Tractable Marginal Inference for Hinge-loss Markov Random Fields

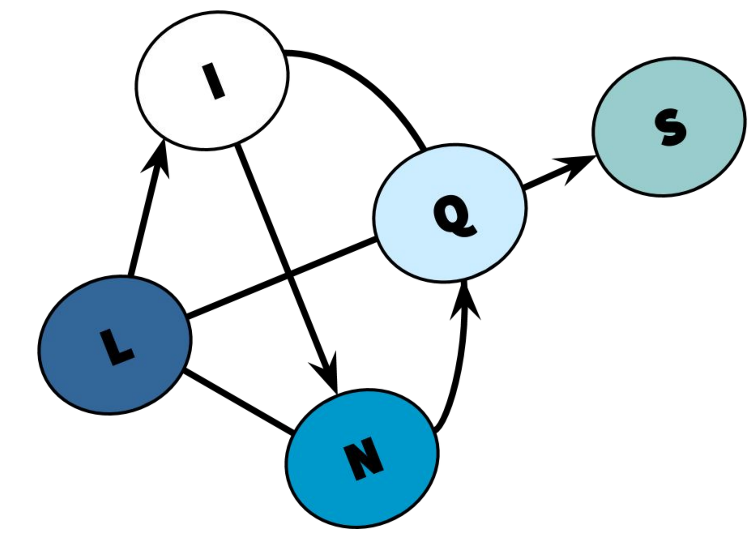

Varun Embar *, Sriram Srinivasan *, and Lise Getoor
University of California, Santa Cruz
*Equal Contribution

Goal

Compute the marginal distribution of a subset of random variables for HL-MRFs

$$p(Y \mid X) \xrightarrow[\text{Intractable}]{\int_{Y_{-c}}} p(Y_c \mid X)$$

Estimate empirically using Gibbs sampler

Challenges

- Hard to sample from conditional distributions
- Small islands of high probability

Contributions

- Metropolis-in-gibbs sampler for HL-MRFs
- Identify islands of high probability using "*association blocks*"
- Estimation of *relational properties* using marginals outperform MAP estimates by upto 60%

Hinge-loss Markov Random Field

- Class of continuous MRF used model richly structured data
- Supports efficient MAP

$$p(Y \mid X) = \frac{1}{Z} exp^{-\sum_{r=1}^{N} w_r \phi_r(Y,X)}$$

Probabilistic Soft Logic [1]

- Templating language for HL-MRFs
- Defined using weighted first-order logic rule

Rule weight → w: LivesTogether(X,Y) ⟶ Friends(X, Y)

\+ Data

Grounding ↓

w: LivesTogether(Alice, Bob) ⟶ Friends(Alice, Bob)

Łukasiewicz logic ↓

$$\phi(x, y) = max\{x - y, 0\}$$

[1] http://psl.linqs.org

Metropolis-in-gibbs sampler for HL-MRF

- Conditional distribution for gibbs sampler

$$p(y_i \mid X, Y_{-i}) \propto exp\{-\sum_{r=1}^{N_i} w_r \phi_r(y_i, X, Y_{-i})\}$$

Hard to sample from

- Single step of Metropolis sampler inside gibbs sampler

$$\alpha = \frac{exp\{-\sum_{r=1}^{N_i} w_r \phi_r(y_i', X, Y_{1:i-1}^{(t+1)}, Y_{i:n}^{(t)})\}}{exp\{-\sum_{r=1}^{N_i} w_r \phi_r(y_i, X, Y_{1:i-1}^{(t+1)}, Y_{i:n}^{(t)})\}}$$

Acceptance ratio

Update state:

with probability α

$y_i^{(t+1)} \sim U[0,1]$

with probability $1 - \alpha$

$y_i^{(t+1)} = y_i^{(t)}$

Association blocks

High rule weights → Correlated random variables → Islands of high probability → Slow convergence of sampler

Association blocks = Clusters of correlated RV

- Identify association blocks from rules using
  - Rule weights
  - Feasible region
- Block sample RV in associated blocks

In O(rules) time

Proposal distribution

For each RV y_i in the block
  With probability β
    $y_i \sim U[0,1]$
  With probability $1 - \beta$
    $y_i \sim$ feasible region of y_i
  Update feasible region for all unsampled RVs

Estimating network properties

Data: Synthetic social network - Small, Medium and Large

Unobserved RVs: Node party affiliations

Approaches: MAP, Mean and Expectation of Gibbs ($Gibbs_{Mean}$ & $Gibbs_{Exp}$); Blocked Gibbs ($ABGibbs_{Mean}$ & $ABGibbs_{Exp}$)

PSL Model:

10: Strong(A, P) ⟶ Party(A,P)
5: Weak(A, P) ⟶ Party(A,P)
5: Party(A, P) ∧ Friends(A,B) ⟶Party(A,P)
1000: Party(A, +P) = 1 ← High rule weight
1: Party(A, P) = 0.5

Approach	Party Affiliation	P1	P2
MAP	0.710	130	121
$Gibbs_{Mean}$	0.707	1001	339
$Gibbs_{Exp}$	0.702	1004	354
$ABGibbs_{Mean}$	0.859	280	105
$ABGibbs_{Mean}$	0.750	484	172
Ground Truth	N/A	595	187

Network properties for the large graph

Network properties:
P1: # of node pairs with different party affiliations
P2: # of nodes that have adjacent nodes affiliated with both parties

Inferred distribution: Gibbs sampler fails to recover the distribution due to correlated RVs

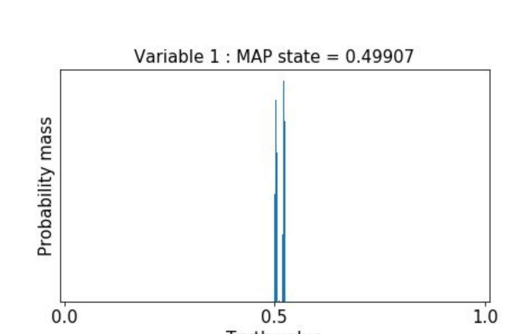

Gibbs

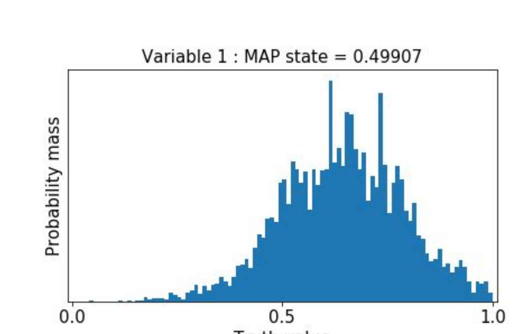

ABGibbs

Conclusion

- Proposed a novel sampling approach to compute the marginal distributions for HL-MRFs
- Using association block, we identify islands of high probability
- Network properties estimated using marginals computed using ABGibbs outperform other approaches by upto 60%

Supported by the NSF grants CCF-1740850 and IIS-1703331 and by AFRL and the DARPA.