

Artificial Neural Networks (Gerstner). Exercises for week 3

Statistical Approaches for Deep Networks

Exercise 1. Maximum Likelihood solution

You have observed P data points, $\{\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^P\}$ where $\mathbf{x}^\mu \in \mathcal{R}^N$.

You have reason to believe that the data set might come from a radially symmetric Gaussian distribution

$$p(\mathbf{x}|\mathbf{x}_c, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}^N} \exp\left[-\frac{(\mathbf{x} - \mathbf{x}_c)^\top (\mathbf{x} - \mathbf{x}_c)}{2\sigma^2}\right] \quad (1)$$

where \mathbf{x}_c is the center Gaussian and σ its width (standard deviation).

- In-class-exercise: Calculate the value of \mathbf{x}_c for which the likelihood that the data could have been generated by a Gaussian is maximal.
- At home: Repeat the same calculation for the width σ .

Exercise 2. Logistic/Sigmoidal unit

In class it was shown that the drive a for the sigmoidal unit in a single-class scenario is given by the log of the probability ratio $a = \ln[p(\mathbf{x}, C_1)/p(\mathbf{x}, \bar{C}_1)]$ where the bar indicates “not in class C_1 ”.

In some textbooks, however, it is argued that the drive a is rather $a = \ln[p(C_1|\mathbf{x})/p(\bar{C}_1|\mathbf{x})]$.

- Are the two expressions for a the same or not?
- Define $b = \ln[p(\mathbf{x}|C_1)/p(\mathbf{x}|\bar{C}_1)]$. Is $b = a$? Always true, never, or under some condition?

Exercise 3. Softmax function

In class we have defined the softmax output for unit k as

$$\hat{y}_k = \frac{\exp(a_k)}{\sum_j \exp(a_j)} \quad (2)$$

The following considerations are important in a practical implementation of the softmax function in order to avoid overflow events.

- Show that the output does not change if you shift, in all components, the activations by the same constant c . In other words, you may work with $a_j + c$ instead of a_j , for all j simultaneously.
- Why is it useful to choose $c = -\max_j[a_j]$?
- Check on your calculator that $\exp(3) \approx 20$. Your friend argues that, by subtracting the constant defined in part b, you can neglect all outputs $a_j - \max_j[a_j] < -10$. Is she correct?
- You have a network with $K = 4$ output units in layer n . We have a multi-class problem with 1-hot coding. For each input \mathbf{x}^μ only one of the $K = 4$ output units has target $t_i^\mu = 1$ and all other target values are zero: $t_k^\mu = 0$ for $k \neq i$. The cross-entropy error function is defined as

$$E(\mathbf{w}) = -\sum_{\mu} \sum_{k=1}^4 t_k^\mu \ln \hat{y}_k^\mu \quad (3)$$

where the sums run over all output units and all patterns. Assume that the output \hat{y}_k is given by the softmax function defined at the beginning of the exercise, with a drive $a_k = \sum_j w_{kj}^{(n)} x_j^{(n-1)}$. Perform stochastic gradient descent. To do so, consider one specific pattern, e.g., $\mu = 12$. Give the update formula for weight $w_{34}^{(n)}$ after the presentation of pattern $\mu = 12$. Hint: exploit that in the 1-hot coding $\sum_k t_k^\mu = 1$ for all μ .

- e. The target output for pattern $\mu = 12$ is $t_1^{12} = 1$ and zero for the other output components. The drives for the output units are $a_1 = 100$; $a_2 = 97$; $a_3 = 92$; $a_4 = 14$. Without using a calculator, which output unit has the largest absolute output error $|\delta| = |dE/da|$? Hint: use the results of (a) - (d).
- f. Show that for the case of two classes and 1-hot coding (1 out of K), the multi-class cross-entropy in (d) reduces to the single-class formula:

$$E(\mathbf{w}) = - \sum_{\mu} [t^\mu \ln \hat{y}^\mu + (1 - t^\mu) \ln (1 - \hat{y}^\mu)] \quad (4)$$

and with the softmax replaced by the sigmoidal.

Hint: Start with the softmax definition for $K = 2$ classes and exploit that $\hat{y}_1 + \hat{y}_2 = 1$.

Exercise 4. Softmax function for multi-class problems.

We have a network with K output units indexed by $1 \leq k \leq K$. We want that output unit k gives the posterior probability for class k


$$\hat{y}_k = P(C_k | \mathbf{x}) \quad (5)$$

Show that we arrive at the softmax function if we set the joint probabilities as $p(\mathbf{x}, C_k) = \frac{1}{Z} e^{a_k}$, where Z is a normalization constant.

Hints: Proceed analogously to the 'derivation' of the sigmoidal unit in class last week. Remember that the softmax function is

$$\hat{y}_k = \frac{\exp(a_k)}{\sum_j \exp(a_j)} \quad (6)$$

Exercise 5. Probability modeling with rectified linear units and threshold units



Please have a look at Figs A and B. Input data \mathbf{x} is two-dimensional. We consider a network with three hidden units and one output unit.

- a. We start with the network in Fig A. All hidden units and the output unit are strict threshold functions $x^{(n)} = \theta(a)$, where θ is the Heaviside function.

All positive data samples \mathbf{x}^μ are uniformly distributed within $x_1 < 3$ and $-1 < x_2 < 1$. All negative examples are outside these bounds. Similar to the construction that was done in class for the XOR problem, your task is to construct by hand a solution such that the output is +1 for all positive examples and zero for all negative examples.

Show that there are 48 different, but equivalent solutions for the weight vectors of the hidden units, if you assume three normalized weight vectors for the hidden units $(\mathbf{w}_1^{(1)}, \mathbf{w}_2^{(1)}, \mathbf{w}_3^{(1)})$; and infinitely many if you drop the normalization condition. Pick one of your solutions and put the weight values for all weights (both layers) on the graphic in A.

Hint: the dashed lines show the three linear separations that you will need for your three hidden units.

- b. In Fig B, all hidden units are rectified linear units. The output unit is the sigmoidal function $g(a) = 1/[1 + \exp(-a)]$. For the sake of numerical simplicity, you should use the following approximation: $g(-3) = 0.05$ and $g(+3) = 0.95$.

- (i) Within the bounds $x_1 < 3$ and $-1 < x_2 < 1$ the probability $p(C_1|\mathbf{x})$ that an example \mathbf{x} is indeed a positive data sample is equal or above 95 percent (and exactly 95 percent on the border, see Fig B).

The probability $p(C_1|\mathbf{x})$ to find a positive example is 0.05 for

- (ii) any input example on the line $x_1 \leq 3$ and $x_2 = 1.5$;
 (iii) any input example on the line $x_1 \leq 3$ and $x_2 = -2.0$;
 (iv) any input example on the line $x_1 = 6$ and $-1 < x_2 < +1$.

Construct a solution with piecewise linear units in the hidden layer and sigmoidal in the output layer, such that your network matches to probabilities summarized under (i) - (iv). Write the weight values for your solution directly into the graph.

Hint: There are many solutions. Pick one of these. Take one of the solutions where weight values have simple numbers. You may find it is easier to assume that in the shaded area the probability is uniformly at 0.95. Always assume that $g(-3) = 0.05$ and $g(+3) = 0.95$.

- c. Optional small programming check:

Generate 10 000 positive and negative training data points with the rules described in (b). Take your calculated weight values as an initial condition for the weights and run a gradient descent optimizer to see whether it finds even better values.