Exercise: Inference in 2D lattice MRFs

Consider an MRF with a 2D $m \times n$ lattice graph structure, so each hidden node, X_{ij}, is connected to its 4 nearest neighbors, as in an Ising model. In addition, each hidden node has its own local evidence, Y_{ij}. Assume all hidden nodes have $K > 2$ states. In general, exact inference in such models is intractable, because the maximum cliques of the corresponding triangulated graph have size $O(\max\{m, n\})$. Suppose $m \ll n$ i.e., the lattice is short and fat.

1. How can one *efficiently* perform exact inference (using a deterministic algorithm) in such models? (By exact inference, I mean computing marginal probabilities $P(X_{ij}|\vec{y})$ exactly, where $\vec{y}$ is all the evidence.) Give a *brief* description of your method.

2. What is the asymptotic complexity (running time) of your algorithm?

3. Now suppose the lattice is large and square, so $m = n$, but all hidden states are binary (ie $K = 2$). In this case, how can one efficiently exactly compute (using a deterministic algorithm) the MAP estimate $\arg\max_x P(x|y)$, where x is the joint assignment to all hidden nodes?