Dynamic Sampling *from* Graphical Models

Yitong Yin
Nanjing University

Joint work with Weiming Feng (Nanjing)
Nisheeth Vishnoi (EPFL)

Graphical Model

instance of graphical model: $\mathcal{I} = (V, E, [q], \Phi)$

- V : variables
- $E \subset 2^V$: constraints
- $[q] = \{0,1, \ldots, q\text{-}1\}$: domain
- $\Phi = (\phi_v)_{v \in V} \cup (\phi_e)_{e \in E}$: factors

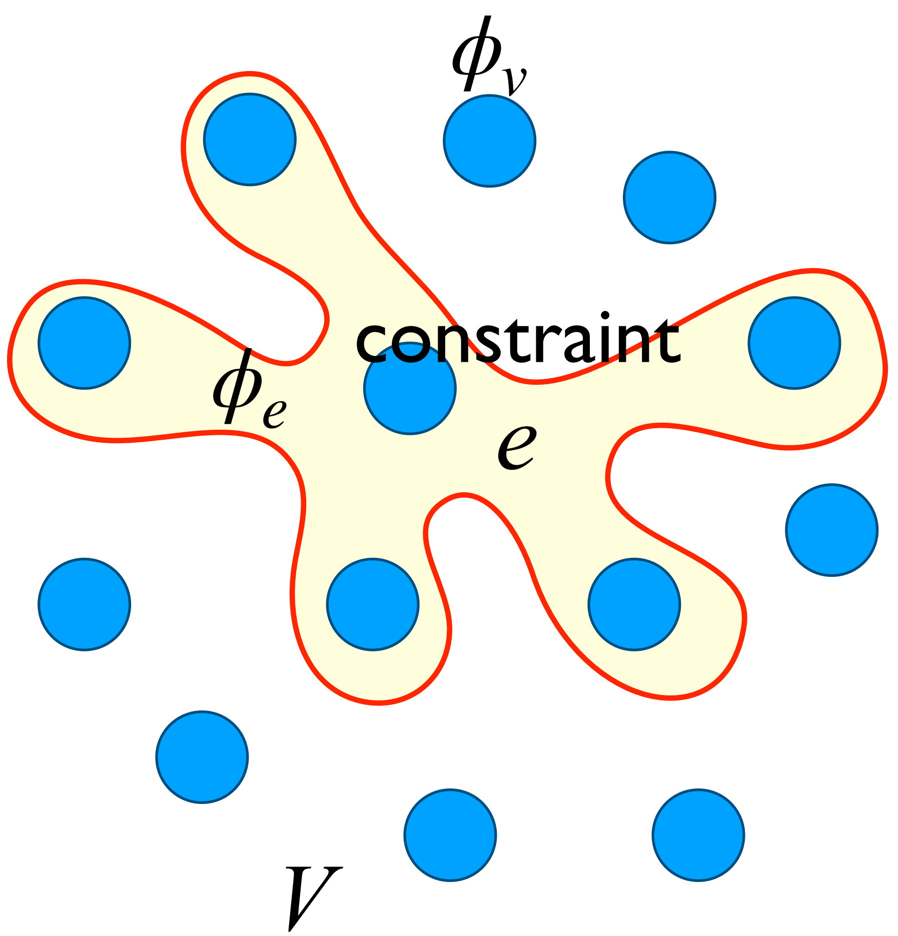

- Gibbs distribution μ over all $\sigma \in [q]^V$:

$$\mu(\sigma) \propto \prod_{v \in V} \phi_v(\sigma_v) \prod_{e \in E} \phi_e(\sigma_e)$$

Graphical Model

instance of graphical model: $\mathcal{I} = (V, E, [q], \Phi)$

- Each $v \in V$ is a *variable* with domain $[q]$ and has a distribution ϕ_v over $[q]$

$$\phi_v : [q] \to [0,1]$$

- Each $e \in E$ is a set of variables and corresponds to a constraint (factor)

$$\phi_e : [q]^e \to [0,1]$$

- Gibbs distribution μ over all $\sigma \in [q]^V$:

$$\mu(\sigma) \propto \prod_{v \in V} \phi_v(\sigma_v) \prod_{e \in E} \phi_e(\sigma_e)$$

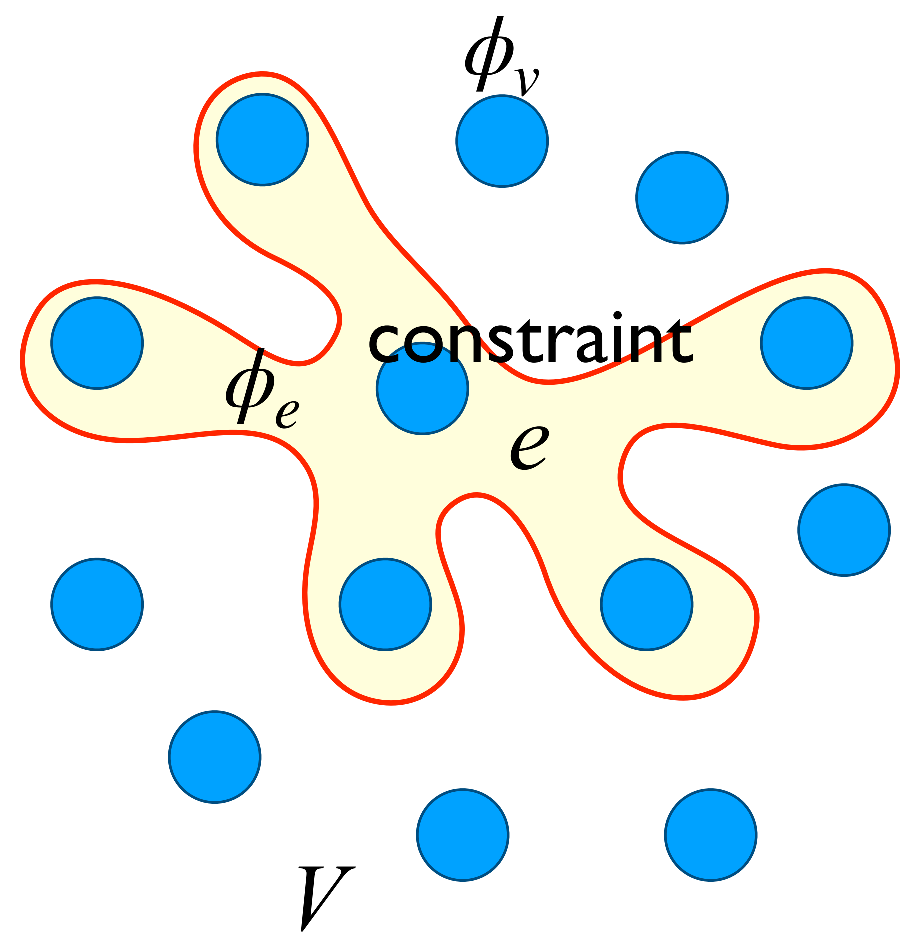

Graphical Model

- Gibbs distribution μ over all $\sigma \in [q]^V$:

$$\mu(\sigma) \propto \prod_{v \in V} \phi_v(\sigma_v) \prod_{e \in E} \phi_e(\sigma_e)$$

ϕ_v is a distribution over $[q]$

$\phi_e : [q]^e \to [0,1]$

- each $v \in V$ independently samples $X_v \in [q]$ according to ϕ_v;
- each $e \in E$ is passed independently with probability $\phi_e(X_e)$;
- X is accepted if all constraints $e \in E$ are passed.

Graphical Model

- Gibbs distribution μ over all $\sigma \in [q]^V$:

$$\mu(\sigma) \propto \prod_{v \in V} \phi_v(\sigma_v) \prod_{e=(u,v)\in E} \phi_e(\sigma_u, \sigma_v)$$

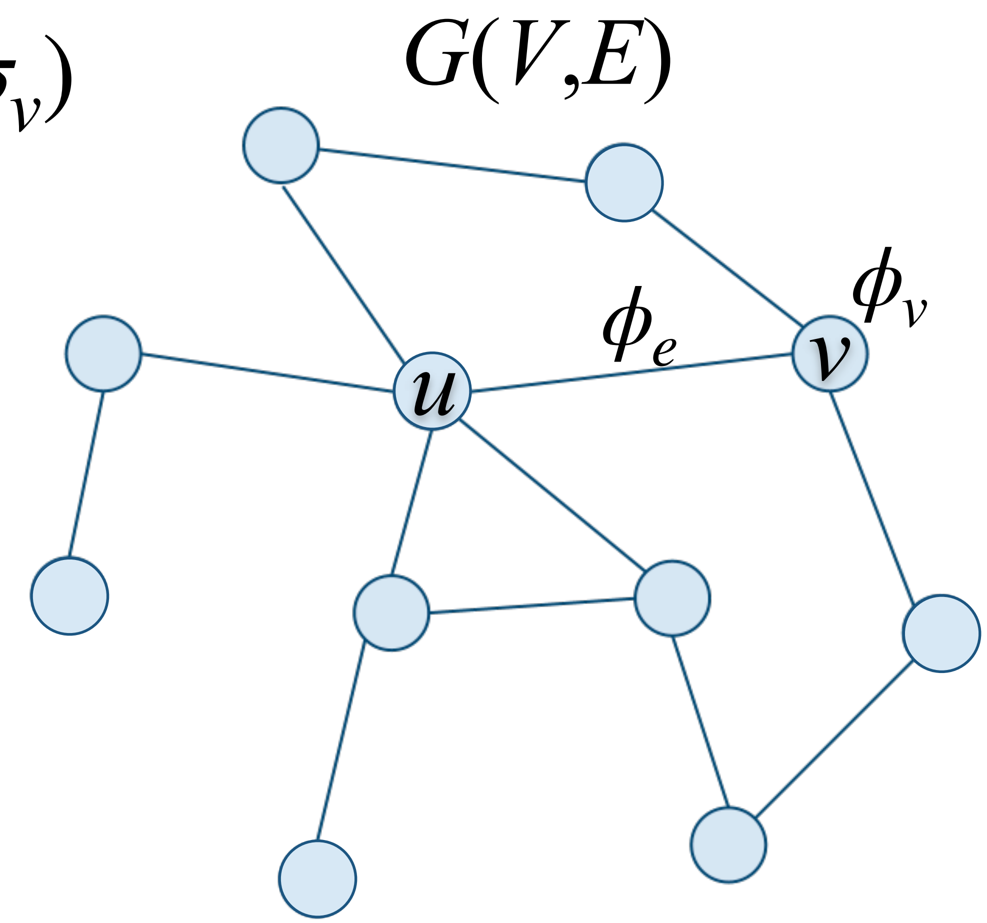

- hardcore morel:

$$[q] = \{0,1\}$$

$$\phi_e(\sigma_u, \sigma_v) = \begin{cases} 0 & \text{if } \sigma_u = \sigma_v = 1 \\ 1 & \text{otherwise} \end{cases}$$

$$\phi_v(\sigma_v) = \begin{cases} \frac{1}{1+\lambda_v} & \text{if } \sigma_v = 0 \\ \frac{\lambda_v}{1+\lambda_v} & \text{if } \sigma_v = 1 \end{cases}$$

$\lambda_v > 0$ is (local) fugacity

Graphical Model

- Gibbs distribution μ over all $\sigma \in [q]^V$:

$$\mu(\sigma) \propto \prod_{v \in V} \phi_v(\sigma_v) \prod_{e=(u,v)\in E} \phi_e(\sigma_u, \sigma_v)$$

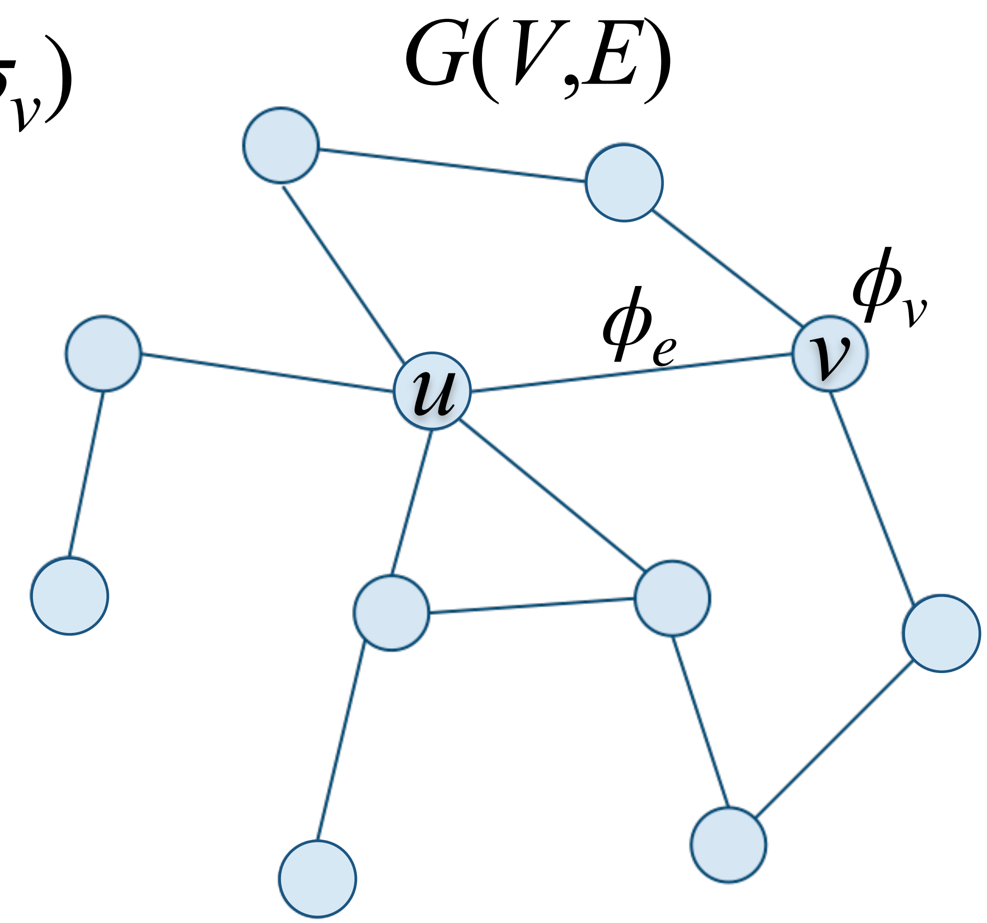

- Ising/Potts model:

or

(ferromagnetic)

$$\phi_e(\sigma_u, \sigma_v) = \begin{cases} 1 & \text{if } \sigma_u = \sigma_v \\ \beta_e \in [0,1] & \text{otherwise} \end{cases}$$

(anti-ferromagnetic)

$$\phi_e(\sigma_u, \sigma_v) = \begin{cases} \beta_e \in [0,1] & \text{if } \sigma_u = \sigma_v \\ 1 & \text{otherwise} \end{cases}$$

ϕ_v is a distribution over $[q]$ (arbitrary local fields)

Dynamic Sampling

- Gibbs distribution μ over all $\sigma\in[q]^V$:

$$\mu(\sigma) \propto \prod_{v\in V}\phi_v(\sigma_v)\prod_{e\in E}\phi_e(\sigma_e)$$

current sample: $X \sim \mu$

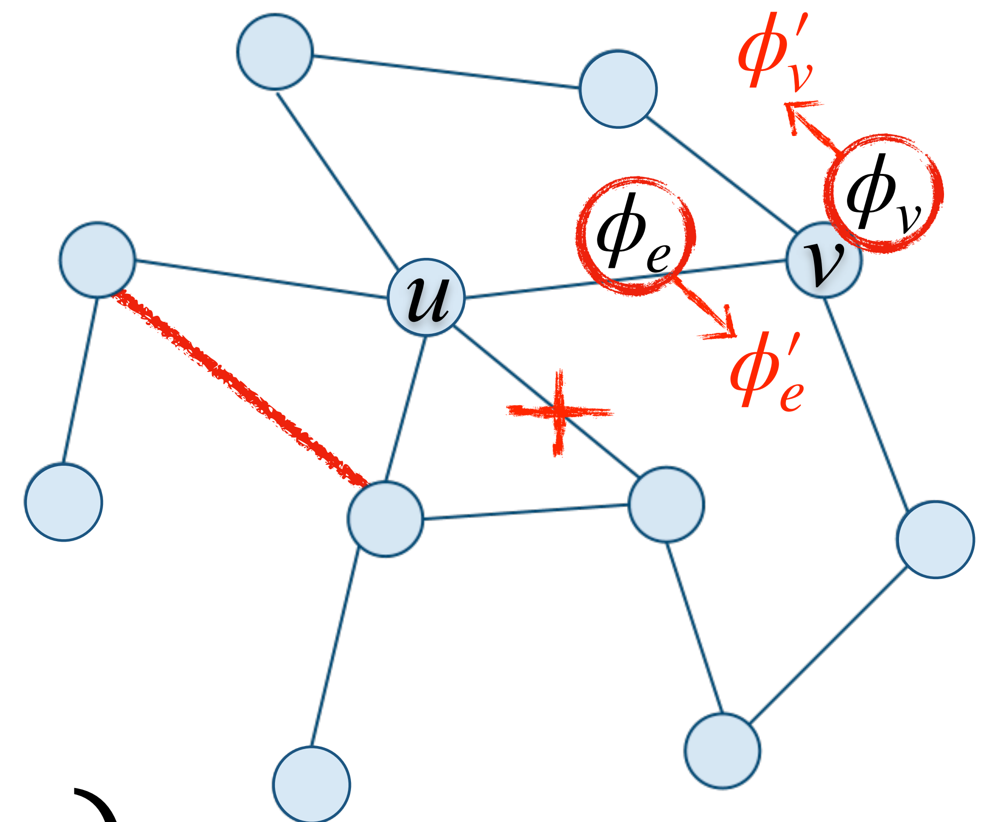

dynamic update:

- adding/deleting a constraint e
- changing a factor ϕ_v or ϕ_e
- adding/deleting an independent variable v

new distribution μ'

Question:

Obtain $X' \sim \mu'$ from $X \sim \mu$ with small *incremental* cost.

Dynamic Sampling

instance of graphical model: $\mathcal{I} = (V, E, [q], \Phi)$

- Gibbs distribution μ over all $\sigma \in [q]^V$:

$$\mu(\sigma) \propto \prod_{v \in V} \phi_v(\sigma_v) \prod_{e \in E} \phi_e(\sigma_e)$$

current sample: $X \sim \mu$

- V : variables
- $E \subset 2^V$: constraints
- $[q] = \{0, 1, \ldots, q\text{-}1\}$: domain
- $\Phi = (\phi_v)_{v \in V} \cup (\phi_e)_{e \in E}$: factors

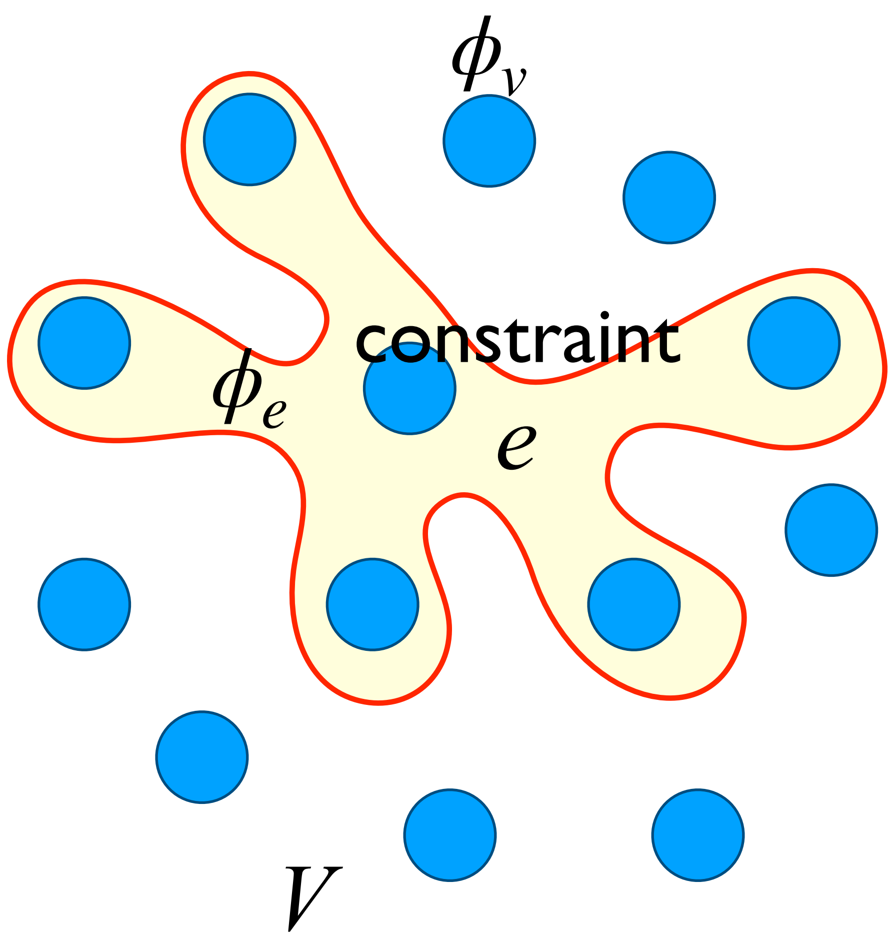

Dynamic Sampling

instance of graphical model: $\mathcal{I} = (V, E, [q], \Phi)$

update: (D, ϕ_D)

$D \subset V \cup 2^V$ is the set of changed variables and constraints

$\Phi_D = (\phi_v)_{v \in V \cap D} \cup (\phi_e)_{e \in 2^V \cap D}$ specifies the new factors

$$(V, E, [q], \Phi) \xrightarrow{(D, \Phi_D)} (V, E', [q], \Phi')$$

$E' = E \cup (2^V \cap D)$

$\Phi' = (\phi'_a)_{a \in V \cup E'}$

where each ϕ'_a is as specified in $\begin{cases} \Phi_D & \text{if } a \in D \\ \Phi & \text{otherwise} \end{cases}$

Dynamic Sampling

instance of graphical model: $\mathcal{I} = (V, E, [q], \Phi)$

update: (D, ϕ_D)

$D \subset V \cup 2^V$ is the set of changed variables and constraints

$\Phi_D = (\phi_v)_{v \in V \cap D} \cup (\phi_e)_{e \in 2^V \cap D}$ specifies the new factors

$$(V, E, [q], \Phi) \xrightarrow{(D, \Phi_D)} (V, E', [q], \Phi')$$

Input: a graphical model with Gibbs distribution μ

a sample $X \sim \mu$, and an update (D, ϕ_D)

Output: $X' \sim \mu'$ where μ' is the new Gibbs distribution

(D, ϕ_D) is fixed by an offline adversary independently of $X \sim \mu$

Dynamic Sampling

Input: a graphical model with Gibbs distribution μ
a sample $X \sim \mu$, and an update (D, ϕ_D)

Output: $X' \sim \mu'$ where μ' is the new Gibbs distribution

- inference/learning tasks where the graphical model is changing dynamically
    - video de-noising
    - online learning with dynamic or streaming data
- sampling/inference/learning algorithms which adaptively and locally change the joint distribution
    - stochastic gradient descent
    - JSV algorithm for perfect matching

Dynamic Sampling

Input: a graphical model with Gibbs distribution μ

a sample $X \sim \mu$, and an update (D, ϕ_D)

Output: $X' \sim \mu'$ where μ' is the new Gibbs distribution

Goal:

transform a $X \sim \mu$ to a $X' \sim \mu'$

by local changes

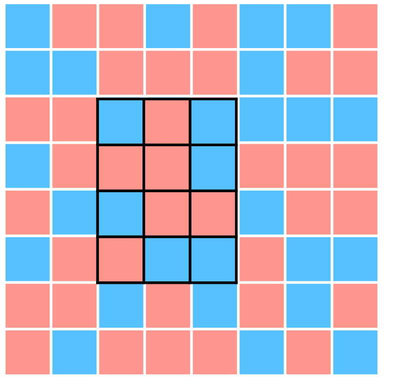

Current sampling techniques are not powerful enough:

- μ may be changed significantly by dynamic updates;
- Monte Carlo sampling does not know when to stop;
- notions such as mixing time give worst-case estimation.

Graphical Model

instance of graphical model: $\mathcal{I} = (V, E, [q], \Phi)$

- V : variables
- $E \subset 2^V$: constraints
- $[q] = \{0,1, \ldots, q\text{-}1\}$: domain
- $\Phi = (\phi_v)_{v\in V} \cup (\phi_e)_{e\in E}$: factors

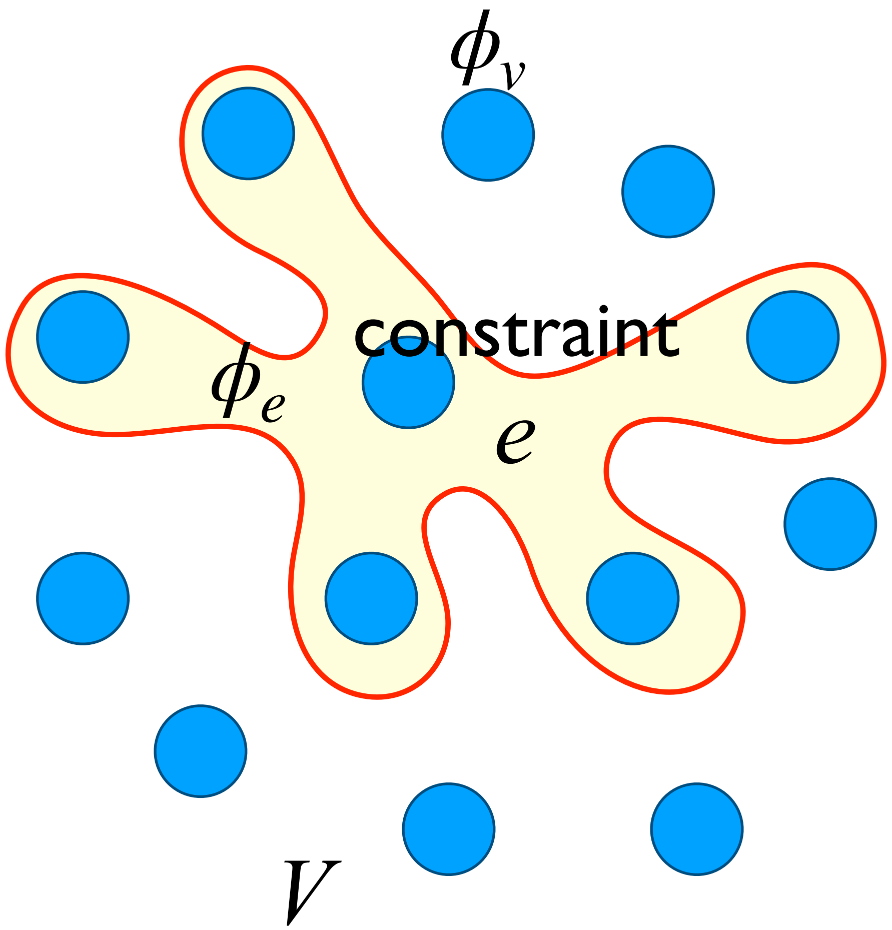

- Gibbs distribution μ over all $\sigma\in[q]^V$:

$$\mu(\sigma) \propto \prod_{v\in V} \phi_v(\sigma_v) \prod_{e\in E} \phi_e(\sigma_e)$$

Notations

instance of graphical model: $\mathcal{I} = (V, E, [q], \Phi)$

for $D \subseteq V \cup 2^V$

$$\mathsf{vbl}(D) \triangleq (V \cap D) \cup \left(\bigcup_{e \in D \cap E} e \right)$$ (involved variables)

for $R \subseteq V$

$$E(R) \triangleq \{e \in E \mid e \subseteq R\}$$ (internal constraints)

$$\delta(R) \triangleq \{e \in E \setminus E(R) \mid e \cap R \neq \varnothing\}$$ (boundary constraints)

$$E^+(R) \triangleq \{e \in E \mid e \cap R \neq \varnothing\}$$ (incident constraints)

$$= E(R) \cup \delta(R)$$

Dynamic Sampler

Input: a graphical model with Gibbs distribution μ
a sample $X \sim \mu$, and an update (D, ϕ_D)

Output: $X' \sim \mu'$ where μ' is the new Gibbs distribution

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \mathrm{vbl}(D) \triangleq (V \cap D) \cup \left(\bigcup_{e \in D \cap E} e \right)$;
- while $R \neq \emptyset$:
    - $(X, R) \leftarrow \texttt{Resample}(X, R)$;

Dynamic Sampler

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow$ vbl$(D) \triangleq (V \cap D) \cup \left(\bigcup_{e \in D \cap E} e \right)$;
- while $R \neq \emptyset$:
  - $(X, R) \leftarrow \texttt{Resample}(X, R)$;

$\texttt{Resample}(X, R)$:

- each $e \in E^+(R)$ computes $\kappa_e = \min_{x_e:\, x_{e \cap R} = X_{e \cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^+(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$;
  (otherwise e is violated)
- $R \leftarrow \bigcup_{e \in E:\ \text{violated } e} e$;

Resampling

$\texttt{Resample}(X, R)$:

- each $e \in E^+(R)$ computes $\kappa_e = \min_{x_e:\, x_{e\cap R} = X_{e\cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^+(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$;
  (otherwise e is violated)
- $R \leftarrow \bigcup_{e \in E:\ \text{violated } e} e$;

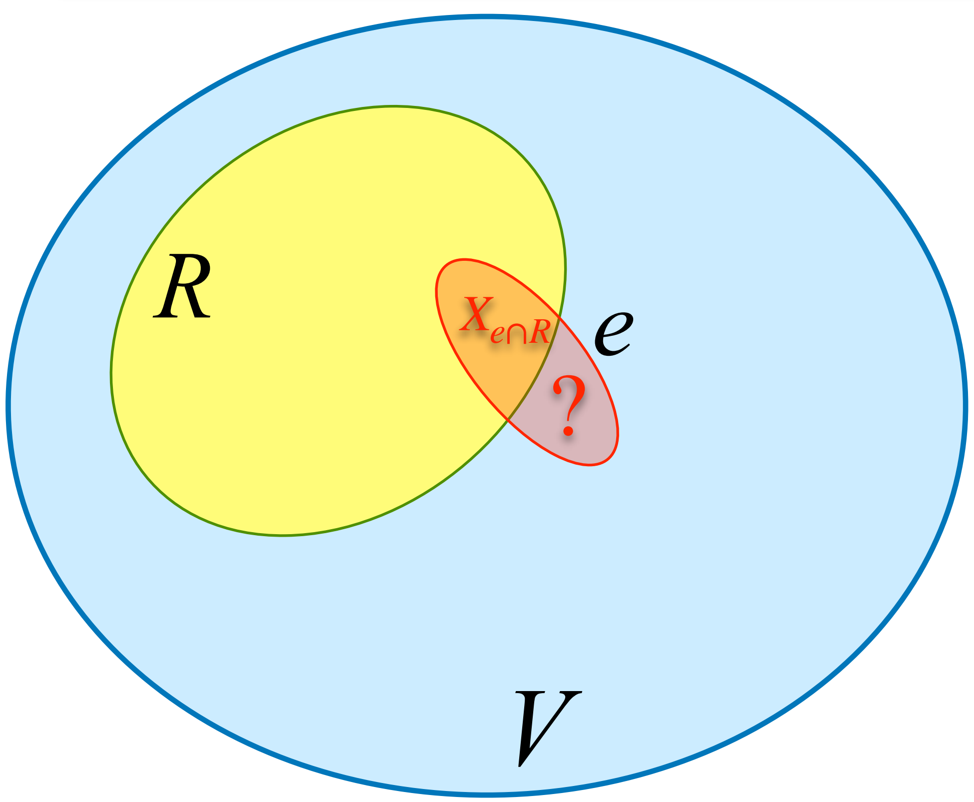

- each *boundary* constraint $e \in \delta(R)$ is violated ind. with prob. $1 - \min_{x_e:\, x_{e\cap R} = X_{e\cap R}} \phi_e(x_e)/\phi_e(X_e)$;
- each $v \in R$ resamples X_v ind. from ϕ_v;
- each non-violated *incident* constraint $e \in E^+(R)$ is violated ind. with prob. $1\text{-}\phi_e(X_e)$;
- all violating variables form the new R;

Resampling

`Resample`(X, R) :

- each $e \in E^+(R)$ computes $\kappa_e = \min_{x_e:\, x_{e\cap R}=X_{e\cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^+(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$;
  (otherwise e is violated)
- $R \leftarrow \bigcup_{e\in E:\ \text{violated } e} e$;

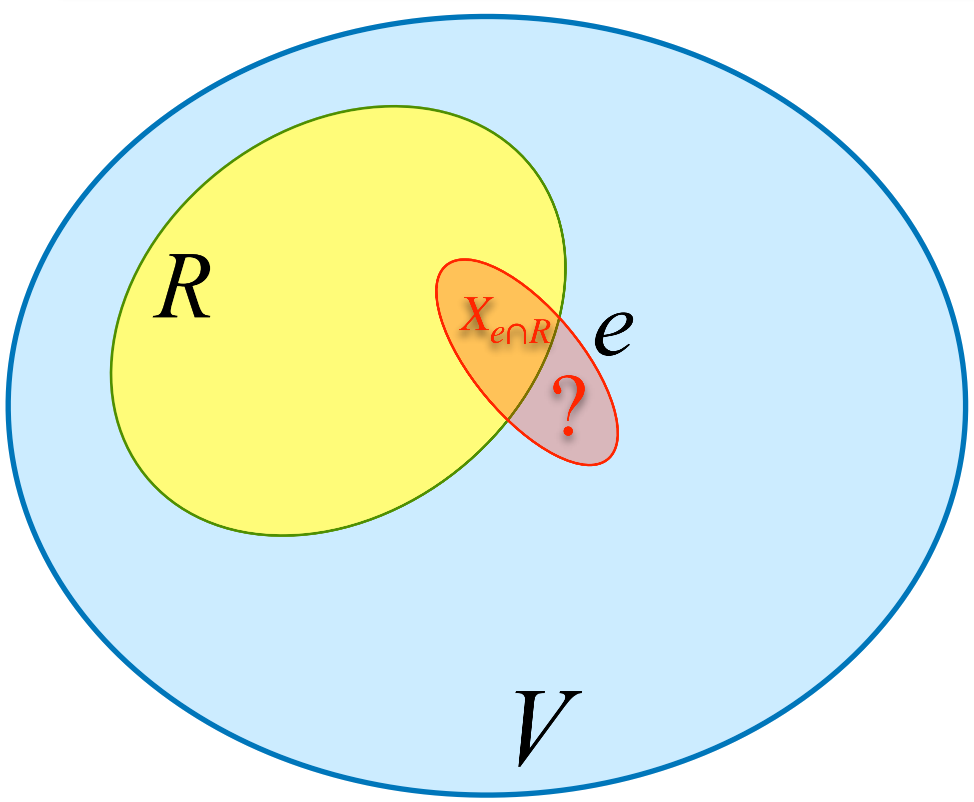

- ~~each boundary constraint~~ $\cancel{e \in \delta(R)}$ ~~is violated~~ ~~ind. with prob.~~ $\cancel{1-\min_{x_e:\, x_{e\cap R}=X_{e\cap R}} \phi_e(x_e)/\phi_e(X_e)}$;
- each $v \in R$ resamples X_v ind. from ϕ_v;
- each ~~non-violated~~ *incident* constraint $e \in E^+(R)$ is violated ind. with prob. $1-\phi_e(X_e)$;
- all violating variables form the new R;

wrong distribution

A more "natural" algorithm?

Dynamic Sampler

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \text{vbl}(D) \triangleq (V \cap D) \cup \left(\bigcup_{e \in D \cap E} e\right)$;
- while $R \neq \emptyset$:
  - $(X, R) \leftarrow \texttt{Resample}(X, R)$;

$\texttt{Resample}(X, R)$:

- each $e \in E^{+}(R)$ computes $\kappa_e = \min_{x_e:\, x_{e \cap R} = X_{e \cap R}} \phi_e(x_e) / \phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^{+}(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$; (otherwise e is violated)
- $R \leftarrow \bigcup_{e \in E:\ \text{violated}\ e} e$;

Correctness of Sampling

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \mathsf{vbl}(D) \triangleq (V \cap D) \cup \left(\bigcup_{e \in D \cap E} e\right)$;
- while $R \neq \varnothing$:
  - $(X, R) \leftarrow \texttt{Resample}(X, R)$;

$\texttt{Resample}(X, R)$:

- each $e \in E^{+}(R)$ computes $\kappa_e = \min_{x_e:\, x_{e \cap R} = X_{e \cap R}} \phi_e(x_e) / \phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^{+}(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$; (otherwise e is violated)
- $R \leftarrow \bigcup_{e \in E:\ \text{violated } e} e$;

Correctness:

Assuming input sample $X \sim \mu$, upon termination, the dynamic sampler returns a sample from the updated distribution μ'.

Fast Convergence

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \mathsf{vbl}(D) \triangleq (V \cap D) \cup \left(\bigcup_{e \in D \cap E} e\right)$;
- while $R \neq \emptyset$:
  - $(X, R) \leftarrow \texttt{Resample}(X, R)$;

$\texttt{Resample}(X, R)$:

- each $e \in E^{+}(R)$ computes $\kappa_e = \min_{x_e:\, x_{e \cap R} = X_{e \cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^{+}(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$;
  (otherwise e is violated)
- $R \leftarrow \bigcup_{e \in E:\ \text{violated } e} e$;

Sufficient Condition for Fast Convergence:

If the followings hold for the updated graphical model:

$$\forall e \in E, \quad \phi_e : [q]^e \to [\beta_e, 1] \quad \text{and} \quad \beta_e > \sqrt{1 - \frac{1}{d+1}}$$

where $d \triangleq \max_{e \in E} |\{e' \in E \setminus \{e\} \mid e' \cap e \neq \emptyset\}|$ is the max-degree of the dependency graph, then:

- \# of iterations is $O(\log |D|)$ in expectation;
- total # of resamplings is $O(kd\,|D|)$ in expectation;

where $k \triangleq \max_{e \in E} |e|$ is the max-size of constraint.

Fast Convergence

- general graphical model with $k = \mathrm{O}(1)$ and $d = \mathrm{O}(1)$:

$$\boxed{\forall e \in E, \quad \phi_e : [q]^e \to [\beta_e, 1] \quad \text{and} \quad \beta_e > \sqrt{1 - \frac{1}{d+1}}}$$

- Ising model of max-degree Δ=O(1): $\boxed{\beta_e > 1 - \frac{1}{2.221\Delta + 1}}$

(ferro-)
$$\phi_e(\sigma_u, \sigma_v) = \begin{cases} 1 & \text{if } \sigma_u = \sigma_v \\ \beta_e \in [0,1] & \text{otherwise} \end{cases}$$

(anti-ferro-)
$$\phi_e(\sigma_u, \sigma_v) = \begin{cases} \beta_e \in [0,1] & \text{if } \sigma_u = \sigma_v \\ 1 & \text{otherwise} \end{cases}$$

uniqueness condition: $\beta > 1 - \frac{2}{\Delta}$

- hardcore model of max-degree Δ=O(1): $\boxed{\lambda_v < \frac{1}{\sqrt{2}\Delta - 1}}$

uniqueness condition: $\lambda < \frac{(\Delta - 1)^{(\Delta-1)}}{(\Delta - 2)^{\Delta}} \approx \frac{\mathrm{e}}{\Delta - 2}$

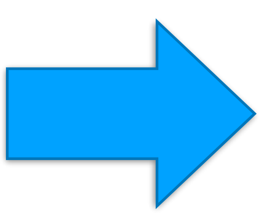

- # of iterations is $\mathrm{O}(\log |D|)$ in expectation;
- total # of resamplings is $\mathrm{O}(|D|)$ in expectation.

Correctness of Sampling

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \mathsf{vbl}(D) \triangleq (V \cap D) \cup \left(\bigcup_{e \in D \cap E} e\right)$;
- while $R \neq \emptyset$:
  - $(X, R) \leftarrow \texttt{Resample}(X, R)$;

$\texttt{Resample}(X, R)$:

- each $e \in E^+(R)$ computes $\kappa_e = \min_{x_e:\, x_{e \cap R} = X_{e \cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^+(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$;
  (otherwise e is violated)
- $R \leftarrow \bigcup_{e \in E:\ \text{violated } e} e$;

Correctness:

Assuming input sample $X \sim \mu$, upon termination, the dynamic sampler returns a sample from the updated distribution μ'.

Resampling

`Resample`(X, R) :

- each $e \in E^+(R)$ computes $\kappa_e = \min\limits_{x_e:\, x_{e\cap R}=X_{e\cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^+(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$;
  (otherwise e is violated)
- $R \leftarrow \bigcup_{e\in E:\ \text{violated } e} e$;

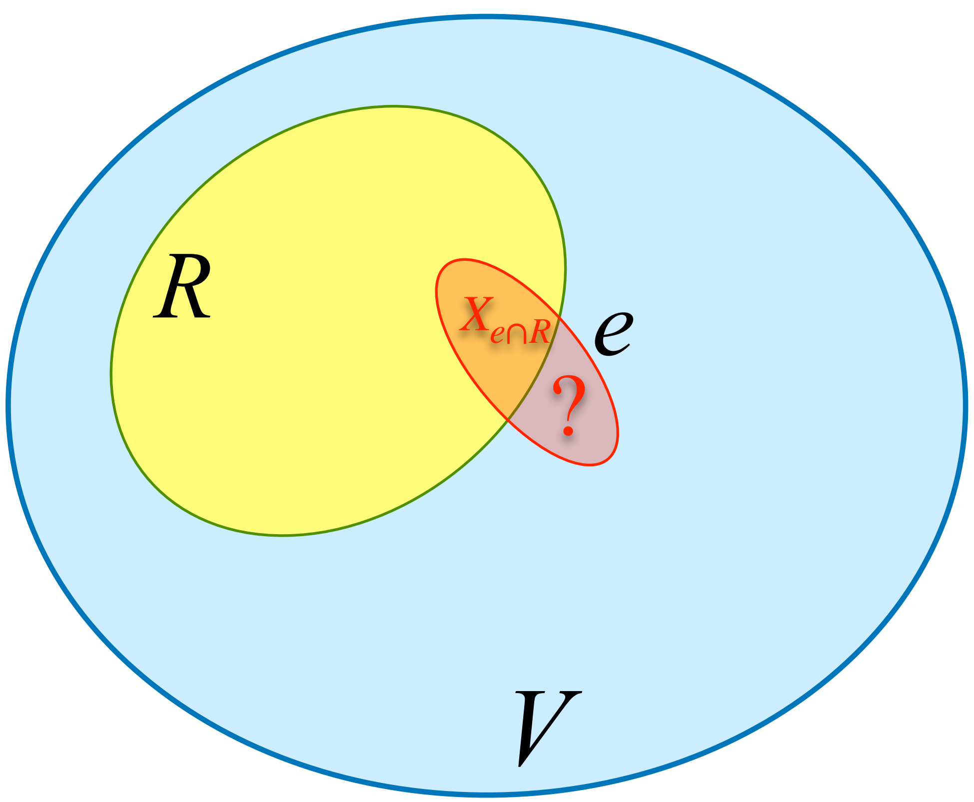

- ~~each *boundary* constraint $e \in \delta(R)$ is violated ind. with prob. $1-\min\limits_{x_e:\, x_{e\cap R}=X_{e\cap R}} \phi_e(x_e)/\phi_e(X_e)$;~~
- each $v \in R$ resamples X_v ind. from ϕ_v;
- each ~~non-violated~~ *incident* constraint $e \in E^+(R)$ is violated ind. with prob. $1-\phi_e(X_e)$;
- all violating variables form the new R;

wrong distribution

A more "natural" algorithm?

Resampling Chain

$\texttt{Resample}(X,R)$:

- each $e \in E^+(R)$ computes $\kappa_e = \min\limits_{x_e:\, x_{e\cap R}=X_{e\cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^+(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$;
  (otherwise e is violated)
- $R \leftarrow \bigcup_{e\in E:\ \text{violated } e} e$;

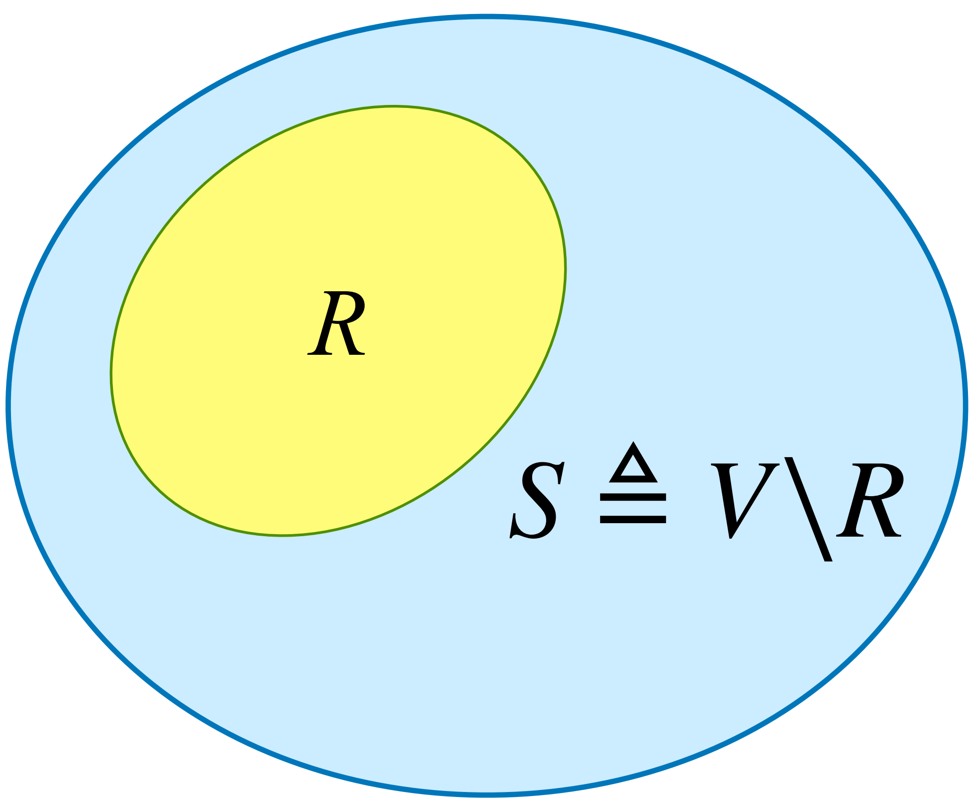

resampling chain $(X,S) \rightarrow (X',S')$:

$S \leftarrow V \setminus R$;

$(X', R') \leftarrow \texttt{Resample}(X,R)$;

$R' \leftarrow V \setminus S'$;

stops when $S = V$

Conditional Gibbs Property

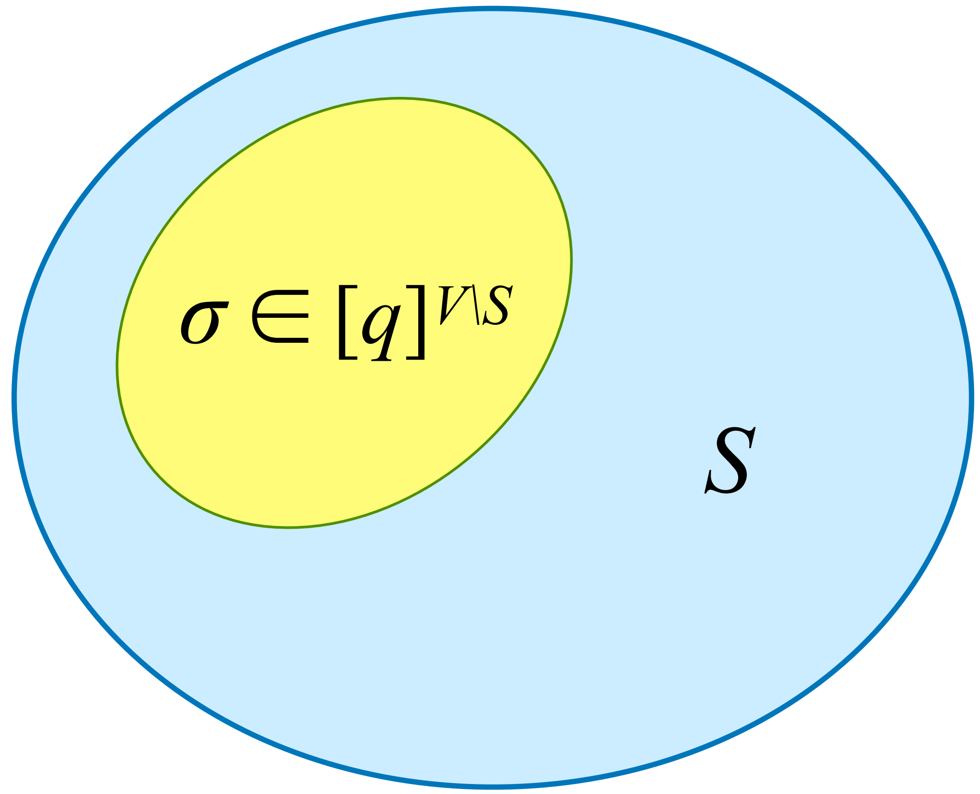

Conditional Gibbs Property:

A random $(X,S) \in [q]^V \times 2^V$ is conditionally Gibbs *w.r.t.* μ if conditioning on any $S \subseteq V$ and any assignment $\sigma \in [q]^{V \setminus S}$ of $X_{V \setminus S}$, the distribution of X_S is precisely μ_S^σ.

μ_S^σ : marginal distribution of μ on S conditioning on σ

Equilibrium

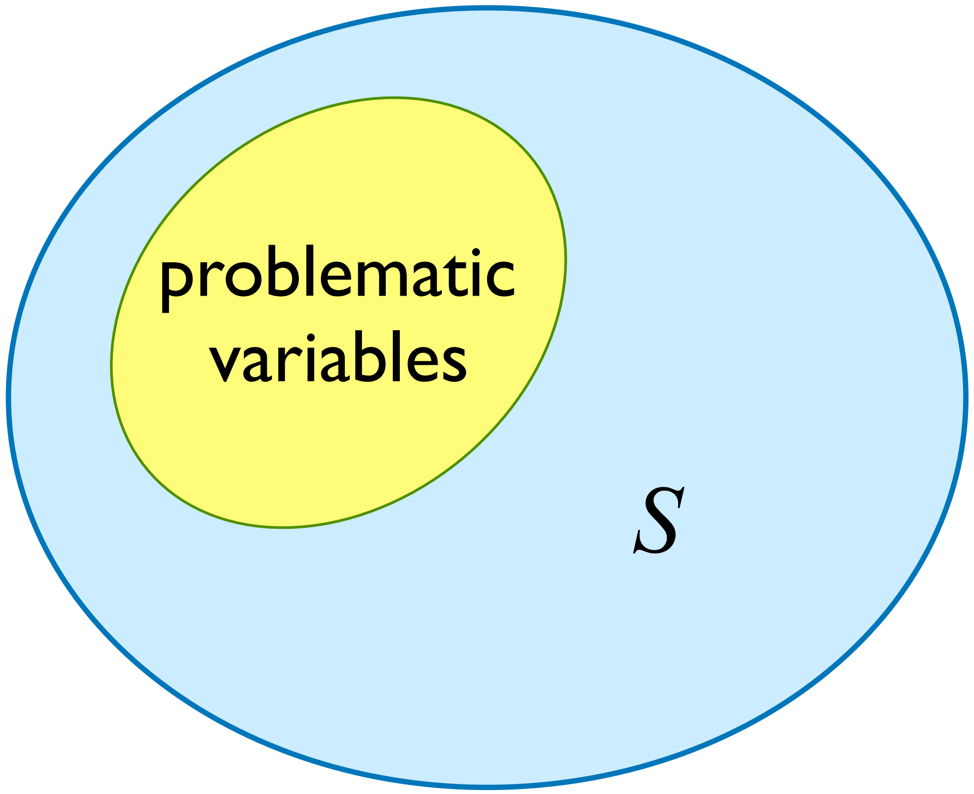

Markov chain $\mathfrak{M}$ on space $[q]^V \times 2^V$:

$$(X, S) \to (X', S')$$

Equilibrium:

If (X,S) is conditionally Gibbs *w.r.t.* μ, then so is (X',S').

Conditional Gibbs Property:

conditioning on any $S \subseteq V$ and any assignment $\sigma \in [q]^{V \setminus S}$ of $X_{V \setminus S}$:
the distribution of X_S is μ_S^σ.

μ_S^σ : marginal distribution of μ on S conditioning on σ

Correctness of Sampling

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \mathsf{vbl}(D) \triangleq (V \cap D) \cup \left(\bigcup_{e \in D \cap E} e\right)$;
- while $R \neq \varnothing$:
  - $(X, R) \leftarrow \mathtt{Resample}(X, R)$;

$\mathtt{Resample}(X, R)$:

- each $e \in E^{+}(R)$ computes $\kappa_e = \min_{x_e :\, x_{e \cap R} = X_{e \cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^{+}(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$;
  (otherwise e is violated)
- $R \leftarrow \bigcup_{e \in E:\ \text{violated } e} e$;

resampling chain

$(X,S) \rightarrow (X',S')$:

$S \leftarrow V \setminus R$;

$(X', R') \leftarrow \mathtt{Resample}(X, R)$;

$R' \leftarrow V \setminus S'$;

Equilibrium:

If (X,S) is conditionally Gibbs w.r.t. μ', then so is (X',S').

Dynamic correctness: Assuming input sample $X \sim \mu$, upon termination, the dynamic sampler returns a sample from the updated distribution μ'.

Equilibrium

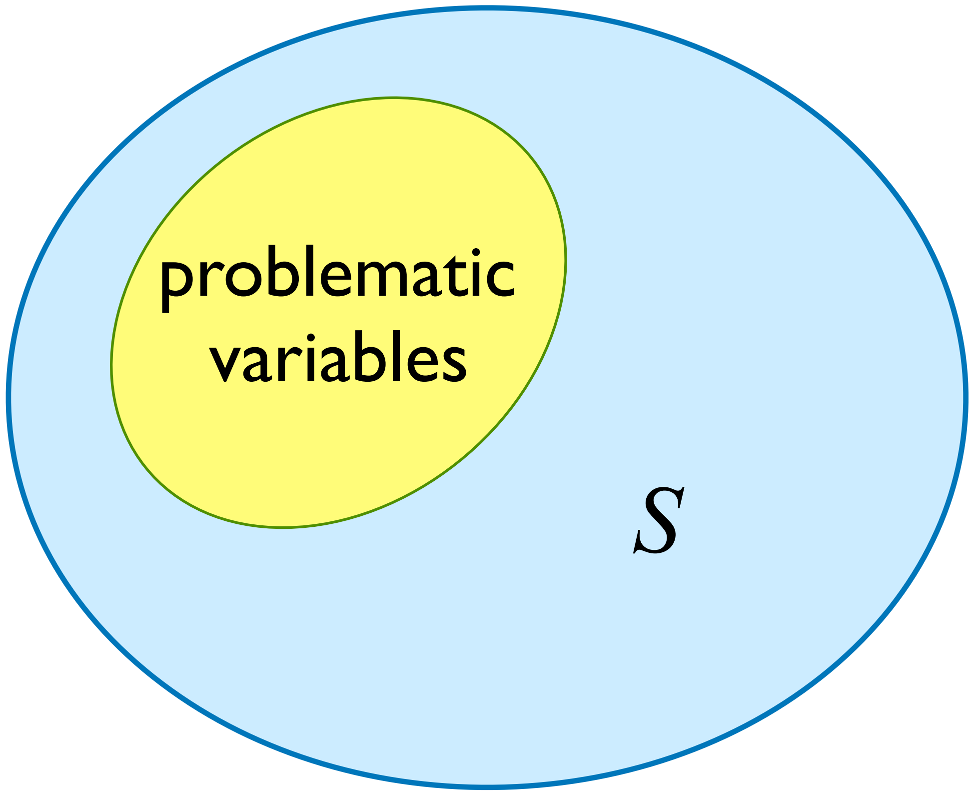

Markov chain $\mathfrak{M}$ on space $[q]^V \times 2^V$:

$$(X, S) \to (X', S')$$

Equilibrium:

If (X,S) is conditional Gibbs *w.r.t.* μ, then so is (X',S').

Conditional Gibbs Property:

conditioning on any $S \subseteq V$ and any assignment $\sigma \in [q]^{V \setminus S}$ of $X_{V \setminus S}$: the distribution of X_S is μ_S^σ.

μ_S^σ : marginal distribution of μ on S conditioning on σ

Equilibria

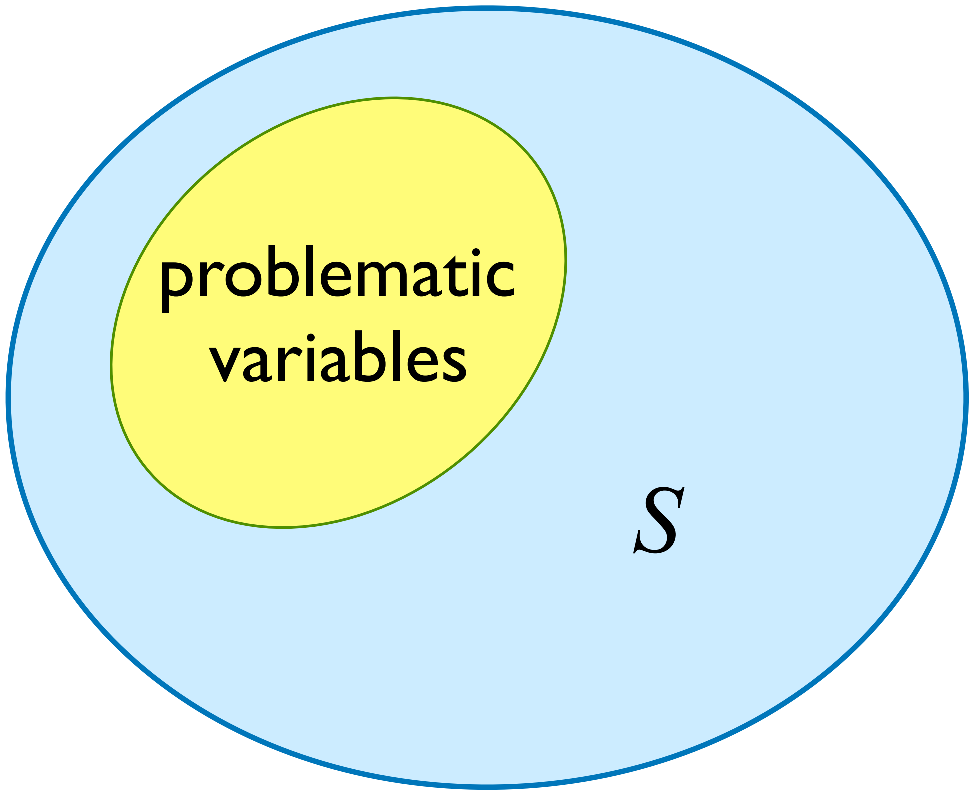

Markov chain $\mathfrak{M}$ on space $[q]^V \times 2^V$:

transition kernel P: $(X, S) \to (X', S')$

Equilibrium:

If (X,S) is conditional Gibbs *w.r.t.* μ,
then so is (X',S').

Refined Equilibrium:

Fix any $S \subseteq V, \sigma \in [q]^{V \setminus S}, T \subseteq V, \tau \in [q]^{V \setminus T}$.

$\forall y \in [q]^V$ where $y_{V \setminus T} = \tau$: $\displaystyle\sum_{\substack{x \in [q]^V \\ x_{V \setminus S} = \sigma}} \mu_S^\sigma(x_S) \cdot P((x, S), (y, T)) \propto \mu_T^\tau(y_T)$

Fixed any $S \subseteq V$ and any assignment $\sigma \in [q]^{V \setminus S}$ of $X_{V \setminus S}$,
the (X',S') is still conditional Gibbs *w.r.t.* μ.

$\texttt{Resample}(X,R)$:

- each $e\in E^{+}(R)$ computes $\kappa_e = \min_{x_e:\, x_{e\cap R}=X_{e\cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v\in R$ resamples $X_v\in[q]$ independently according to ϕ_v;
- each $e\in E^{+}(R)$ is passed independently with prob. $\kappa_e\cdot\phi_e(X_e)$; (otherwise e is violated)
- $R \leftarrow \bigcup_{e\in E:\ \text{violated } e} e$;

resampling chain $(X,S)\rightarrow(X',S')$:

$S \leftarrow V\setminus R$;

$(X',R') \leftarrow \texttt{Resample}(X,R)$;

$R' \leftarrow V\setminus S'$;

transition matrix P

<u>Refined Equilibrium</u>:

Fix any $S\subseteq V, \sigma\in[q]^{V\setminus S}, T\subseteq V, \tau\in[q]^{V\setminus T}$.

$\forall y\in[q]^V$ where $y_{V\setminus T}=\tau$: $\boxed{\sum_{\substack{x\in[q]^V\\ x_{V\setminus S}=\sigma}} \mu_S^{\sigma}(x_S)\cdot P((x,S),(y,T))} \propto \mu_T^{\tau}(y_T)$

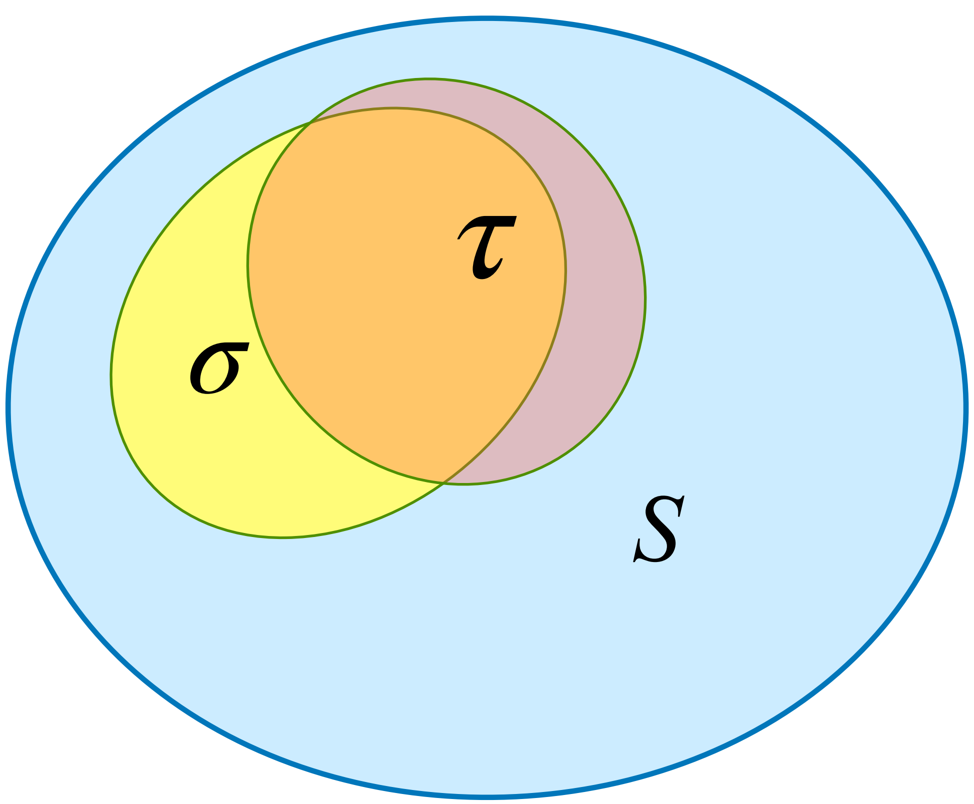

$\mu_S^{\sigma}(x_S)\cdot P((x,S),(y,T))$

where $x\in[q]^V$ is constructed as :

$$x_v = \begin{cases} \sigma_v & v\in V\setminus S \\ y_v & v\in S \end{cases}$$

`Resample`(X,R) :

- each $e\in E^+(R)$ computes $\kappa_e = \min_{x_e:\, x_{e\cap R}=X_{e\cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v\in R$ resamples $X_v\in[q]$ independently according to ϕ_v;
- each $e\in E^+(R)$ is passed independently with prob. $\kappa_e\cdot\phi_e(X_e)$; (otherwise e is violated)
- $R\leftarrow \bigcup_{e\in E:\ \text{violated } e} e$;

resampling chain $(X,S)\to(X',S')$:

$S\leftarrow V\setminus R$;
$(X',R')\leftarrow$ `Resample`(X,R);
$R'\leftarrow V\setminus S'$;

transition matrix P

Refined Equilibrium:

Fix any $S\subseteq V, \sigma\in[q]^{V\setminus S}, T\subseteq V, \tau\in[q]^{V\setminus T}$.

$\forall y\in[q]^V$ where $y_{V\setminus T}=\tau$ and $x\in[q]^V$ defined as $x_v=\begin{cases}\sigma_v & v\in V\setminus S\\ y_v & v\in S\end{cases}$

$$\mu_S^\sigma(x_S)\cdot P((x,S),(y,T)) \propto \mu_T^\tau(y_T)$$

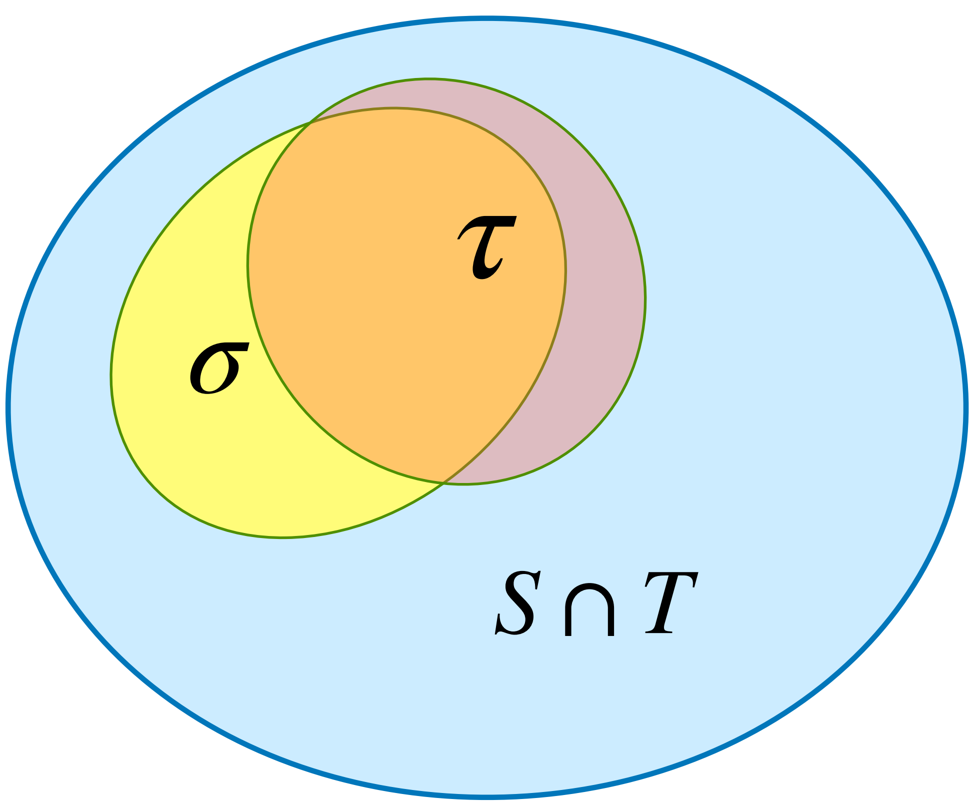

$$\mu_S^\sigma(x_S)\propto\prod_{v\in S}\phi_v(x_v)\prod_{e\in E^+(S)}\phi_e(x_e)$$

$$\propto\prod_{v\in S\cap T}\phi_v(x_v)\prod_{e\in E^+(S)\cap E^+(T)}\phi_e(x_e)$$

$$=\prod_{v\in S\cap T}\phi_v(y_v)\prod_{e\in\delta(S)\cap E^+(T)}\phi_e(x_e)\prod_{e\in E(S)\cap E^+(T)}\phi_e(y_e)$$

`Resample`(X,R) :

- each $e \in E^+(R)$ computes $\kappa_e = \min_{x_e:\, x_{e\cap R}=X_{e\cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^+(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$; (otherwise e is violated)
- $R \leftarrow \bigcup_{e \in E:\text{ violated}} e$;

resampling chain $(X,S) \rightarrow (X',S')$:

$S \leftarrow V \setminus R$;

$(X', R') \leftarrow \texttt{Resample}(X, R)$;

$R' \leftarrow V \setminus S'$;

transition matrix P

Refined Equilibrium:

Fix any $S \subseteq V, \sigma \in [q]^{V \setminus S}, T \subseteq V, \tau \in [q]^{V \setminus T}$.

$\forall y \in [q]^V$ where $y_{V \setminus T} = \tau$ and $x \in [q]^V$ defined as $x_v = \begin{cases} \sigma_v & v \in V \setminus S \\ y_v & v \in S \end{cases}$

$$\mu_S^\sigma(x_S) \cdot P((x,S),(y,T)) \propto \mu_T^\tau(y_T)$$

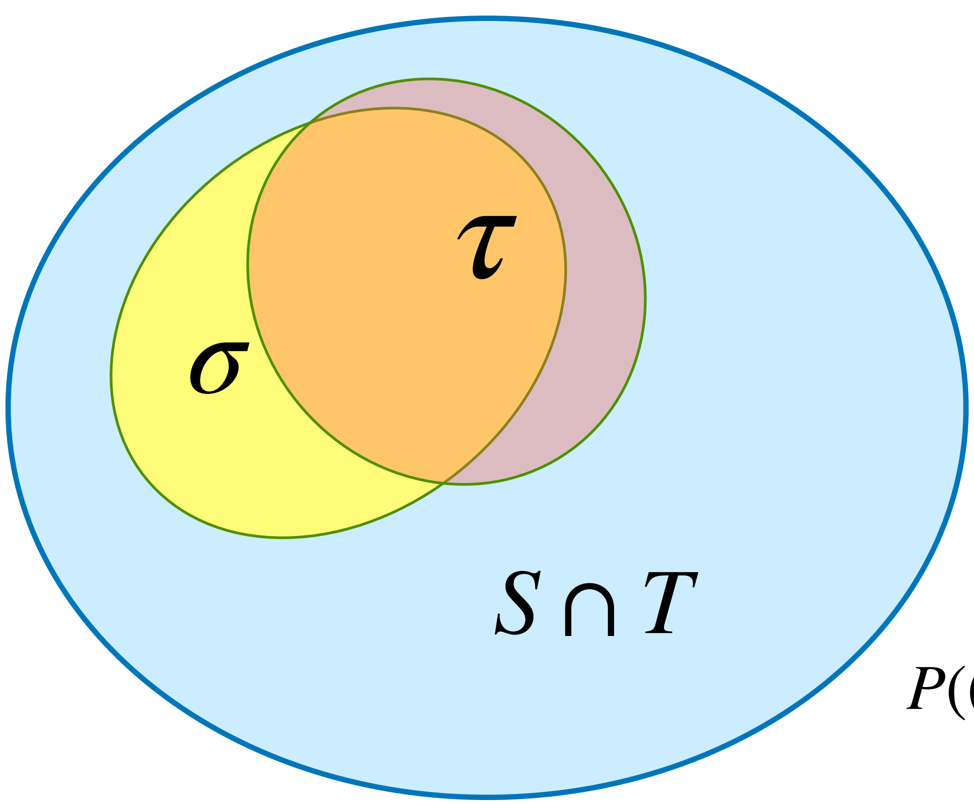

$$\mu_S^\sigma(x_S) \propto \prod_{v \in S \cap T} \phi_v(y_v) \prod_{e \in \delta(S) \cap E^+(T)} \phi_e(x_e) \prod_{e \in E(S) \cap E^+(T)} \phi_e(y_e)$$

$$\mu_T^\tau(y_T) \propto \prod_{v \in T} \phi_v(y_v) \prod_{e \in E^+(T)} \phi_e(y_e)$$

only need:

$$P((x,S),(y,T)) \propto \prod_{v \in T \setminus S} \phi_v(y_v) \prod_{e \in \delta(S) \cap E^+(T)} \frac{\phi_e(y_e)}{\phi_e(x_e)} \prod_{e \in E(V \setminus S) \cap E^+(T)} \phi_e(y_e)$$

$\texttt{Resample}(X,R)$:

- each $e \in E^+(R)$ computes $\kappa_e = \min_{x_e:\, x_{e\cap R} = X_{e \cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^+(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$; (otherwise e is violated)
- $R \leftarrow \bigcup_{e \in E:\text{ violated } e} e$;

resampling chain $(X,S) \to (X',S')$:

$S \leftarrow V \setminus R$;
$(X', R') \leftarrow \texttt{Resample}(X, R)$;
$R' \leftarrow V \setminus S'$;

transition matrix P

Fix any $S \subseteq V, \sigma \in [q]^{V\setminus S}, T \subseteq V, \tau \in [q]^{V \setminus T}$.

$\forall y \in [q]^V$ where $y_{V\setminus T} = \tau$ and $x \in [q]^V$ defined as $x_v = \begin{cases} \sigma_v & v \in V\setminus S \\ y_v & v \in S \end{cases}$

only need: $$P((x,S),(y,T)) \propto \prod_{v \in T\setminus S} \phi_v(y_v) \prod_{e \in \delta(S) \cap E^+(T)} \frac{\phi_e(y_e)}{\phi_e(x_e)} \prod_{e \in E(V\setminus S)\cap E^+(T)} \phi_e(y_e)$$

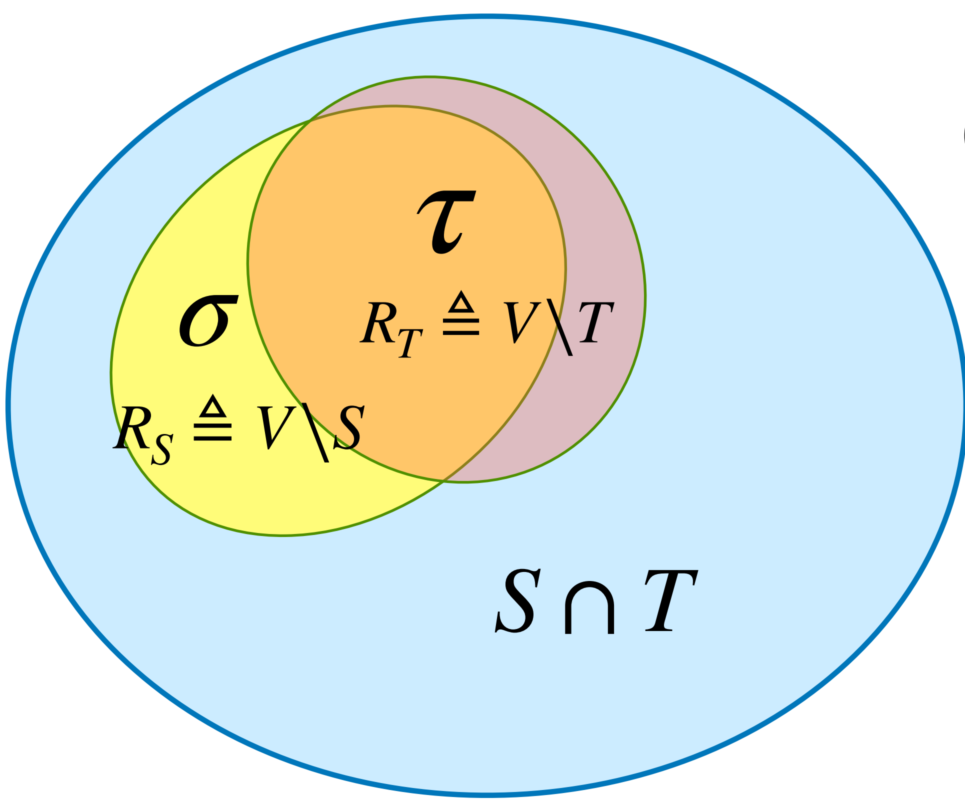

$(x,S) \to (y,T)$ if all these events occur:

A_1 : the resampled configuration is y

A_2 : $\exists F \subseteq E^+(R_S)$ s.t. $\bigcup_{e\in F} e = R_T$ and all $e \in F$ are violated

A_3 : all $e \in E^+(R_S) \setminus E(R_T)$ are passed

$$P((x,S),(y,T)) = \Pr[A_1 \wedge A_2 \wedge A_3]$$

Resample(X,R) :

- each $e \in E^+(R)$ computes $\kappa_e = \min_{x_e:\, x_{e\cap R}=X_{e\cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^+(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$; (otherwise e is violated)
- $R \leftarrow \bigcup_{e\in E:\ \text{violated } e} e$;

resampling chain $(X,S)\to(X',S')$:

$S \leftarrow V\setminus R$;
$(X', R') \leftarrow \texttt{Resample}(X, R)$;
$R' \leftarrow V\setminus S'$;

transition matrix P

Fix any $S \subseteq V, \sigma \in [q]^{V\setminus S}, T \subseteq V, \tau \in [q]^{V\setminus T}$.

$\forall y \in [q]^V$ where $y_{V\setminus T} = \tau$ and $x \in [q]^V$ defined as $x_v = \begin{cases} \sigma_v & v \in V\setminus S \\ y_v & v \in S \end{cases}$

only need: $P((x,S),(y,T)) \propto \prod_{v\in T\setminus S} \phi_v(y_v) \prod_{e\in\delta(S)\cap E^+(T)} \frac{\phi_e(y_e)}{\phi_e(x_e)} \prod_{e\in E(V\setminus S)\cap E^+(T)} \phi_e(y_e)$

$(x,S)\to(y,T)$ if all these events occur:

A_1 : the resampled configuration is y $\qquad \Pr[A_1] \propto \prod_{v\in T\setminus S} \phi_v(y_v)$

A_2 : $\exists F \subseteq E^+(R_S)$ s.t. $\bigcup_{e\in F} e = R_T$ and all $e\in F$ are violated $\qquad \Pr[A_2 \mid A_1] \propto 1$

A_3 : all $e \in E^+(R_S)\setminus E(R_T)$ are passed $\qquad \Pr[A_3 \mid A_1] \propto \prod_{e\in\delta(S)\cap E^+(T)} \frac{\phi_e(y_e)}{\phi_e(x_e)} \prod_{e\in E(V\setminus S)\cap E^+(T)} \phi_e(y_e)$

$$P((x,S),(y,T)) = \Pr[A_1 \wedge A_2 \wedge A_3]$$

`Resample`(X, R) :

- each $e \in E^+(R)$ computes $\kappa_e = \min\limits_{x_e:\, x_{e\cap R} = X_{e\cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^+(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$;
  (otherwise e is violated)
- $R \leftarrow \bigcup_{e \in E:\ \text{violated } e} e$;

resampling chain $(X,S) \to (X',S')$:

$S \leftarrow V \setminus R$;
$(X', R') \leftarrow \texttt{Resample}(X, R)$;
$R' \leftarrow V \setminus S'$;

transition matrix P

Refined Equilibrium:

Fix any $S \subseteq V, \sigma \in [q]^{V\setminus S}, T \subseteq V, \tau \in [q]^{V \setminus T}$.

$\forall y \in [q]^V$ where $y_{V\setminus T} = \tau$ and $x \in [q]^V$ defined as $x_v = \begin{cases} \sigma_v & v \in V \setminus S \\ y_v & v \in S \end{cases}$

$$\mu_S^\sigma(x_S) \cdot P((x, S), (y, T)) \propto \mu_T^\tau(y_T)$$

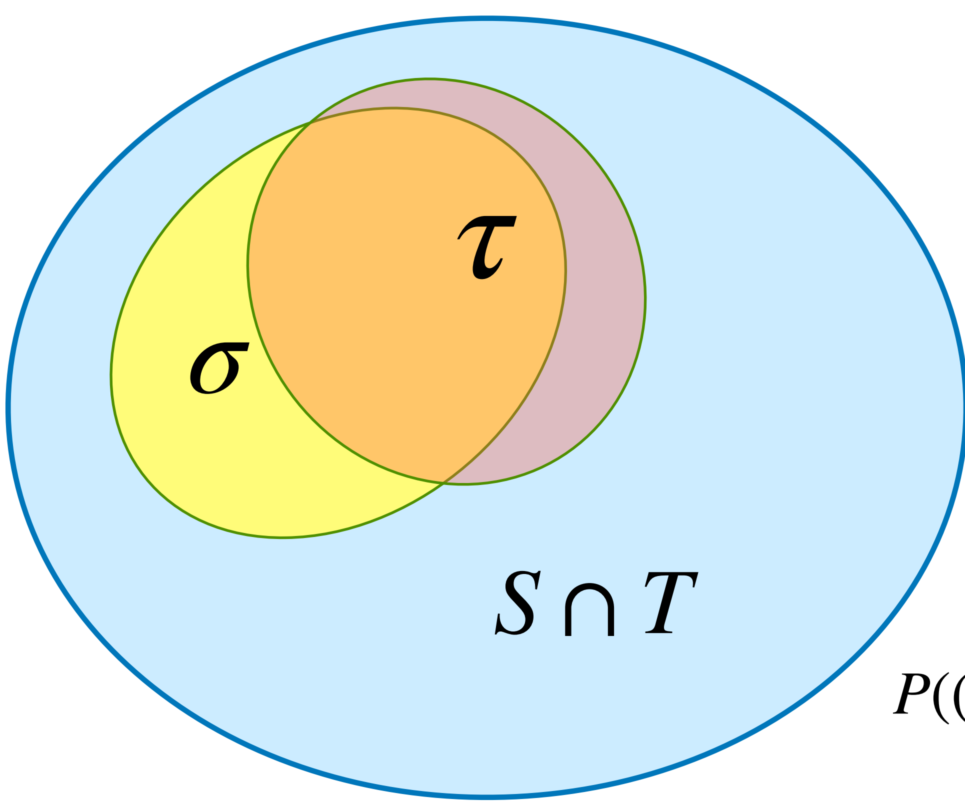

$$\mu_S^\sigma(x_S) \propto \prod_{v \in S\cap T} \phi_v(y_v) \prod_{e \in \delta(S) \cap E^+(T)} \phi_e(x_e) \prod_{e \in E(S) \cap E^+(T)} \phi_e(y_e)$$

$$\mu_T^\tau(y_T) \propto \prod_{v \in T} \phi_v(y_v) \prod_{e \in E^+(T)} \phi_e(y_e)$$

only need:

$$P((x,S),(y,T)) \propto \prod_{v \in T\setminus S} \phi_v(y_v) \prod_{e \in \delta(S)\cap E^+(T)} \frac{\phi_e(y_e)}{\phi_e(x_e)} \prod_{e \in E(V\setminus S)\cap E^+(T)} \phi_e(y_e)$$

Equilibria

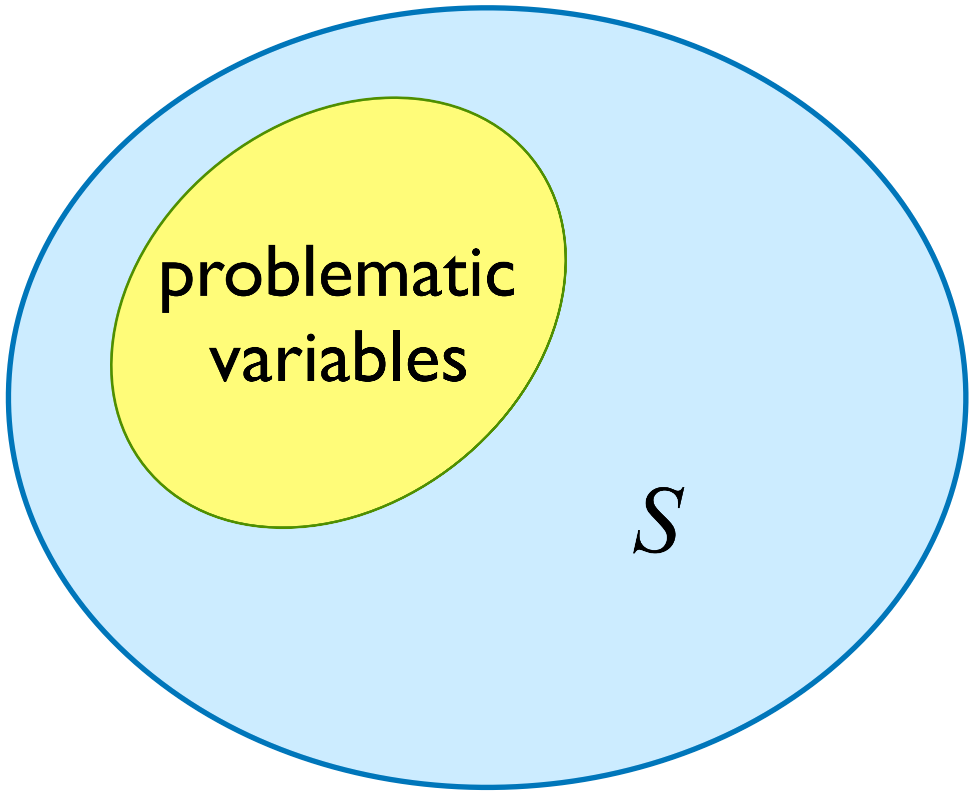

Markov chain $\mathfrak{M}$ on space $[q]^V \times 2^V$:

transition kernel P: $(X, S) \to (X', S')$

Equilibrium:

If (X,S) is conditional Gibbs *w.r.t.* μ,
then so is (X',S').

Refined Equilibrium:

Fix any $S \subseteq V, \sigma \in [q]^{V \setminus S}, T \subseteq V, \tau \in [q]^{V \setminus T}$.

$\forall y \in [q]^V$ where $y_{V \setminus T} = \tau$: $\displaystyle\sum_{\substack{x \in [q]^V \\ x_{V \setminus S} = \sigma}} \mu_S^\sigma(x_S) \cdot P((x, S), (y, T)) \propto \mu_T^\tau(y_T)$

Fixed any $S \subseteq V$ and any assignment $\sigma \in [q]^{V \setminus S}$ of $X_{V \setminus S}$,
the (X',S') is still conditional Gibbs *w.r.t.* μ.

Correctness of Sampling

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \mathsf{vbl}(D) \triangleq (V \cap D) \cup \left(\bigcup_{e \in D \cap E} e\right)$;
- while $R \neq \varnothing$:
  - $(X, R) \leftarrow \texttt{Resample}(X, R)$;

$\texttt{Resample}(X, R)$:

- each $e \in E^{+}(R)$ computes $\kappa_e = \min_{x_e:\, x_{e \cap R} = X_{e \cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^{+}(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$;
  (otherwise e is violated)
- $R \leftarrow \bigcup_{e \in E:\ \text{violated } e} e$;

resampling chain

$(X,S) \rightarrow (X',S')$:

$S \leftarrow V \setminus R$;

$(X', R') \leftarrow \texttt{Resample}(X, R)$;

$R' \leftarrow V \setminus S'$;

Equilibrium:

If (X,S) is conditional Gibbs w.r.t. μ', then so is (X',S').

Dynamic correctness: Assuming input sample $X \sim \mu$, upon termination, the dynamic sampler returns a sample from the updated distribution μ'.

Stronger Adversary

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow$ vbl$(D) \triangleq (V \cap D) \cup \left(\bigcup_{e \in D \cap E} e\right)$;
- while $R \neq \emptyset$:
  - $(X, R) \leftarrow \texttt{Resample}(X, R)$;

(D, ϕ_D) can be adaptive to $X \sim \mu$ as long as $(X, V \setminus R)$ is conditional Gibbs w.r.t. μ'

resampling chain $(X,S) \rightarrow (X',S')$:

(what've been "seen" by the adversary must be resampled)

$S \leftarrow V \setminus R$;

$(X', R') \leftarrow \texttt{Resample}(X, R)$;

$R' \leftarrow V \setminus S'$;

Equilibrium:

If (X,S) is conditional Gibbs w.r.t. μ', then so is (X',S').

Dynamic correctness: Assuming input sample $X \sim \mu$, upon termination, the dynamic sampler returns a sample from the updated distribution μ'.

Fast Convergence

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \mathsf{vbl}(D) \triangleq (V \cap D) \cup \left(\bigcup_{e \in D \cap E} e\right)$;
- while $R \neq \emptyset$:
  - $(X, R) \leftarrow \texttt{Resample}(X, R)$;

$\texttt{Resample}(X, R)$:

- each $e \in E^{+}(R)$ computes $\kappa_e = \min_{x_e:\, x_{e \cap R} = X_{e \cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^{+}(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$; (otherwise e is violated)
- $R \leftarrow \bigcup_{e \in E:\text{ violated } e} e$;

Sufficient Condition for Fast Convergence:

If the followings hold for the updated graphical model:

$$\forall e \in E, \quad \phi_e : [q]^e \to [\beta_e, 1] \quad \text{and} \quad \beta_e > \sqrt{1 - \frac{1}{d+1}}$$

where $d \triangleq \max_{e \in E} |\{e' \in E \setminus \{e\} \mid e' \cap e \neq \emptyset\}|$ is the max-degree of the dependency graph, then:

- # of iterations is $O(\log |D|)$ in expectation;
- total # of resamplings is $O(kd\,|D|)$ in expectation;

where $k \triangleq \max_{e \in E} |e|$ is the max-size of constraint.

Fast Convergence

- general graphical model with $k = O(1)$ and $d = O(1)$:

$$\forall e \in E, \quad \phi_e : [q]^e \to [\beta_e, 1] \quad \text{and} \quad \beta_e > \sqrt{1 - \frac{1}{d+1}}$$

- Ising model of max-degree Δ=O(1):

$$\boxed{\beta_e > 1 - \frac{1}{2.221\Delta + 1}}$$

(ferro-)

$$\phi_e(\sigma_u, \sigma_v) = \begin{cases} 1 & \text{if } \sigma_u = \sigma_v \\ \beta_e \in [0,1] & \text{otherwise} \end{cases}$$

(anti-ferro-)

$$\phi_e(\sigma_u, \sigma_v) = \begin{cases} \beta_e \in [0,1] & \text{if } \sigma_u = \sigma_v \\ 1 & \text{otherwise} \end{cases}$$

uniqueness condition: $\beta > 1 - \frac{2}{\Delta}$

- hardcore model of max-degree Δ=O(1):

$$\lambda_v < \frac{1}{\sqrt{2\Delta - 1}}$$

uniqueness condition: $\lambda < \frac{(\Delta-1)^{(\Delta-1)}}{(\Delta-2)^{\Delta}} \approx \frac{\mathrm{e}}{\Delta - 2}$

- # of iterations is $O(\log |D|)$ in expectation;
- total # of resamplings is $O(|D|)$ in expectation.

Ising Model

- Gibbs distribution μ over all $\sigma \in [q]^V$:

$$\mu(\sigma) \propto \prod_{v \in V} \phi_v(\sigma_v) \prod_{e=(u,v)\in E} \phi_e(\sigma_u, \sigma_v)$$

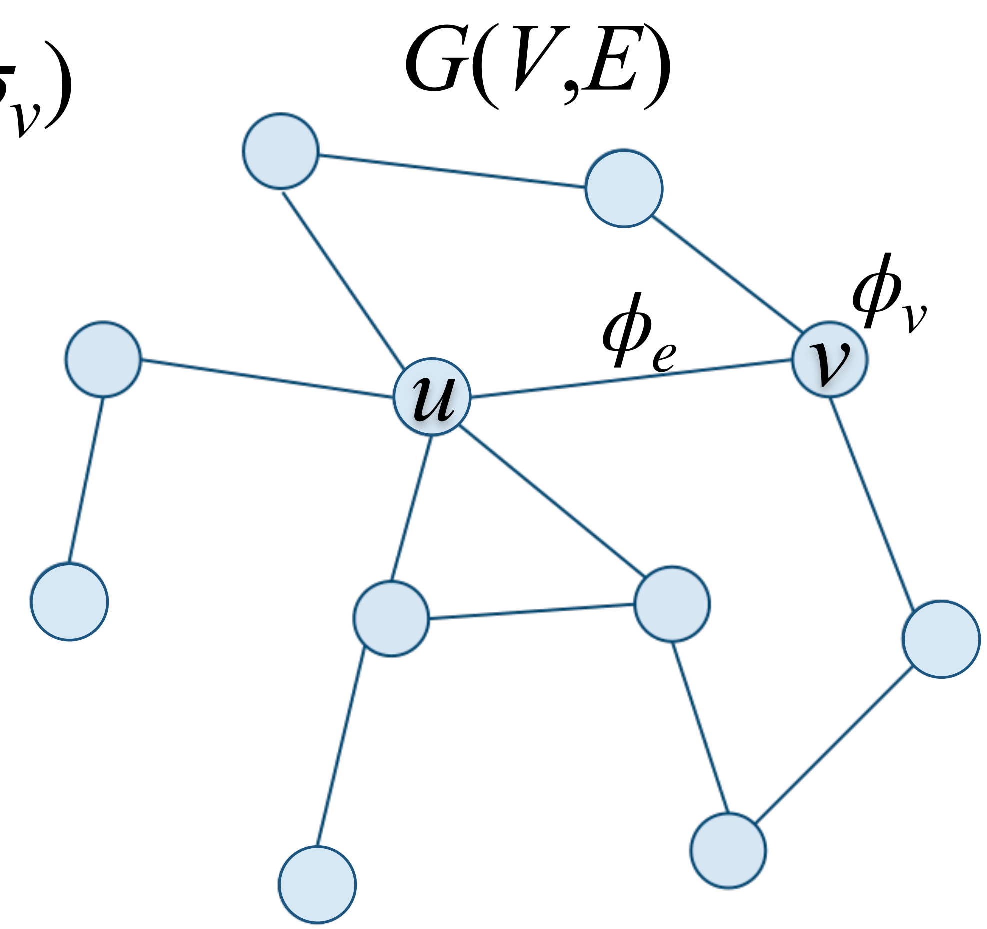

- Ising model: $[q] = \{0,1\}$

or

(ferromagnetic)

$$\phi_e(\sigma_u, \sigma_v) = \begin{cases} 1 & \text{if } \sigma_u = \sigma_v \\ \beta_e \in [0,1] & \text{otherwise} \end{cases}$$

(anti-ferromagnetic)

$$\phi_e(\sigma_u, \sigma_v) = \begin{cases} \beta_e \in [0,1] & \text{if } \sigma_u = \sigma_v \\ 1 & \text{otherwise} \end{cases}$$

ϕ_v is the uniform distribution over $\{0,1\}$

(zero field)

update (D, ϕ_D)

$$D \subseteq \binom{V}{2}$$

Dynamic Ising Sampler

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \mathsf{vbl}(D) \triangleq \bigcup_{e \in E} e$;
- while $R \neq \emptyset$:
  - each $e=(u,v) \in \delta(R)$ is violated independently with prob. $1\text{-}\beta_e / \phi_e(X_u,X_v)$;
  - each $v \in R$ resamples $X_v \in \{0,1\}$ uniformly and independently;
  - each non-violated $e=(u,v) \in E^+(R)$ is violated ind. with prob. $1\text{-}\phi_e(X_u,X_v)$;
  - $R \leftarrow \bigcup_{e \in E:\ \text{violated in the current iteration}} e$;

(ferromagnetic)

$$\phi_e(\sigma_u, \sigma_v) = \begin{cases} 1 & \text{if } \sigma_u = \sigma_v \\ \beta_e \in [0,1] & \text{otherwise} \end{cases}$$

(anti-ferromagnetic)

$$\phi_e(\sigma_u, \sigma_v) = \begin{cases} \beta_e \in [0,1] & \text{if } \sigma_u = \sigma_v \\ 1 & \text{otherwise} \end{cases}$$

Dynamic Ising Sampler

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \mathsf{vbl}(D) \triangleq \bigcup_{e\in E} e$;
- while $R \neq \emptyset$:
  - each $e=(u,v) \in \delta(R)$ is violated independently with prob. $1\text{-}\beta_e / \phi_e(X_u,X_v)$;
  - each $v \in R$ resamples $X_v \in \{0,1\}$ uniformly and independently;
  - each non-violated $e=(u,v) \in E^+(R)$ is violated ind. with prob. $1\text{-}\phi_e(X_u,X_v)$;
  - $R \leftarrow \bigcup_{e\in E:\text{ violated in the current iteration}} e$;

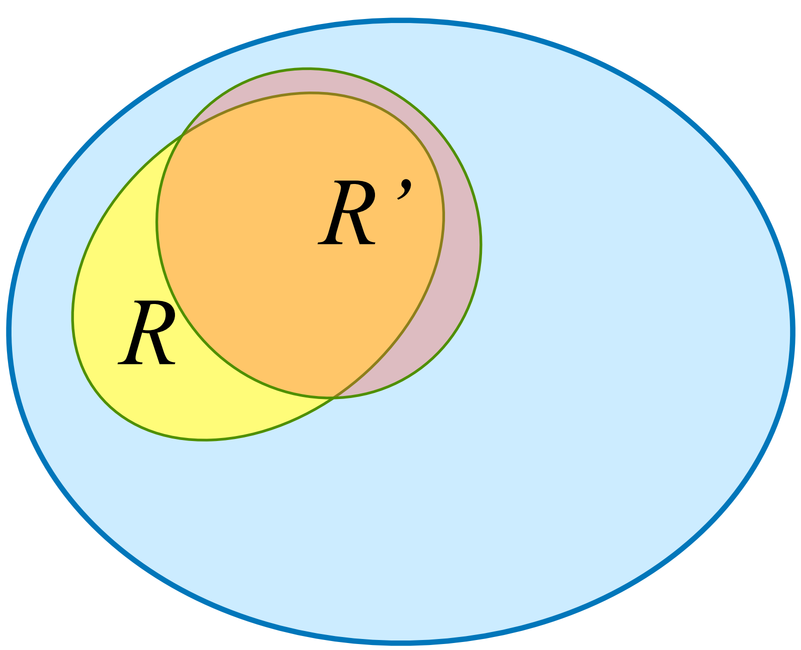

in each iteration: $(X,R) \rightarrow (X',R')$

$\forall R \subseteq V$:

$$\mathbb{E}[H(R') \mid R] \leq (1-\delta)H(R)$$

for some potential function H

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \mathsf{vbl}(D) \triangleq \bigcup_{e\in E} e$;
- while $R \neq \varnothing$:
  - each $e=(u,v) \in \delta(R)$ is violated independently with prob. $1-\beta_e / \phi_e(X_u,X_v)$;
  - each $v \in R$ resamples $X_v \in \{0,1\}$ uniformly and independently;
  - each non-violated $e=(u,v) \in E^+(R)$ is violated ind. with prob. $1-\phi_e(X_u,X_v)$;
  - $R \leftarrow \bigcup_{e\in E:\text{ violated in the current iteration}} e$;

in each iteration: $(X,R) \rightarrow (X',R')$

$\forall R \subseteq V:$

$$\mathbb{E}[H(R') \mid R] \le (1-\delta)H(R)$$

for some potential function H

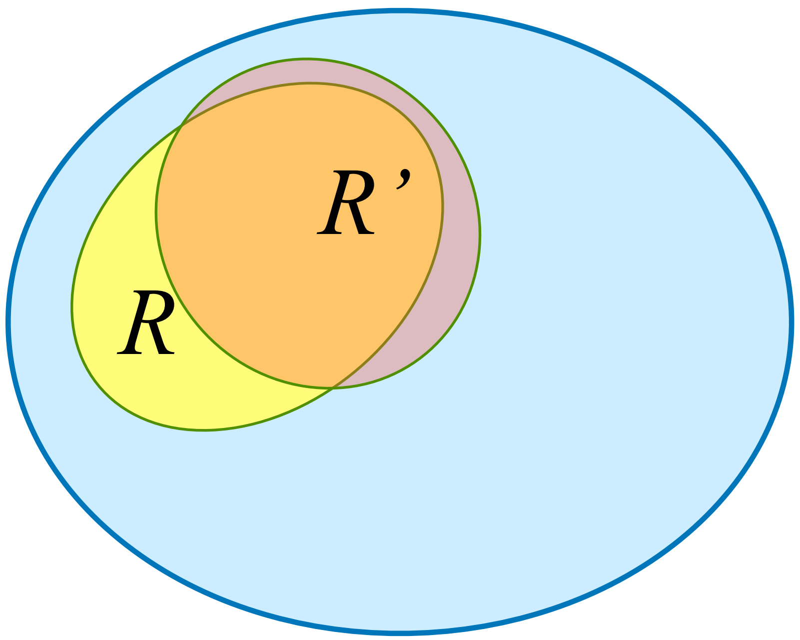

$$H(R) = \min\left\{ |C| \mid C \subseteq E \wedge R \subseteq \bigcup_{e\in C} e \right\}$$

size of minimum edge cover of R

$$\mathbb{E}[H(R') \mid R] \le \sum_{e\in E^+(R)} \Pr[e \text{ is violated}]$$

$$\forall e \in E(R): \quad \Pr[e \text{ is violated}] = \frac{1-\beta_e}{2}$$

$$\forall e = (u,v) \in \delta(R) \text{ where } u \in R, v \notin R: \quad (X, V\backslash R) \text{ is conditional Gibbs}$$

$$\Pr[e \text{ is violated} \mid X, X'] = 1 - \beta_e \phi_e(X'_u, X_v)/\phi_e(X_u, X_v)$$

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \mathrm{vbl}(D) \triangleq \bigcup_{e\in E} e$;
- while $R \neq \emptyset$:
  - each $e=(u,v) \in \delta(R)$ is violated independently with prob. $1-\beta_e / \phi_e(X_u,X_v)$;
  - each $v \in R$ resamples $X_v \in \{0,1\}$ uniformly and independently;
  - each non-violated $e=(u,v) \in E^+(R)$ is violated ind. with prob. $1-\phi_e(X_u,X_v)$;
  - $R \leftarrow \bigcup_{e\in E:\ \text{violated in the current iteration}} e$;

in each iteration: $(X,R) \rightarrow (X',R')$

$$\forall R \subseteq V:$$

$$\mathbb{E}[H(R') \mid R] \le (1-\delta)H(R)$$

for some potential function H

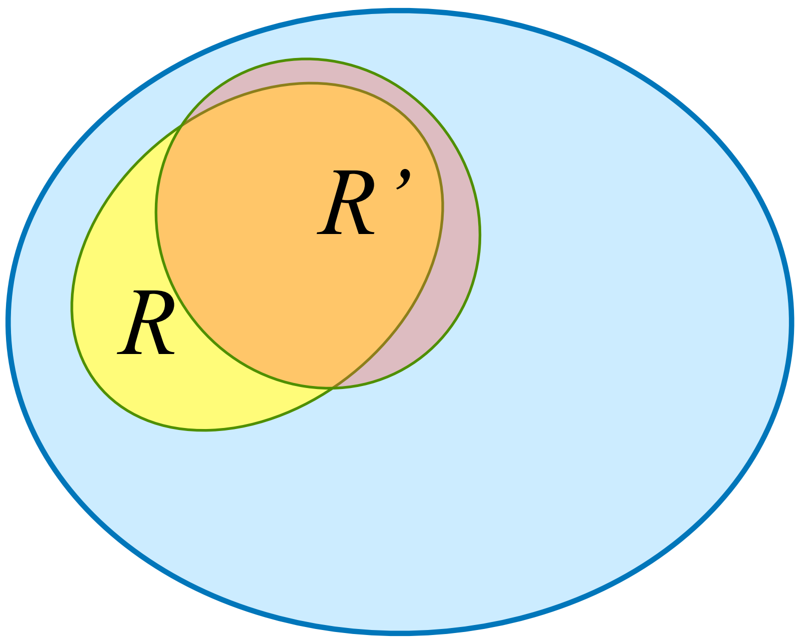

$$H(R) = \min\left\{ |C| \mid C \subseteq E \wedge R \subseteq \bigcup_{e\in C} e \right\}$$

size of minimum edge cover of R

$$\mathbb{E}[H(R') \mid R] \le \sum_{e\in E^+(R)} \Pr[e \text{ is violated}]$$

$$\forall e \in E(R): \quad \Pr[e \text{ is violated}] = \frac{1-\beta_e}{2}$$

$$\forall e = (u,v) \in \delta(R): \quad \Pr[e \text{ is violated}] \le \frac{1-\beta}{2}\left(1 + \frac{1+\beta}{1+\beta^{\Delta}}\right)$$

where $\beta \triangleq \max_{e\in E} \beta_e$

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \mathsf{vbl}(D) \triangleq \bigcup_{e\in E} e$;
- while $R \neq \emptyset$:
  - each $e=(u,v) \in \delta(R)$ is violated independently with prob. $1-\beta_e / \phi_e(X_u,X_v)$;
  - each $v \in R$ resamples $X_v \in \{0,1\}$ uniformly and independently;
  - each non-violated $e=(u,v) \in E^+(R)$ is violated ind. with prob. $1-\phi_e(X_u,X_v)$;
  - $R \leftarrow \bigcup_{e\in E\text{: violated in the current iteration}} e$;

in each iteration: $(X,R)\rightarrow(X',R')$

$\forall R \subseteq V:$

$$\mathbb{E}[H(R') \mid R] \leq (1-\delta)H(R)$$

for some potential function H

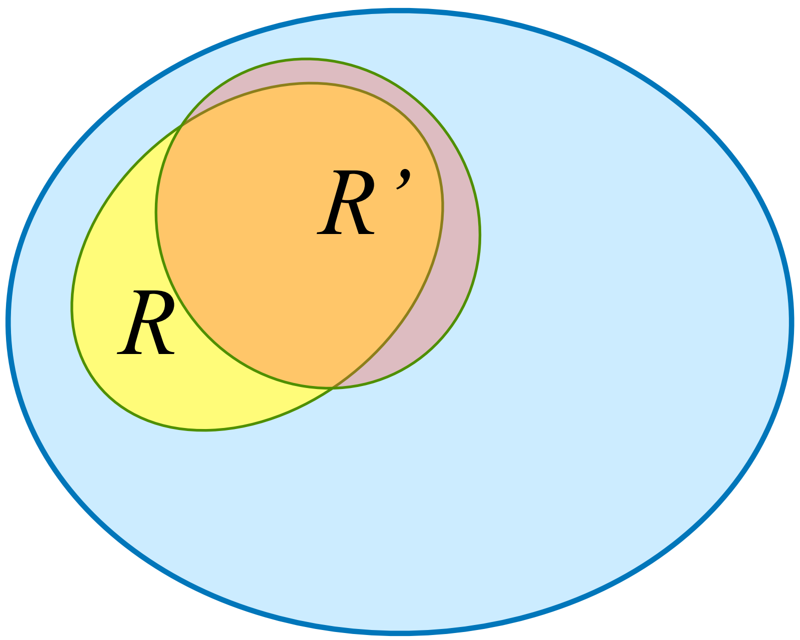

$$H(R) = \min\left\{ |C| \mid C \subseteq E \wedge R \subseteq \bigcup_{e\in C} e \right\}$$

size of minimum edge cover of R

$$\mathbb{E}[H(R') \mid R] \leq \sum_{e\in E^+(R)} \Pr[e \text{ is violated}]$$

where $\beta \triangleq \max_{e\in E} \beta_e$

$$\leq \frac{1-\beta}{2}\left(1+\frac{1+\beta}{1+\beta^\Delta}\right)|E^+(R)|$$

when $\boxed{\beta \geq 1 - \frac{1}{\alpha\Delta+1}}$

$$\leq \frac{(2\Delta-1)(1-\beta)}{2}\left(1+\frac{1+\beta}{1+\beta^\Delta}\right)|H(R)|$$

where $\alpha \approx 2.22$ is the root of $\alpha = 1 + \frac{2}{1+\mathrm{e}^{-1/\alpha}}$

$$\leq (1-\delta)H(R)$$

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \mathsf{vbl}(D) \triangleq \bigcup_{e \in E} e$;
- while $R \neq \emptyset$:
  - each $e=(u,v) \in \delta(R)$ is violated independently with prob. $1\text{-}\beta_e / \phi_e(X_u,X_v)$;
  - each $v \in R$ resamples $X_v \in \{0,1\}$ uniformly and independently;
  - each non-violated $e=(u,v) \in E^+(R)$ is violated ind. with prob. $1\text{-}\phi_e(X_u,X_v)$;
  - $R \leftarrow \bigcup_{e \in E:\text{ violated in the current iteration}} e$;

in each iteration: $(X,R) \rightarrow (X',R')$

$\forall R \subseteq V:$

$$\mathbb{E}[H(R') \mid R] \leq (1-\delta)H(R)$$

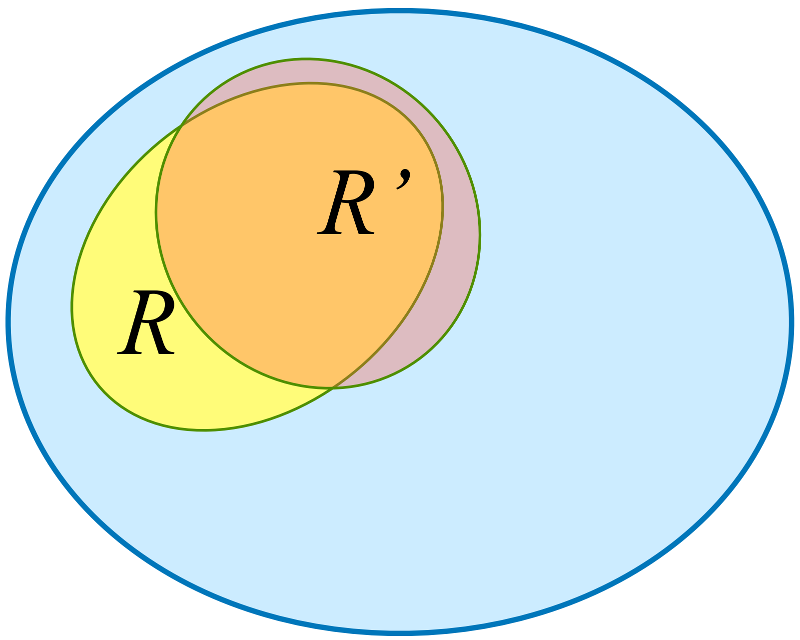

$$H(R) = \min\left\{ |C| \mid C \subseteq E \wedge R \subseteq \bigcup_{e \in C} e \right\}$$

size of minimum edge cover of R

when $\boxed{\beta \geq 1 - \frac{1}{\alpha\Delta + 1}}$ where $\beta \triangleq \max_{e \in E} \beta_e$

where $\alpha \approx 2.22$ is the root of $\alpha = 1 + \frac{2}{1 + \mathrm{e}^{-1/\alpha}}$

$\Delta = O(1)$

- # of iterations is $O(\log |D|)$ in expectation;
- total # of resamplings is $O(|D|)$ in expectation.

Fast Convergence

- general graphical model with $k = O(1)$ and $d = O(1)$:

$$\forall e \in E, \quad \phi_e : [q]^e \to [\beta_e, 1] \quad \text{and} \quad \beta_e > \sqrt{1 - \frac{1}{d+1}}$$

- Ising model of max-degree $\Delta = O(1)$:

$$\boxed{\beta_e > 1 - \frac{1}{2.221\Delta + 1}}$$

(ferro-)

$$\phi_e(\sigma_u, \sigma_v) = \begin{cases} 1 & \text{if } \sigma_u = \sigma_v \\ \beta_e \in [0,1] & \text{otherwise} \end{cases}$$

(anti-ferro-)

$$\phi_e(\sigma_u, \sigma_v) = \begin{cases} \beta_e \in [0,1] & \text{if } \sigma_u = \sigma_v \\ 1 & \text{otherwise} \end{cases}$$

uniqueness condition: $\beta > 1 - \frac{2}{\Delta}$

- hardcore model of max-degree $\Delta = O(1)$:

$$\boxed{\lambda_v < \frac{1}{\sqrt{2\Delta - 1}}}$$

uniqueness condition: $\lambda < \frac{(\Delta-1)^{(\Delta-1)}}{(\Delta-2)^{\Delta}} \approx \frac{e}{\Delta - 2}$

- \# of iterations is $O(\log |D|)$ in expectation;
- total # of resamplings is $O(|D|)$ in expectation.

Fast Convergence

- general graphical model with $k = O(1)$ and $d = O(1)$:

$$\forall e \in E, \quad \phi_e : [q]^e \to [\beta_e, 1] \quad \text{and} \quad \beta_e > \sqrt{1 - \frac{1}{d+1}}$$

- Ising model of max-degree Δ=O(1): $\beta_e > 1 - \frac{1}{2.221\Delta + 1}$

(ferro-) $\phi_e(\sigma_u, \sigma_v) = \begin{cases} 1 & \text{if } \sigma_u = \sigma_v \\ \beta_e \in [0,1] & \text{otherwise} \end{cases}$

(anti-ferro-) $\phi_e(\sigma_u, \sigma_v) = \begin{cases} \beta_e \in [0,1] & \text{if } \sigma_u = \sigma_v \\ 1 & \text{otherwise} \end{cases}$

uniqueness condition: $\beta > 1 - \frac{2}{\Delta}$

- hardcore model of max-degree Δ=O(1): $\lambda_v < \frac{1}{\sqrt{2}\Delta - 1}$

uniqueness condition: $\lambda < \frac{(\Delta-1)^{(\Delta-1)}}{(\Delta-2)^{\Delta}} \approx \frac{e}{\Delta - 2}$

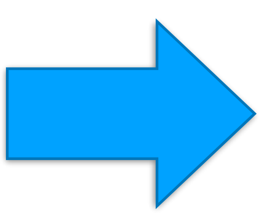

- # of iterations is $O(\log |D|)$ in expectation;
- total # of resamplings is $O(|D|)$ in expectation.

Dynamic Hardcore Sampler

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \mathsf{vbl}(D) \triangleq (V \cap D) \cup \left(\bigcup_{e \in D \cap E} e\right)$;
- while $R \neq \emptyset$:
  - each $v \in R$ with $X_v=1$ adds all neighbors to R;
  - each $v \in R$ resamples $X_v \in \{0,1\}$ ind. with prob. $\Pr[X_v=1]=\lambda_v/(1+\lambda_v)$;
  - $R \leftarrow \bigcup_{e=(u,v)\in E:\ X_u=X_v=1} e$;

in each iteration: $(X,R)\longrightarrow(X',R')$

$$\forall R \subseteq V: \quad \mathbb{E}[H(R') \mid R] \leq (1-\delta)H(R)$$

potential function $H(R) \triangleq |E(R)|$

Fast Convergence

- general graphical model with $k = \mathrm{O}(1)$ and $d = \mathrm{O}(1)$:

$$\boxed{\forall e \in E, \quad \phi_e : [q]^e \to [\beta_e, 1] \quad \text{and} \quad \beta_e > \sqrt{1 - \frac{1}{d+1}}}$$

- Ising model of max-degree Δ=O(1):

$$\boxed{\beta_e > 1 - \frac{1}{2.221\Delta + 1}}$$

(ferro-)

$$\phi_e(\sigma_u, \sigma_v) = \begin{cases} 1 & \text{if } \sigma_u = \sigma_v \\ \beta_e \in [0,1] & \text{otherwise} \end{cases}$$

(anti-ferro-)

$$\phi_e(\sigma_u, \sigma_v) = \begin{cases} \beta_e \in [0,1] & \text{if } \sigma_u = \sigma_v \\ 1 & \text{otherwise} \end{cases}$$

uniqueness condition: $\beta > 1 - \frac{2}{\Delta}$

- hardcore model of max-degree Δ=O(1):

$$\boxed{\lambda_v < \frac{1}{\sqrt{2}\Delta - 1}}$$

uniqueness condition: $\lambda < \frac{(\Delta - 1)^{(\Delta - 1)}}{(\Delta - 2)^{\Delta}} \approx \frac{\mathrm{e}}{\Delta - 2}$

- # of iterations is $\mathrm{O}(\log |D|)$ in expectation;
- total # of resamplings is $\mathrm{O}(|D|)$ in expectation.

Dynamic Sampling

Input: a graphical model with Gibbs distribution μ
a sample $X \sim \mu$, and an update (D, ϕ_D)

Output: $X' \sim \mu'$ where μ' is the new Gibbs distribution

Upon receiving update (D, ϕ_D):

- apply changes (D, ϕ_D) to the current graphical model;
- $R \leftarrow \mathsf{vbl}(D) \triangleq (V \cap D) \cup \left(\bigcup_{e \in D \cap E} e\right)$;
- while $R \neq \emptyset$:
  - $(X, R) \leftarrow \mathtt{Resample}(X, R)$;

$\mathtt{Resample}(X, R)$:

- each $e \in E^+(R)$ computes $\kappa_e = \min_{x_e :\, x_{e \cap R} = X_{e \cap R}} \phi_e(x_e)/\phi_e(X_e)$
- each $v \in R$ resamples $X_v \in [q]$ independently according to ϕ_v;
- each $e \in E^+(R)$ is passed independently with prob. $\kappa_e \cdot \phi_e(X_e)$;
  (otherwise e is violated)
- $R \leftarrow \bigcup_{e \in E:\ \text{violated}\ e} e$;

Summary

- A dynamic sampler for graphical models.
- The algorithm
  - is *Las Vegas*: good for simulation;
  - is *parallel* & *distributed*: good for systems;
  - can handle each local update in *constant* time.
- Equilibrium conditions for resampling.
- **Open problems**:
  - Dynamic sampler for colorings.
  - Better convergence regimes.
  - Extend to continuous variables & global constraints.

Thank you!