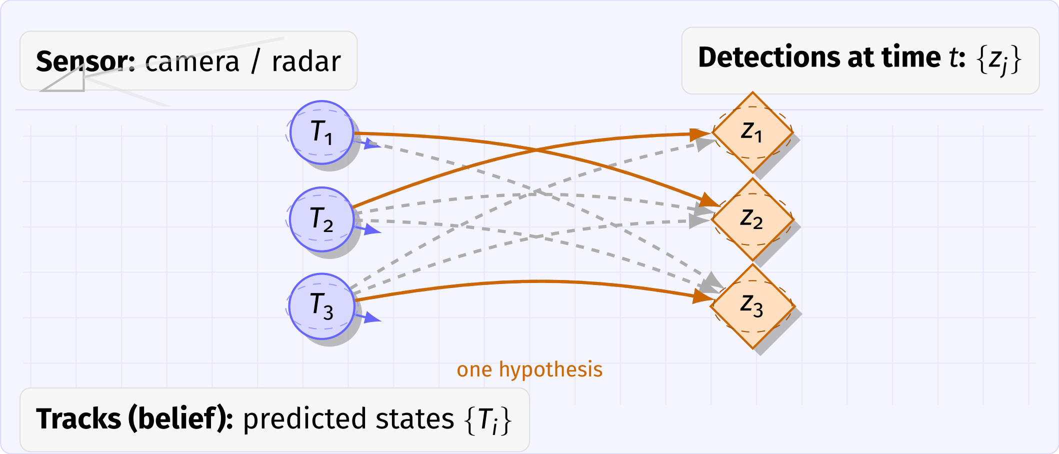

Core question: Which detection came from which track?

Data association (multi-target tracking) $\to$ permanent

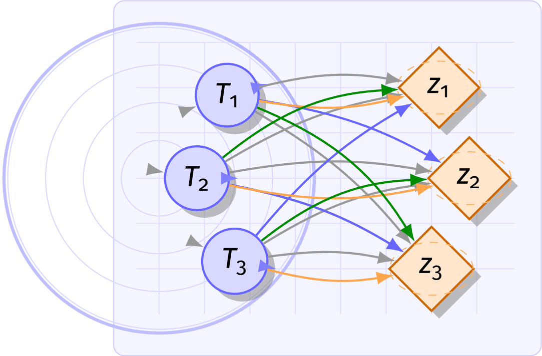

Likelihood matrix characterizing tracks and detections

$$\mathbf{L} \triangleq \begin{pmatrix} \ell(T_1 \to z_1) & \ell(T_1 \to z_2) & \ell(T_1 \to z_3) \\ \ell(T_2 \to z_1) & \ell(T_2 \to z_2) & \ell(T_2 \to z_3) \\ \ell(T_3 \to z_1) & \ell(T_3 \to z_2) & \ell(T_3 \to z_3) \end{pmatrix}$$

Uncertainty of hypothesis (matching) σ

$$p(\sigma) = \frac{\prod_i L(i, \sigma(i))}{\sum_{\sigma \in S_n} \prod_i L(i, \sigma(i))} = \frac{\prod_i L(i, \sigma(i))}{\mathrm{perm}(\mathbf{L})}$$

Joint hypotheses (example matchings): σ_1 σ_2 σ_3 (many more)

Core question: Which detection came from which track?
To get association **uncertainty**: $\mathrm{perm}(\mathbf{L})$.

Topic 1: certified counting for the matrix permanent

Problem

For a **square nonnegative matrix** $\boldsymbol{\theta} \in \mathbb{R}_{\geq 0}^{n \times n}$, compute the **matrix permanent** $\mathrm{perm}(\boldsymbol{\theta})$.

Complexity class: **#P-complete**, even for $\boldsymbol{\theta} \in \{0, 1\}^{n \times n}$.

Approach (graphical-model viewpoint)

1. Build a **standard factor graph** (S-FG) N with partition function

$$Z(\mathsf{N}) \; = \; \mathrm{perm}(\boldsymbol{\theta}).$$

2. Run **sum-product alg. (SPA)** on N to obtain
   the **Bethe approximation** $\mathrm{perm}_{\mathrm{B}}(\boldsymbol{\theta})$ (a fast estimate).

Goal in this talk

Turn a **fast SPA / Bethe estimate** into a **certified output** (with provable bounds).

Topic 1: certified counting for the matrix permanent

Known deterministic bounds for Bethe permanent (more details later)

$$1 \leq \frac{\mathrm{perm}(\boldsymbol{\theta})}{\mathrm{perm}_{\mathrm{B}}(\boldsymbol{\theta})} \leq 2^{n/2}, \qquad \forall \boldsymbol{\theta} \in \mathbb{R}_{\geq 0}^{n \times n}$$

Our contribution: degree-*M* certificate

For any $M \in \mathbb{Z}_{\geq 1}$, it holds that

$$1 \leq \frac{\mathrm{perm}(\boldsymbol{\theta})}{\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})} \leq \left(2^{n/2}\right)^{\frac{M-1}{M}} < 2^{n/2}, \qquad \forall \boldsymbol{\theta} \in \mathbb{R}_{\geq 0}^{n \times n},$$

$\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})$: **degree-*M* Bethe permanent** (defined via **finite graph covers**).

How to read the bound (what to remember)

$$1 \leq \frac{\mathrm{perm}(\boldsymbol{\theta})}{\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})} \leq \left(2^{n/2}\right)^{\frac{M-1}{M}} < 2^{n/2}, \qquad \forall M \in \mathbb{Z}_{\geq 1},\ \forall \boldsymbol{\theta} \in \mathbb{R}_{\geq 0}^{n \times n},$$

Degree-M bound provides a **possible controllable tradeoff:** with $(M-1)/M < 1$,

higher complexity for obtaining $\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta}) \Rightarrow$ **tighter** certificates.

What the bound says

- $\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})$ is a **certified approximation** of $\mathrm{perm}(\boldsymbol{\theta})$.
- $M = 1$: **exact** $\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta}) = \mathrm{perm}(\boldsymbol{\theta})$.
- finite $M > 1$: **improved** certificate vs known bound;
- $M \to \infty$: **recovers** known Bethe bound.

Why it matters

- **Trustworthy:** certificates add bounds guarantees.
- **Application:** data association likelihood / matching problems.

Permanent $\to$ probabilistic matching / assignment (soft one-to-one)

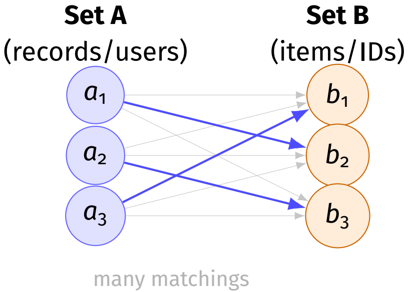

Score matrix $W(i,j) \geq 0$

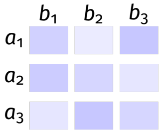

Probability of matching σ

$$\Pr(\sigma) = \frac{\prod_i W(i,\sigma(i))}{\sum_{\sigma'} \prod_i W(i,\sigma'(i))}$$

Takeaway: permanent $=$ **sum** over weighted perfect matchings

Permanent → clustered event modeling (permanental point process)

Event data: locations of incidents

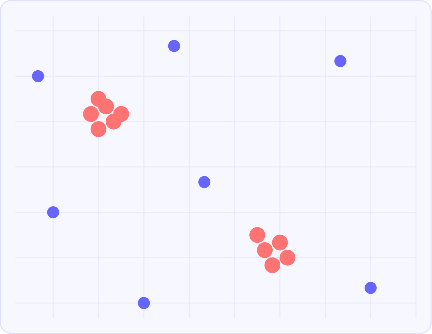

Clustered event modeling

1. event data $\Rightarrow$ similarity kernel $K(x, y)$
2. subset $S \Rightarrow$ kernel submatrix K_S
3. $\Pr(S)$ **increases** with $\mathrm{perm}(K_S)$
$\Rightarrow$ nearby / similar points get
higher weight(hotspots)

Takeaway: permanents are not niche; they show up across modeling tasks.

Intro and Overview

Introduce **permanent** and our **contribution**.

- Introduce **contingency table counting** and our **contribution**.

Introduce the **SPA** and **research thesis**.

Contingency tables → online experimentation (exact tests)

condition on margins

A / B test exposure + outcomes → **2×2 table fixed margins** → **Exact *p*-value** (Fisher's exact test.)

	A	B	Total
Click	42	30	**72**
No click	58	70	**128**
Total	**100**	**100**	

row totals fixed

column totals fixed

Exact test idea
enumerate all 2×2 tables with the **same margins**

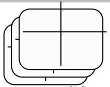
many tables

Key message: fixed margins ⇒ **count compatible tables** ⇒ exact *p*-value
Counting becomes **challenging** as the table size grows.

Overview of Topic 2: binary contingency tables

Setup: Consider **arbitrary** $\boldsymbol{\theta} \in \{0,1\}^{m \times n}$, prescribed **row sums** $r_1, \ldots, r_m \in \mathbb{Z}_{\geq 1}$ and **column sums** $c_1, \ldots, c_n \in \mathbb{Z}_{\geq 1}$.

Problem: **Count** the number of contingency tables $\mathbf{X} \in \{0,1\}^{m \times n}$ such that

1. it satisfies prescribed **row and column sums**:
$$\sum_{j=1}^{n} X(i,j) = r_i \in \mathbb{Z}_{\geq 1}, \quad \forall i = 1, \ldots, m, \qquad \sum_{i=1}^{m} X(i,j) = c_j \in \mathbb{Z}_{\geq 1}, \quad \forall j = 1, \ldots, n.$$
2. it respects the **support constraint** induced by $\boldsymbol{\theta}$:
$$X(i,j) \leq \theta(i,j), \qquad \forall i = 1, \ldots, m, \ \forall j = 1, \ldots, n.$$

The **total number** of such binary contingency table $\mathbf{X}$ is denoted by $\mathrm{BCT}_{\mathbf{r},\mathbf{c}}(\boldsymbol{\theta})$

Generalization of matrix permanent:
$$\mathrm{BCT}_{\mathbf{r},\mathbf{c}}(\boldsymbol{\theta}) = \mathrm{perm}(\boldsymbol{\theta}), \quad \text{if } r_1 = \cdots = r_m = c_1 = \cdots = c_n = 1.$$

Overview of Topic 2: binary contingency tables

Graphical model-based approximation:

1. Construct an S-FG N such that $Z(\mathsf{N}) = \mathrm{BCT}_{\mathbf{r},\mathbf{c}}(\boldsymbol{\theta})$
2. Run the **SPA** on N to get the the value of the **Bethe approximation** $Z_{\mathrm{B}}(\mathsf{N})$.

Our main contributions:

1. Prove that the SPA finds the value of $Z_{\mathrm{B}}(\mathsf{N})$ **exponentially fast**.
2. **Bounds:**

$$1 \leq \frac{Z(\mathsf{N})}{Z_{\mathrm{B}}(\mathsf{N})} \leq \alpha_{\mathsf{Bethe}}, \quad 1 \leq \frac{Z_{\mathrm{Cap}}(\mathsf{N})}{Z(\mathsf{N})} \leq \alpha_{\mathsf{Cap}}, \quad \alpha_{\mathsf{Bethe}} \leq \alpha_{\mathsf{Cap}},$$

where $Z_{\mathrm{Cap}}(\mathsf{N})$ is **the state-of-the-art** poly-time deterministic approximation method proposed in [Adv. Math, 2016] and [Israel J. Math., 2023].

Impact: Validates the **SPA** as a **leading poly-time deterministic approximation method**.

Contingency tables → privacy risk (fixed totals)

privacy

Scenario: publish totals, hide some cells

	Group A	Group B	Total
Category 1	**?**	30	**72**
Category 2	58	70	**128**
Total	**100**	**100**	

row totals given

column totals given

Privacy question
with the totals released,
how much can we guess the hidden cell?

Counting helps
how many tables fit the same totals?
more tables ⇒ safer release

Key message: fixed totals create many possible tables.
Counting (or bounding) how many tells us the **privacy risk** and helps choose a safer release.

Intro and Overview

Introduce **permanent** and our **contribution**.

Introduce **contingency table counting** and our **contribution**.

- Introduce the **SPA** and **research thesis**.

Sum–Product Algorithm (SPA) / Belief Propagation (BP)

What it is

- **Local message passing** on **probabilistic graphical models** $\Rightarrow$ approximating marginal prob. dist.
- **Trees:** **exact** inference in two passes.
- **Loops:** **scalable** heuristic (often accurate in practice).

Dynamic programming in disguise

Kalman filter	state-space
Viterbi / BCJR	hidden Markov chain
Value iteration	Markov decision processes

Industry application: decoding error correction code in **5G/6G** standard

Trustworthy: SPA / Bethe $\Rightarrow$ **certified** bounds (this talk)

Research thesis

Turn **fast estimate** into **decision-ready** analytics

by attaching certificates.

Talk roadmap

Intro and Overview

Exemplar I (Certified counting): Degree-M Bethe bounds for the **permanent**

Exemplar II (Broader class): Bethe approximation for **contingency-table counting**

Future directions: toward a **certified inference toolkit**

Finite-graph-covers-based bounds for the permanent of a non-negative square matrix

- **Setup**

The Bethe approximation method

Finite graph covers

Analyzing the permanent and its degree-M Bethe permanent

Bounding the permanent via its approximations

Conclusion

Permanent vs determinant: same terms, different signs

Example: permanent vs determinant

Key difference: $\det(\boldsymbol{\theta})$ uses **signs** of permutations; $\mathrm{perm}(\boldsymbol{\theta})$ uses **all plus**.

$$\boldsymbol{\theta} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \qquad \boldsymbol{\theta} = \begin{pmatrix} a & b & c \\ d & e & g \\ h & i & j \end{pmatrix}$$

$$\det(\boldsymbol{\theta}) = ad - bc \qquad \det(\boldsymbol{\theta}) = (aej + bgh + cdi) - (agi + bdj + ceh),$$

$$\mathrm{perm}(\boldsymbol{\theta}) = ad + bc \qquad \mathrm{perm}(\boldsymbol{\theta}) = (aej + bgh + cdi) + (agi + bdj + ceh).$$

Takeaway: permanent = “determinant without cancellations.”

Determinant vs permanent: same sum, different signs $\Rightarrow$ huge complexity gap.

Let $\mathcal{S}_{[n]}$ be the set of all $n!$ **permutations** of the set $[n] \triangleq \{1, 2, \ldots, n\}$.

Define $\boldsymbol{\theta} \triangleq \begin{pmatrix} \theta(1,1) & \cdots & \theta(1,n) \\ \vdots & \ddots & \vdots \\ \theta(n,1) & \cdots & \theta(n,n) \end{pmatrix} \in \mathbb{R}_{\geq 0}^{n \times n}$, a **square nonnegative matrix**.

Determinant and **permanent**:

$$\det(\boldsymbol{\theta}) \triangleq \sum_{\sigma \in \mathcal{S}_{[n]}} \mathrm{sgn}(\sigma) \cdot \prod_{i \in [n]} \theta\big(i, \sigma(i)\big),$$ Compl.: **poly-time** $\sim O(n^3)$ (Gaussian elimination)

$$\mathrm{perm}(\boldsymbol{\theta}) \triangleq \sum_{\sigma \in \mathcal{S}_{[n]}} \prod_{i \in [n]} \theta\big(i, \sigma(i)\big),$$ Compl.: **#P-complete**, even for 0–1 matrices.

Why we care: **classical** (Binet/Cauchy, 1812) $\rightarrow$ provably **hard** (Valiant, 1979)

$\rightarrow$ **poly-time randomized approximations** (JSV, 2004)

$\rightarrow$ still inspiring **new bounds and algorithms** today.

Finite-graph-covers-based bounds for the permanent of a non-negative square matrix

Setup

- **The Bethe approximation method**

Finite graph covers

Analyzing the permanent and its degree-M Bethe permanent

Bounding the permanent via its approximations

Conclusion

Motivation:

Among all known deterministic poly-time approximations at the **same** complexity scale, the Bethe approximation offers the **strongest certified bound**.

An S-FG representation of the permanent

The **standard factor graph (S-FG)** N for $\boldsymbol{\theta}$ consists of

1. **edges**: $(1,1), (1,2), (2,1), (2,2)$;
2. **variables** in the matrix
$$\gamma \triangleq \begin{pmatrix} \gamma(1,1) & \gamma(1,2) \\ \gamma(2,1) & \gamma(2,2) \end{pmatrix} \in \{0,1\}^{2\times 2}.$$
3. **nonnegative-valued** local functions $f_{\mathrm{r},1}, f_{\mathrm{r},2}, f_{\mathrm{c},1}, f_{\mathrm{c},2}$;
4. the **partition function**:
$$Z(\boldsymbol{\theta}) = \sum_{\gamma \in \{0,1\}^{2\times 2}} f_{\mathrm{r},1}\big(\gamma(1,:)\big) \cdot f_{\mathrm{r},2}\big(\gamma(2,:)\big) \cdot f_{\mathrm{c},1}\big(\gamma(:,1)\big) \cdot f_{\mathrm{c},2}\big(\gamma(:,2)\big)$$
$$= a \cdot d + b \cdot c$$
$$= \mathrm{perm}(\boldsymbol{\theta}).$$

$$\boldsymbol{\theta} = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

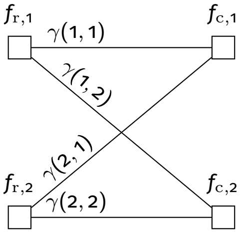

Graphical-model-based approximation method

3. The **Bethe approximation** of the permanent, *i.e.*, the **Bethe partition function**:

$$\mathrm{perm}_{\mathrm{B}}(\boldsymbol{\theta}) \triangleq \exp\Big(- \min_{\gamma \in \Gamma_n} F_{\mathrm{B},\boldsymbol{\theta}}(\gamma)\Big),$$

where $F_{\mathrm{B},\boldsymbol{\theta}}$ is the **Bethe free energy (BFE)** function,

where Γ_n is the set of **doubly stochastic matrices** of size $n \times n$

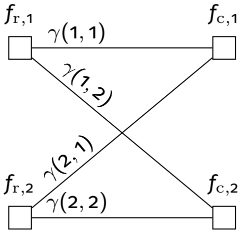

Algorithmic link: run the **SPA** on the S-FG to get $\mathrm{perm}_{\mathrm{B}}(\boldsymbol{\theta})$ **exponentially fast**.

Takeaway: The SPA / Bethe gives a **fast estimate**; our goal is to add **certificates** (bounds).

Graphical-model-based approximation method

We can make **similar definitions** for a more general case:

$$\boldsymbol{\theta} = \begin{pmatrix} \theta(1,1) & \cdots & \theta(1,4) \\ \vdots & \ddots & \vdots \\ \theta(4,1) & \cdots & \theta(4,4) \end{pmatrix} \in \mathbb{R}_{\geq 0}^{4\times 4}.$$

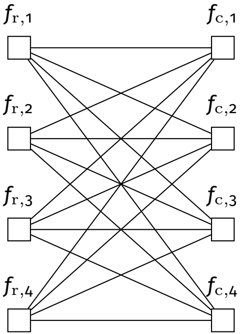

The **S-FG** for $\boldsymbol{\theta}$.

Graphical-model-based approximation method (permanent)

1. **Model:** build a **standard factor graph (S-FG)** N such that

$$Z(\mathsf{N}) = \mathrm{perm}(\boldsymbol{\theta}) = \exp\left(-\min_{\boldsymbol{p}\in\Pi_{\mathcal{A}(\boldsymbol{\theta})}} F_{\mathrm{G},\boldsymbol{\theta}}(\boldsymbol{p})\right),$$

   where $F_{\mathrm{G},\boldsymbol{\theta}}$ is the **Gibbs free energy function**.

2. **Relax + approximate:** replace exact counting by a **Bethe / variational** objective

$$\mathrm{perm}_{\mathrm{B}}(\boldsymbol{\theta}) \triangleq Z_{\mathrm{B}}(\mathsf{N}) = \exp\left(-\min_{\beta\in\Gamma_n} F_{\mathrm{B},\boldsymbol{\theta}}(\beta)\right),$$

   where $F_{\mathrm{B},\boldsymbol{\theta}}$ is the **Bethe free energy function**
   and Γ_n is the set of **doubly-stochastic matrices**.

3. **Compute:** run the **SPA** on N to obtain $\mathrm{perm}_{\mathrm{B}}(\boldsymbol{\theta})$.

Graphical-model-based approximation method

The Bethe approximation offers the **strongest certified bound** we are aware of.

Bounding the **permanent** in terms of the **Bethe permanent**:

$$1 \leq \frac{\mathrm{perm}(\boldsymbol{\theta})}{\mathrm{perm}_{\mathrm{B}}(\boldsymbol{\theta})} \leq 2^{n/2}.$$

- The first inequality was **proven** by Gurvits in [Electron. Colloq. Comput. Complex., 2011] with the help of an **inequality** by Schrijver in [Comb. Theory, Ser. B, 1998].
- The second inequality was **conjectured** by Gurvits and **proven** by Anari and Rezaei in [FOCS 2019].

Vontobel in [IEEE T-IT, Nov. 2013]:

The **sum-product algorithm (SPA)** finds the value of $\mathrm{perm}_{\mathrm{B}}(\boldsymbol{\theta})$ **exponentially fast**.

Exact vs Bethe view (and what comes next)

- **Exact:** $\mathrm{perm}(\boldsymbol{\theta})$ is a partition function
- **Bethe:** $\mathrm{perm}_{\mathrm{B}}(\boldsymbol{\theta})$ **replaces** exact counting by a **variational** objective is **computable by the SPA**.
- **Next step (this work):**
  use **finite graph covers** to derive tighter **degree-M certificates** for $\mathrm{perm}(\boldsymbol{\theta})$.

Finite-graph-covers-based bounds for the permanent of a non-negative square matrix

Setup

The Bethe approximation method

- **Finite graph covers**

Analyzing the permanent and its degree-M Bethe permanent

Bounding the permanent via its approximations

Conclusion

Why graph covers (lifts) show up here

Graph covers (graph lifts) appear across CS/Math:

- **Distributed computing / networks** by Angluin [STOC, 1980]:
  local vs global properties under limited views.
- **Random lifts / expanders** by Linial [Combinatorica, 2002, 2006]:
  constructions and typical properties of lifted graphs.
- **Coding theory** by Koetter and Vontobel in [Adv. Math., 2007]:
  pseudocodewords and behavior of message-passing decoders via covers.

Why we use them today:

Graph-cover viewpoint gives a **combinatorial characterization** of Bethe partition function, which further provides **degree-M** tightening bound.

Finite graph covers: the object we will use

Focus: covers of the S-FG N whose partition function is the permanent.

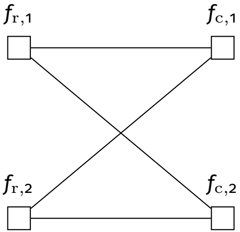

$$\boldsymbol{\theta} = \begin{pmatrix} a & b \\ c & d \end{pmatrix}, \qquad Z(\mathsf{N}) = \mathrm{perm}(\boldsymbol{\theta}) = ad + bc.$$

Definition (informal): an M-cover "replicates" each node M times while preserving local neighborhoods.

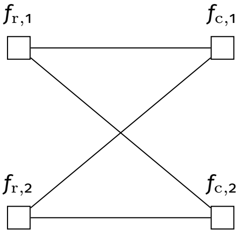

Original graph

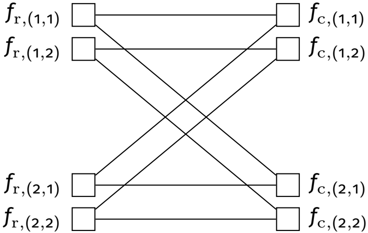

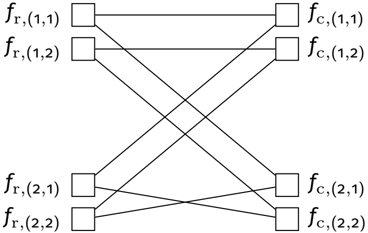

Possible 2-covers

Covers $\longleftrightarrow$ matrix liftings (key equivalence)

Key point: analyzing the **degree-M finite graph covers** $\hat{\mathsf{N}}$ is **equivalent** to analyzing the **P_M-liftings** of $\boldsymbol{\theta}$.

For general $\boldsymbol{\theta} \in \mathbb{R}^{n\times n}$, define a **$P_M$-lifting** of $\boldsymbol{\theta}$:

$$\boldsymbol{\theta}^{\uparrow \boldsymbol{P}_M} \triangleq \begin{pmatrix} \theta(1,1)\cdot \boldsymbol{P}^{(1,1)} & \cdots & \theta(1,n)\cdot \boldsymbol{P}^{(1,n)} \\ \vdots & \ddots & \vdots \\ \theta(n,1)\cdot \boldsymbol{P}^{(n,1)} & \cdots & \theta(n,n)\cdot \boldsymbol{P}^{(n,n)} \end{pmatrix} \in \mathbb{R}_{\geq 0}^{Mn\times Mn},$$

where $\boldsymbol{P}_M \triangleq (\boldsymbol{P}^{(i,j)})_{i,j\in[n]}$ and $\boldsymbol{P}^{(i,j)}$ is a **permutation** matrix of size $M \times M$.

Consequence: for each cover $\hat{\mathsf{N}}$ there exists $\boldsymbol{P}_M$ such that $Z(\hat{\mathsf{N}}) = \mathrm{perm}(\boldsymbol{\theta}^{\uparrow \boldsymbol{P}_M})$.

Degree-M Bethe permanent (Vontobel)

When N corresponds to $\mathrm{perm}(\boldsymbol{\theta})$, define

$$\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta}) \triangleq \sqrt[M]{\frac{1}{|\tilde{\Psi}_M|} \sum_{\boldsymbol{P}_M \in \tilde{\Psi}_M} \mathrm{perm}(\boldsymbol{\theta}^{\uparrow \boldsymbol{P}_M})},$$

where $\tilde{\Psi}_M$ is the set of all $\boldsymbol{P}_M$-liftings of $\boldsymbol{\theta}$.

Interpretation:

$\mathrm{perm}_{\mathrm{B},M}$ **averages** permanents over all degree-M lifts and takes an M-th root.

Graph-cover theorem (Vontobel, [IEEE T-IT 2013])

Takeaway:

The Bethe partition function admits a **combinatorial characterization** via finite graph covers.

$$\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})\Big|_{M\to\infty} = \mathrm{perm}_{\mathrm{B}}(\boldsymbol{\theta})$$

$$\Big|$$

$$\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})$$

$$\Big|$$

$$\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})\Big|_{M=1} = \mathrm{perm}(\boldsymbol{\theta})$$

A **combinatorial characterization** of the Bethe permanent.

Our main contribution for Topic 1

We bound $\mathrm{perm}(\boldsymbol{\theta})$ via $\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})$.

Finite-graph-covers-based bounds for the permanent of a non-negative square matrix

Setup

The Bethe approximation method

Finite graph covers

- **Analyzing the permanent and its degree-M Bethe permanent**

Bounding the permanent via its approximations

Conclusion

Analyzing $\mathrm{perm}(\boldsymbol{\theta})$ vs. $\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})$: discrete types

Setup: $\boldsymbol{\theta} \in \mathbb{R}_{\geq 0}^{n \times n}$, $M \in \mathbb{Z}_{\geq 2}$.

- Γ_n: **doubly-stochastic** $n \times n$ matrices.
- $\Gamma_{M,n} \subset \Gamma_n$: entries are **multiples** of $1/M$ $\big(\gamma(i,j) \in \{0, 1/M, 2/M, \ldots, 1\}\big)$.

Monomial associated with $\gamma \in \Gamma_{M,n}$**:**

$$\boldsymbol{\theta}^{M \cdot \gamma} \triangleq \prod_{i,j \in [n]} \big(\theta(i,j)\big)^{M \cdot \gamma(i,j)}.$$

Key viewpoint: both $\big(\mathrm{perm}(\boldsymbol{\theta})\big)^M$ and $\big(\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})\big)^M$ can be expanded as

$$\sum_{\gamma \in \Gamma_{M,n}} \boldsymbol{\theta}^{M \cdot \gamma} \cdot (\textbf{coefficient} \text{ depending on } \gamma),$$

and we **compare** these coefficients.

Toy case $(n = 2, M = 2)$: same monomials, different coefficients

$$\boldsymbol{\theta} = \begin{pmatrix} a & b \\ c & d \end{pmatrix}, \qquad \mathrm{perm}(\boldsymbol{\theta}) = ad + bc.$$

Three discrete types in $\Gamma_{2,2}$:

$$\gamma^{(1,0)} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \gamma^{(1,1)} = \tfrac{1}{2}\begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}, \quad \gamma^{(0,1)} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

type γ	**monomial** $\theta^{2\gamma}$	**coefficient**
$\gamma^{(1,0)}$	$(ad)^2$	$C_{M,2}(\gamma) = 1,\ C_{\mathrm{B},M,2}(\gamma) = 1$
$\gamma^{(1,1)}$	$abcd$	$C_{M,2}(\gamma) = 2,\ C_{\mathrm{B},M,2}(\gamma) = 1$
$\gamma^{(0,1)}$	$(bc)^2$	$C_{M,2}(\gamma) = 1,\ C_{\mathrm{B},M,2}(\gamma) = 1$

Takeaway: $1 \le \dfrac{C_{M,2}(\gamma)}{C_{\mathrm{B},M,2}(\gamma)} \le 2$ for all types $\Rightarrow 1 \le \dfrac{\left(\mathrm{perm}(\boldsymbol{\theta})\right)^2}{\left(\mathrm{perm}_{\mathrm{B},2}(\boldsymbol{\theta})\right)^2} < 2.$

Toy case $n = 2$ and arbitrary $M \in \mathbb{Z}_{\geq 1}$: same monomials, different coefficients

Generalizing the above result to the case where $n = 2$ and $M \in \mathbb{Z}_{\geq 1}$, the **coefficients** in $\left(\mathrm{perm}(\boldsymbol{\theta})\right)^M$ satisfy

$$C_{M,n}\left(\gamma^{(k,M-k)}\right) = \binom{M}{k}.$$

Note that the **recursion**

$$C_{M+1,n}\left(\gamma^{(k,M+1-k)}\right) = C_{M,n}\left(\gamma^{(k-1,M+1-k)}\right) + C_{M,n}\left(\gamma^{(k,M-k)}\right),$$

is **equivalent** to

$$\binom{M+1}{k} = \binom{M}{k-1} + \binom{M}{k}.$$

Toy case $n = 2$ and arbitrary $M \in \mathbb{Z}_{\geq 1}$: same monomials, different coefficients

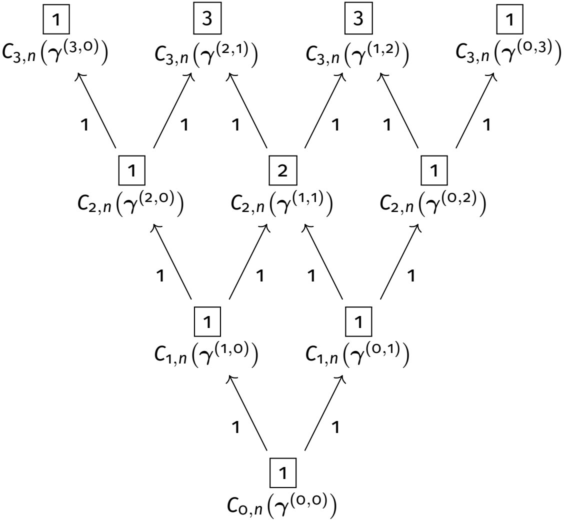

Pascal's triangle visualizing the **recursion** for $C_{M,n}$.

Toy case $n = 2$ and arbitrary $M \in \mathbb{Z}_{\geq 1}$: same monomials, different coefficients

For this **special setup**, the **coefficients** in $\left(\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})\right)^M$ satisfy

$$c_{\mathrm{B},M,n}\left(\gamma^{(k,M-k)}\right) = 1.$$

We have the **recursion**

$$c_{\mathrm{B},M+1,n}\left(\gamma^{(k,M+1-k)}\right) = \begin{cases} c_{\mathrm{B},M,n}\left(\gamma^{(k,M-k)}\right) & k = 0 \\ c_{\mathrm{B},M,n}\left(\gamma^{(k-1,M+1-k)}\right) & k = M+1 \\ \frac{1}{2} \cdot c_{\mathrm{B},M,n}\left(\gamma^{(k-1,M+1-k)}\right) + \frac{1}{2} \cdot c_{\mathrm{B},M,n}\left(\gamma^{(k,M-k)}\right) & 1 \leq k \leq M \end{cases}$$

Toy case $n = 2$ and arbitrary $M \in \mathbb{Z}_{\geq 1}$: same monomials, different coefficients

$$
\begin{tikzcd}
\overset{\boxed{1}}{C_{\mathrm{B},3,n}\left(\gamma^{(3,0)}\right)} & & \overset{\boxed{1}}{C_{\mathrm{B},3,n}\left(\gamma^{(2,1)}\right)} & & \overset{\boxed{1}}{C_{\mathrm{B},3,n}\left(\gamma^{(1,2)}\right)} & & \overset{\boxed{1}}{C_{\mathrm{B},3,n}\left(\gamma^{(0,3)}\right)} \\
& \overset{\boxed{1}}{C_{\mathrm{B},2,n}\left(\gamma^{(2,0)}\right)} \arrow[ul, "1"] \arrow[ur, "\frac{1}{2}"'] & & \overset{\boxed{1}}{C_{\mathrm{B},2,n}\left(\gamma^{(1,1)}\right)} \arrow[ul, "\frac{1}{2}"] \arrow[ur, "\frac{1}{2}"'] & & \overset{\boxed{1}}{C_{\mathrm{B},2,n}\left(\gamma^{(0,2)}\right)} \arrow[ul, "\frac{1}{2}"] \arrow[ur, "1"'] & \\
& & \overset{\boxed{1}}{C_{\mathrm{B},1,n}\left(\gamma^{(1,0)}\right)} \arrow[ul, "1"] \arrow[ur, "\frac{1}{2}"'] & & \overset{\boxed{1}}{C_{\mathrm{B},1,n}\left(\gamma^{(0,1)}\right)} \arrow[ul, "\frac{1}{2}"] \arrow[ur, "1"'] & & \\
& & & \overset{\boxed{1}}{C_{\mathrm{B},0,n}\left(\gamma^{(0,0)}\right)} \arrow[ul, "1"] \arrow[ur, "1"'] & & &
\end{tikzcd}
$$

Generalization of Pascal's triangle visualizing the recursion $C_{\mathrm{B},M,n}$.

Toy case $n = 2$ and arbitrary $M \in \mathbb{Z}_{\geq 1}$: same monomials, different coefficients

$$\begin{tikzcd}[column sep=tiny]
\overset{\boxed{1}}{C_{3,n}(\gamma^{(3,0)})} & & \overset{\boxed{3}}{C_{3,n}(\gamma^{(2,1)})} & & \overset{\boxed{3}}{C_{3,n}(\gamma^{(1,2)})} & & \overset{\boxed{1}}{C_{3,n}(\gamma^{(0,3)})} \\
 & \overset{\boxed{1}}{C_{2,n}(\gamma^{(2,0)})} \arrow[ul, "1"] \arrow[ur, "1"'] & & \overset{\boxed{2}}{C_{2,n}(\gamma^{(1,1)})} \arrow[ul, "1"] \arrow[ur, "1"'] & & \overset{\boxed{1}}{C_{2,n}(\gamma^{(0,2)})} \arrow[ul, "1"] \arrow[ur, "1"'] & \\
 & & \overset{\boxed{1}}{C_{1,n}(\gamma^{(1,0)})} \arrow[ul, "1"] \arrow[ur, "1"'] & & \overset{\boxed{1}}{C_{1,n}(\gamma^{(0,1)})} \arrow[ul, "1"] \arrow[ur, "1"'] & & \\
 & & & \overset{\boxed{1}}{C_{0,n}(\gamma^{(0,0)})} \arrow[ul, "1"] \arrow[ur, "1"'] & & &
\end{tikzcd}$$

$$\begin{tikzcd}[column sep=tiny]
\overset{\boxed{1}}{C_{\mathrm{B},3,n}(\gamma^{(3,0)})} & & \overset{\boxed{1}}{C_{\mathrm{B},3,n}(\gamma^{(2,1)})} & & \overset{\boxed{1}}{C_{\mathrm{B},3,n}(\gamma^{(1,2)})} & & \overset{\boxed{1}}{C_{\mathrm{B},3,n}(\gamma^{(0,3)})} \\
 & \overset{\boxed{1}}{C_{\mathrm{B},2,n}(\gamma^{(2,0)})} \arrow[ul, "1"] \arrow[ur, "\frac{1}{2}"'] & & \overset{\boxed{1}}{C_{\mathrm{B},2,n}(\gamma^{(1,1)})} \arrow[ul, "\frac{1}{2}"] \arrow[ur, "\frac{1}{2}"'] & & \overset{\boxed{1}}{C_{\mathrm{B},2,n}(\gamma^{(0,2)})} \arrow[ul, "\frac{1}{2}"] \arrow[ur, "1"'] & \\
 & & \overset{\boxed{1}}{C_{\mathrm{B},1,n}(\gamma^{(1,0)})} \arrow[ul, "1"] \arrow[ur, "\frac{1}{2}"'] & & \overset{\boxed{1}}{C_{\mathrm{B},1,n}(\gamma^{(0,1)})} \arrow[ul, "\frac{1}{2}"] \arrow[ur, "1"'] & & \\
 & & & \overset{\boxed{1}}{C_{\mathrm{B},0,n}(\gamma^{(0,0)})} \arrow[ul, "1"] \arrow[ur, "1"'] & & &
\end{tikzcd}$$

Visualizing the **recursions** of $C_{M,n}$ and $C_{\mathrm{B},M,n}$.

General case: same monomials, different coefficients

Setup: $n \in \mathbb{Z}_{\geq 1}$, $M \in \mathbb{Z}_{\geq 2}$, and $\Gamma_{M,n} \subseteq \Gamma_n$ (discrete types).

Lemma (type expansion).

There exist **nonnegative** coefficients $\{C_{M,n}(\gamma)\}_{\gamma \in \Gamma_{M,n}}$ and $\{C_{\mathrm{B},M,n}(\gamma)\}_{\gamma \in \Gamma_{M,n}}$ such that

$$\left(\mathrm{perm}(\boldsymbol{\theta})\right)^M = \sum_{\gamma \in \Gamma_{M,n}} \boldsymbol{\theta}^{M \cdot \gamma} \cdot C_{M,n}(\gamma),$$

$$\left(\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})\right)^M = \sum_{\gamma \in \Gamma_{M,n}} \boldsymbol{\theta}^{M \cdot \gamma} \cdot C_{\mathrm{B},M,n}(\gamma).$$

Interpretation:

the **only difference** is the **coefficients** $C_{M,n}(\gamma)$ vs $C_{\mathrm{B},M,n}(\gamma)$.

So a **bound** of $C_{M,n}(\gamma)/C_{\mathrm{B},M,n}(\gamma)$ yields a **bound** of $\mathrm{perm}(\boldsymbol{\theta})/\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})$.

General case: recursion $\Rightarrow$ coefficient comparison

Lemma (recursion, sketch). For $M \geq 2$ and $\gamma \in \Gamma_{M,n}$,

$$C_{M,n}(\gamma) = \sum_{\sigma_1 \in \mathcal{S}_{[n]}(\gamma)} C_{M-1,n}\big(\gamma_{\sigma_1}\big),$$

$$C_{\mathrm{B},M,n}(\gamma) = \frac{1}{\operatorname{perm}(\hat{\gamma}_{\mathcal{R},\mathcal{C}})} \sum_{\sigma_1 \in \mathcal{S}_{[n]}(\gamma)} C_{\mathrm{B},M-1,n}\big(\gamma_{\sigma_1}\big).$$

Main idea of the degree-M bound:

Bounds on $\operatorname{perm}(\hat{\gamma}_{\mathcal{R},\mathcal{C}})$ & recursion $\Rightarrow$ bounds between $C_{M,n}(\gamma)$ and $C_{\mathrm{B},M,n}(\gamma)$.

Finite-graph-covers-based bounds for the permanent of a non-negative square matrix

Setup

The Bethe approximation method

Finite graph covers

Analyzing the permanent and its degree-M Bethe permanent

- **Bounding the permanent via its approximations**

Conclusion

Bounding the permanent via its approximations

Lemma: We **bound** $C_{M,n}$ via $C_{\mathrm{B},M,n}$:

$$1 \leq \frac{C_{M,n}(\gamma)}{C_{\mathrm{B},M,n}(\gamma)} \leq (2^{n/2})^{M-1},$$

where the lower bound **resolves a conjecture** by Vontobel [IEEE T-IT, Nov. 2013].

Theorem: Based on

$$\big(\mathrm{perm}(\boldsymbol{\theta})\big)^{M} = \sum_{\gamma \in \Gamma_{M,n}} \boldsymbol{\theta}^{M \cdot \gamma} \cdot C_{M,n}(\gamma), \quad \big(\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})\big)^{M} = \sum_{\gamma \in \Gamma_{M,n}} \boldsymbol{\theta}^{M \cdot \gamma} \cdot C_{\mathrm{B},M,n}(\gamma),$$

we **bound** the permanent $\mathrm{perm}(\boldsymbol{\theta})$ via its **degree-M Bethe permanent**:

$$1 \leq \frac{\mathrm{perm}(\boldsymbol{\theta})}{\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})} < \left(2^{n/2}\right)^{\frac{M-1}{M}},$$

where the lower bound **resolves another conjecture** by Vontobel [IEEE T-IT, Nov. 2013].

Bounding the permanent via its approximations

As $M \to \infty$,

$$1 \leq \liminf_{M\to\infty} \frac{\mathrm{perm}(\boldsymbol{\theta})}{\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})} \leq \lim_{M\to\infty} \left(2^{n/2}\right)^{\frac{M-1}{M}},$$

we **recover the bounds**

$$1 \leq \frac{\mathrm{perm}(\boldsymbol{\theta})}{\mathrm{perm}_{\mathrm{B}}(\boldsymbol{\theta})} \leq 2^{n/2}$$

where

- the **lower bound** proven by Gurvits in [Electron. Colloq. Comput. Complex., 2011].
- the **upper bound** proven by Anari and Rezaei in [FOCS, SIAM J. Comput., 2019].

Big picture: a tightening knob

Fact (tightening by M). For any $\boldsymbol{\theta} \in \mathbb{R}_{\geq 0}^{n \times n}$ and $M > 1$,

$$\underbrace{1 \leq \frac{\mathrm{perm}(\boldsymbol{\theta})}{\mathrm{perm}_{\mathrm{B}}(\boldsymbol{\theta})} \leq 2^{n/2}}_{\text{Bethe bound}} \qquad \text{and} \qquad \underbrace{1 \leq \frac{\mathrm{perm}(\boldsymbol{\theta})}{\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})} \leq \left(2^{n/2}\right)^{\frac{M-1}{M}}}_{\text{degree-}M\text{ bound}}.$$

Because $\frac{M-1}{M} < 1$, we have $\left(2^{n/2}\right)^{\frac{M-1}{M}} < 2^{n/2}$: the degree-$M$ certificate is **strictly tighter** than the Bethe bound.

Interpretation. M is a **knob**: large $M \Rightarrow$ **tighter** certificates

but computing $\mathrm{perm}_{\mathrm{B},M}$ might be **more expensive / less explicit**.

Research direction Can we compute $\mathrm{perm}_{\mathrm{B},M}(\boldsymbol{\theta})$ in $\mathrm{poly}(n)$ time for **moderately large** M so we obtain **systematically tighter** polynomial-time certificates?

Finite-graph-covers-based bounds for the permanent of a non-negative square matrix

- Setup
- The Bethe approximation method
- Finite graph covers
- Analyzing the permanent and its degree-M Bethe permanent
- Bounding the permanent via its approximations
- **Conclusion**

Conclusion (Topic 1)

- **Certified bounds for permanents:**
  degree-M Bethe permanents yield **certificates** that **tighten** the classical Bethe bound.

- **Conjecture resolved:**
  we prove a conjectured degree-M Bethe-permanent lower bound posed by Vontobel [IEEE T-IT, 2013].

Takeaway: fast SPA / Bethe-style inference can be made **trustworthy** via **certificates**.

The Bethe Partition Function and the SPA for Factor Graphs based on Homogeneous Real Stable Polynomials

An introductory example

Problem formulation and related work

Main results

Numerical results

Intro example: binary contingency tables (3×3, margins $= 2$)

Question. How many binary 3×3 matrices have every row sum and column sum equal to 2?

Examples (✓ **satisfies margins,** $\times$ **does not):**

$$\underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{pmatrix}}_{\times} \qquad \underbrace{\begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \end{pmatrix}}_{\checkmark} \qquad \underbrace{\begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 1 & 0 \end{pmatrix}}_{\checkmark}$$

Key observation.

Row/column sum $= 2$ means each row/column has exactly one zero

$\Rightarrow$ the **zeros form a permutation matrix** $\Rightarrow$ number of such tables is $3! = 6$.

Problem: binary contingency table counting with support constraints

Setup: Consider **arbitrary** $\boldsymbol{\theta} \in \{0,1\}^{m\times n}$, prescribed **row sums** $r_1, \dots, r_m \in \mathbb{Z}_{\geq 1}$ and **column sums** $c_1, \dots, c_n \in \mathbb{Z}_{\geq 1}$.

Problem: **Count** the number of contingency tables $\mathbf{X} \in \{0,1\}^{m\times n}$ such that

1. it satisfies prescribed **row and column sums**:
$$\sum_{j=1}^{n} X(i,j) = r_i \in \mathbb{Z}_{\geq 1}, \quad \forall i = 1, \dots, m, \qquad \sum_{i=1}^{m} X(i,j) = c_j \in \mathbb{Z}_{\geq 1}, \quad \forall j = 1, \dots, n.$$
2. it respects the **support constraint** induced by $\boldsymbol{\theta}$:
$$X(i,j) \leq \theta(i,j), \qquad \forall i = 1, \dots, m, \ \forall j = 1, \dots, n.$$

The **total number** of such binary contingency table $\mathbf{X}$ is denoted by $\mathrm{BCT}_{\mathbf{r},\mathbf{c}}(\boldsymbol{\theta})$

Generalization of matrix permanent:

$$\mathrm{BCT}_{\mathbf{r},\mathbf{c}}(\boldsymbol{\theta}) = \mathrm{perm}(\boldsymbol{\theta}), \quad \text{if } r_1 = \cdots = r_m = c_1 = \cdots = c_n = 1.$$

Graphical-model-based approximation (high level)

Complexity class: #-P complete.

Step 1 (model). Build a **standard factor graph** (S-FG) N with

$$Z(\mathsf{N}) = \mathrm{BCT}_{\mathbf{r},\mathbf{c}}(\boldsymbol{\theta}).$$

Step 2 (compute). Run the **SPA** on N to obtain the **Bethe estimate**

$$Z_{\mathrm{B}}(\mathsf{N}) = \exp\Big(-\min_{\beta \in \mathcal{L}} F_{\mathrm{B}}(\beta)\Big).$$

where $F_{\mathrm{B}}(\beta)$ is the Bethe free energy function.

Where we stand: deterministic poly-time approaches

State-of-the-art **deterministic polynomial-time** approximation methods.

Group / Author	**Method family**	**Result / Limitation**
A. Barvinok [Adv. Math 2016]	Convex / entropy programs	Can be **heavy** for large scale
L. Gurvits [Inform. and Comput. 2015] **Brändén *et al.*** [Israel J. Math. 2023]	Capacity bounds (stable/Lorentzian polynomials)	Can be **loose**. $1 \leq Z_{\mathrm{Cap}}/\mathrm{BCT}_{\mathbf{r},\mathbf{c}} \leq \alpha_{\mathsf{Cap}}$
D. Straszak and **N. K. Vishnoi** [IEEE T-IT, 2019]	Bethe approximation	$\mathrm{BCT}_{\mathbf{r},\mathbf{c}}/Z_{\mathrm{B}} \geq 1$

The Bethe Partition Function and the SPA for Factor Graphs based on Homogeneous Real Stable Polynomials

An introductory example

Problem formulation and related work

- **Main results**

Numerical results

Main results (high level)

Our main contribution (three pieces):

1. **Convex view:** under homogeneous real-stable polynomials, the Bethe partition function $Z_{\mathrm{B}}(\mathsf{N})$ admits a **convex reformulation**.

2. **Fast computation:** the SPA converges to the Bethe value $Z_{\mathrm{B}}(\mathsf{N})$ **exponentially fast**.

3. **Certificates (deterministic bounds):**

$$1 \leq \frac{Z(\mathsf{N})}{Z_{\mathrm{B}}(\mathsf{N})} \leq \alpha_{\mathsf{Bethe}}, \qquad 1 \leq \frac{Z_{\mathrm{Cap}}(\mathsf{N})}{Z(\mathsf{N})} \leq \alpha_{\mathsf{Cap}}, \qquad \alpha_{\mathsf{Bethe}} \leq \alpha_{\mathsf{Cap}}.$$

(Here Z_{Cap} is the **capacity-based approx.** and α_{Cap} is the **state-of-the-art bound** among all deterministic approx. methods by [Inform. and Comput. 2015].)

Takeaway: SPA / Bethe is a **leading deterministic** approach here because it is **fast** and comes with **tighter certificates**.

The Bethe Partition Function and the SPA for Factor Graphs based on Homogeneous Real Stable Polynomials

An introductory example

Problem formulation and related work

Main results

- **Numerical results**

Numerical results

Setup

- $n = m = 6$, $r_i = c_j = 2$;
- Generate 3000 random instances.
- Compare $Z(\mathsf{N})$ vs $Z_{\mathrm{B}}(\mathsf{N})$.

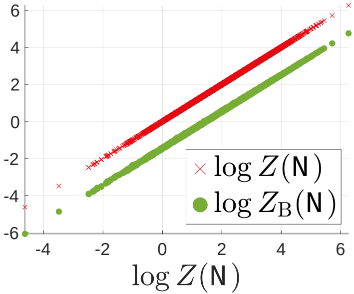

Observation

- $Z_{\mathrm{B}}(\mathsf{N}) \leq Z(\mathsf{N})$ proven by D. Straszak and N. K. Vishnoi in [IEEE T-IT, 2019].
- $Z_{\mathrm{B}}(\mathsf{N})$ provides a **good estimate** of $Z(\mathsf{N})$ in this case.

Numerical results

Setup

Consider **the same setup** as the previous case, but with $n = m = 6$ **replaced** by $n = m = 7$.

Observation

The distribution of $\log Z - \log Z_{\mathrm{B}}$ remains **concentrated**.

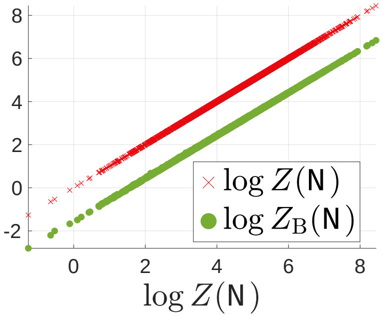

Conclusions of the whole presentation

Conclusions

Takeaway from this talk:

estimate with the SPA / Bethe $+$ **certificate** $\Rightarrow$ decision-ready outputs.

Two instances in this talk:

- **Permanents (matching):** degree-M Bethe $\Rightarrow$ a **tightening** certificate.
- **Contingency tables (A / B tests):**
  The SPA / Bethe $\Rightarrow$ provable convergence + deterministic bounds.

Why this matters for DSA:

a unifying toolkit that connects **statistics / ML / optimization** to **reliable** analytics.

1. **Statistics & machine learning:** probabilistic graphical models & factor graph.
2. **Inference & optimization:** Bethe / variational inference via the SPA.

Future directions: generalize → apply → deliver

1) AI4Science

- **extend** to **physics**-inspired factor-graph models
- produce **practical error bars** (tight bounds + convergence)

2) ML pipelines

- provide **certified** data association primitives in **ML pipelines**.
- **exact-test** workflows: provide reliable p-values in **A / B test**

3) Engineering

- **scalable** SPA **implementations** with clean APIs
- reproducible **benchmarks** + reference **baselines**

Near-term research goal

Open-source certified inference toolkit & **benchmark suite** (accuracy / runtime).

List of presented publications

1. **Y. Huang** and P. O. Vontobel,
   Bounding the permanent of a non-negative matrix via its degree-M Bethe and Sinkhorn permanents
   in Proc. IEEE Int. Symp. Inf. Theory (ISIT), Jun. 2023.

2. **Y. Huang**, N. Kashyap, and P. O. Vontobel,
   Degree-M Bethe and Sinkhorn permanent based bounds on the permanent of a non-negative matrix
   in IEEE Trans. Inf. Theory (T-IT), Jul. 2024.

List of presented publications

3. **Y. Huang** and P. O. Vontobel,
   The Bethe Partition Function and the SPA for Factor Graphs based on Homogeneous Real Stable Polynomials
   in Proc. IEEE Int. Symp. Inf. Theory (ISIT), Jul. 2024.

4. **Y. Huang** and P. O. Vontobel,
   Bethe Approximation of Partition Function based on Homogeneous Real Stable Polynomials
   submitted to IEEE Trans. Inf. Theory (T-IT).

Thank you!