TU München
Fakultät für Informatik
PD Dr. Rudolph Triebel
John Chiotellis

Machine Learning for Computer Vision
Winter term 2016

November 28, 2016
Deep Learning, Boosting

Neural Networks and Deep Learning

Please read `http://neuralnetworksanddeeplearning.com/chap1.html` and `chap2.html`. You don't have to implement the network yourself nor do the exercises. Please first complete the reading above. We are going to follow Michael Nielsen's notation.

Exercise 1: Back Propagation

- Suppose we modify a single neuron in a feedforward network so that the output from the neuron is given by $f(\sum_j w_j x_j + b)$, where f is some function other than the sigmoid.
  How should we modify the backpropagation algorithm (from chapter 2 of the above reading) in this case?

- Compute the gradient of the cost function C respect to w_5 $\left(\frac{\partial C}{\partial w_5}\right)$ given the following network:

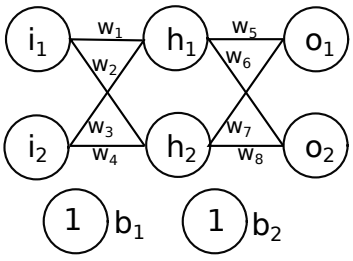

  $C = \sum_{i=1}^{2}(t_i - a_i)^2$ where t_i is the target value for the respective output neuron o_i and a_i is the output of the neuron o_i. Input/output of h_1, h_2, o_1, o_2 is computed as $z_i^l = \sum_j w_j \cdot a_j + b_l$, $a_i^l = \sigma(z_i^l)$ where $\sigma(\cdot)$ is any activation function.

Exercise 2: Convolutional Layer Arithmetic

Consider a very simple convolutional neural network that just consists of one convolutional layer. It has the following parameters:

- number of kernels: $num = 64$
- size of kernels: $k = 3 \times 5$
- stride: $s = 2$
- padding: $p = 1$

Assume, the input to this layer is an a batch of RGB images. There are 10 images in one batch and the images have a dimension of 123×81.

a) What is the shape of the input blob to the convolutional layer? Hint: it's a tensor with four axes.

b) What is the shape of the output blob of the convolutional layer?

Boosting

Exercise 3: Adaboost (Programming)

Download the file 'banknote_auth.zip' available at the course's website. The data are features of banknotes and the labels indicate whether a banknote is forged or not. The dataset is taken from `https://archive.ics.uci.edu/ml/datasets/banknote+authentication` with some duplicate entries removed. Implement the AdaBoost algorithm with decision stumps as weak classifiers.

a) To begin train on 50% of the data with 20 weak classifiers.

b) Generate a plot of the classification error with respect to the number of weak classifiers. What do you observe?

c) Add more weak classifiers. Does the error still change? What's the optimal number of weak classifiers to use?

d) Now keep the number of weak classifiers fixed and try different training/testing set sizes. How does it affect the classification accuracy?

The next exercise class will take place on **December 9th, 2016**.

For downloads of slides and of homework assignments and for further information on the course see

`https://vision.in.tum.de/teaching/ws2016/mlcv16`
