

Deep Learning and Artificial Intelligence
WS 2018/19

Exercise 7: Representation and Distributional Learning

Exercise 7-1 Local and Distributed Representations

- (a) Consider the sequence s of words “*I like cats; cats like to catch mice*” and a set of symbols $D = \{\textit{cats, dogs, mice, like, hate, to, catch, '.'}\}$. Create a local representation of D .
- (b) Represent s as a sequence of vectors and a bag-of-symbols respectively.
- (c) What are the advantages of the bag-of-symbols representation compared to the sequence of vectors representation? What are the disadvantages?
- (d) Can local representation encode word semantics? How do distributed representations encode word semantics and relatedness between words?

Exercise 7-2 Autoencoders

Given a dataset $X \in \mathbb{R}^{N \times d}$ containing N samples of d features. We want to use an undercomplete autoencoder with one hidden layer of k units ($k \ll d$) in order to capture the most salient features of X . Let W and b_e denote the weights and bias of the encoder E and let U and b_d denote the weights and bias of the decoder D .


- (a) Assume that the encoder uses sigmoid as activation and the decoder is linear (i.e. the activation is the identity function). Write down the loss function (using mean squared error) that has to be minimized!
- (b) Assume that the bias terms are zero and the encoder is also linear. Write down the loss for this setting in a vectorized way.
- (c) It can be shown that the optimal weights of an autoencoder with one hidden layer of size k and no non-linearity leads to the same result as PCA, i.e. the weight vectors learned by the autoencoder span the same subspace as the k first principal components of PCA.¹ On the course website, you can find an ipython notebook that applies PCA to a small data matrix. Implement the missing parts of the method `autoencoder` and apply it to the same data. Compare the results to the ones of PCA.
- (d) Apply the same autoencoder to all MNIST images that contain one digit of your choice (e.g. 3). How could you improve this autoencoder?

Exercise 7-3 Restricted Boltzmann Machines

Restricted Boltzmann Machines (RBM) can be used to learn representations of inputs. The canonical RBM is an energy-based model with binary visible v_i and hidden h_i units².

¹For those who are interested: You can read more about this relationship in this paper

²Source and further reading: “Deep Learning” from Goodfellow et. al. chapter 16.7.1 and 20.



Its energy function is

$$E(v, h) = -b^T v - c^T h - v^T W h,$$

where b , c and W are unconstrained, real valued, learnable parameters. The joint probability distribution is specified by its energy function as

$$p(v, h) = \frac{1}{Z} e^{-E(v, h)},$$

where the normalizing constant Z is defined as

$$Z = \sum_v \sum_h e^{-E(v, h)}.$$

- Explain intuitively what the energy function tries to achieve.
- What is the conditional probability $P(h = \mathbb{1} | v)$ and $P(v = \mathbb{1} | h)$, whereas $\mathbb{1}$ is a vector whose values are all 1?
- Assume we have a restricted Boltzmann machine with the following parameters:

$$W = \begin{pmatrix} -0.384 & -2.295 & -0.872 & -1.219 \\ -1.258 & -2.219 & -0.154 & -1.581 \end{pmatrix}$$

$$b = \begin{pmatrix} 1.350 \\ 2.266 \\ 0.299 \\ 1.545 \end{pmatrix}$$

$$c = \begin{pmatrix} 3.077 \\ 3.205 \end{pmatrix}$$

Compute $P(h_j = 1 | v)$ for all $h_j \in h$ for the following v :

$$v = \begin{pmatrix} 1 \\ 1 \\ 0 \\ 1 \end{pmatrix}$$

- Use the same parameters as in the exercise before. Compute the probability $P(v_i = 1 | h)$ for each $v_i \in v$ for the following h :

$$h = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

- (Optional) You can take a look at the “RBM” Jupyter notebook at the course web-site. This notebook is generating a random dataset. This dataset will be used to train a RBM. Then we extract the parameters of the RBM and we compute $P(h = \mathbb{1} | v)$ and $P(v = \mathbb{1} | h)$.

Hint: You can use the notebook as generator and validator of paper and pencil exercises similar to the exercise above for your exam preparation.