

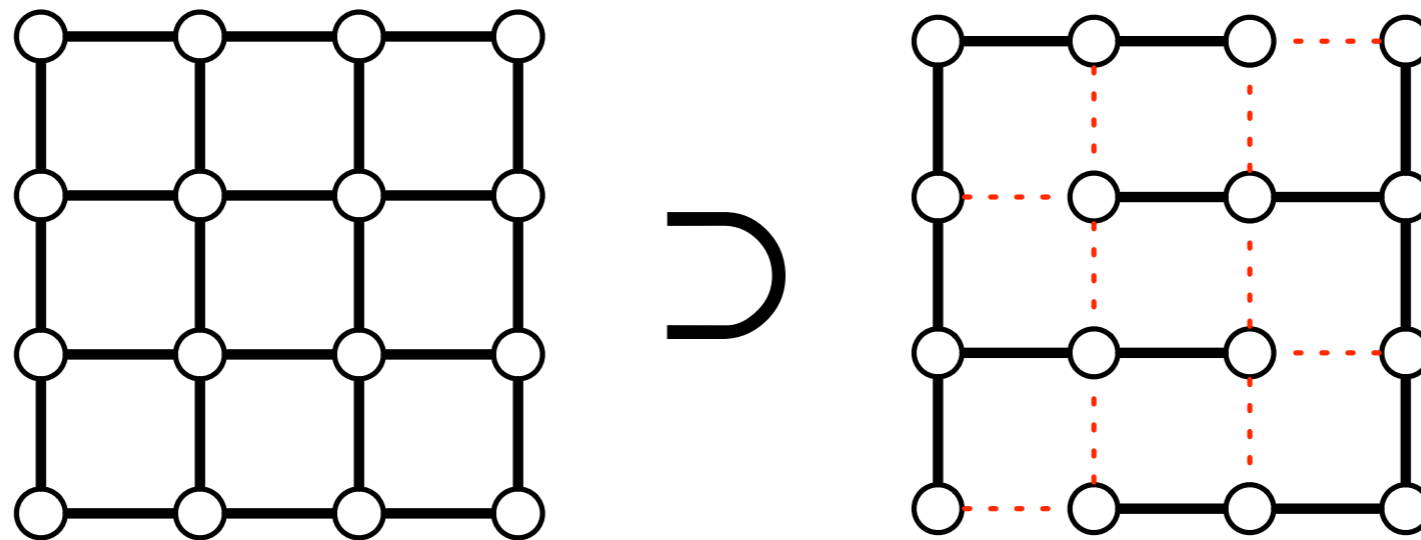
Optimization of Structured Mean Field Objectives

Alexandre Bouchard-Côté* Michael I. Jordan*,†

* Computer Science Division † Department of Statistics
University of California at Berkeley


Structured mean field

- A well-known method for doing approximate inference in intractable probabilistic models
- In Markov random fields, the approximation is usually based on an acyclic subgraph




Picking a subgraph

- Using more edges increases quality
- What is the impact on computational complexity?




$n = \#nodes$

Preview of our results



v -acyclic



b -acyclic

ν -acyclic


- Computationally easy
- Approximations in the literature fall into this category
- Connection with block Gibbs sampling

0.1	0.3			
0.9		0.8		
			0.4	0.9
	0.2		0.6	
	0.5	0.2		0.1

0	1			
1		1		
			0	1
	0		1	
	0	0		0

b-acyclic

- More accurate but computationally harder
- We improve on the direct method by using a technique based on *auxiliary exponential families*



$$\mathcal{O}(n^3) \rightarrow \mathcal{O}(n^2)$$