00103335: Deep Learning and Reinforcement Learning
Homework 4 *Due*: December 21, 2022

1. Prove that an undercomplete autoencoder with linear decoder and MSE loss learns to span the principal subspace of the training data.
2. Derive (19.56) in the DL book for the mean field approximation.
3. Consider the Boltzmann machine with state vector $\boldsymbol{x} \in \{0,1\}^d$ and energy function

$$E(\boldsymbol{x}) = -\boldsymbol{x}^T \boldsymbol{U} \boldsymbol{x} - \boldsymbol{b}^T \boldsymbol{x}.$$

   (a) Derive the conditional distributions $p(x_i \mid x_{-i})$.
   (b) Do the conditional distributions in part (a) uniquely determine the joint distribution of $\boldsymbol{x}$? Prove or disprove.
4. Describe an algorithm for training a generative adversarial network and comment on its convergence properties.
5. RL Exercise 9.5
6. RL Exercise 10.6 (the two-state problem, which was Exercise 10.7 in the first printing).
7. RL Exercise 13.1
8. RL Exercise 13.5