CS 228T QUIZ 3

1. Briefly explain how to parameterize the multinomial distribution as a linear exponential family, and comment on the form of the sufficient statistics.

2. Discuss the differences between the M-projection and I-projection of a mixture of Gaussians onto the family of Gaussian distributions.

3. The exact inference problem involved maximizing the free energy functional over the marginal polytope. Meanwhile, the loopy BP problem involved approximating this problem by modifying both the objective and the feasible set: we optimized the factored free energy functional over the local consistency polytope, a superset of the marginal polytope.

(a) In mean field methods, do we approximate the exact objective, the marginal polytope, or both?
(b) What is the form of the feasible set in the naive mean field method?
(c) Is the feasible set in the mean field problem a subset or superset of the marginal polytope? Explain briefly.

4. Consider two approximating families $\mathcal{Q}_1$ and $\mathcal{Q}_2$ such that $\mathcal{Q}_1$ is strictly more expressive than $\mathcal{Q}_2$ (so that any distribution expressible in the latter family is expressive in the former, but not vice versa).

(a) What can we say about the relation between the optimal values of the optimization problem on page 449 for these two families?
(b) For these two families, what can we say about the results of the iterative algorithm that cycles through the fixed point equations?

Date: September 29, 2020.