Assignment 6

We are interested in modeling the financial performance of companies across sectors. Each company's stock will either rise or fall in value in the next quarter. Each company has a particular sector $1 \dots K$ (agriculture, health, energy, etc). We observe M binary public attributes of each company (sales increasing/decreasing, public/private, etc).

We will use a naive-Bayes-like model which imagines that the public attributes depend on both stock and sector, and that assumes that the attributes are independent of one another given stock and sector. We will not assume any particular structure of this dependence (i.e. we will use totally different parameters for $P(attribute|rise, agriculture)$ and $P(attribute|fall, agriculture)$.

Notation:

- use S_i to represent sector of company i, ($k \in \{1, ..., K\}$).
- use R_i to represent rise or fall in company i's stock, where $R_i \in \{-1, +1\}$.
- use $A_{i,j}$ to represent public attribute j of company i, where $i \in \{1, ..., N\}$ and $j \in \{1, ..., M\}$.

1. Draw a BN, MRF and factor graph that represent this situation. Use plate notation. Propose a way to parameterize each distribution, and propose reasonable priors for each.
2. Assume that stock and sector are observed. Derive a posterior distribution for all parameters.
3. For the remainder of the problem, we will assume that we never observe the stock or sector. We will build an unsupervised model. Of course, our learned clusters may or may not correspond to the known types.
   (a) Explain why belief propagation would not give exact posterior estimates for this model.
   (b) Describe how to use Gibbs sampling for inference in this model. Derive the update probabilities.
   (c) Suppose we use the Metropolis-Hastings for inference, using a transition that chooses both R_i and S_i (simultaneously) uniformly at random. Derive the acceptance probability of this transition.
   (d) Describe how to use EM to learn the model parameters. Derive the updates.