Information Mining - winter semester 2019

Exercise sheet 10

Exercise 1: Neuron architecture

What happens within a single neuron?

Exercise 2: Activation functions

Mention 2 activation functions by name, give detailed formula and draw a plot presenting the output ranges. What is the motivation of using activation functions?

Exercise 3: RNN

What is the motivation of RNNs?

Exercise 4: LSTM

What is the advantage of LSTM over simple RNNs?

Exercise 5: One-hot-Encoding

What do you understand under one-hote-encoding. Give an example.

Exercise 6: Embedding

What is a sparse and what is a dense vector? Can you motivate the use of embeddings.

Exercise 7: Cross-Entropy

Suppose a classifier predicts each possible class with equal probability. If there are 10 classes, what will the cross-entropy error be on a single example?

- - log(10)
- - 0.1 log(1)
- - log(0.1)
- - 10 log(0.1)

Exercise 8: Understanding questions

- Suppose that you are training a neural network for classification, but you notice that the training loss is much lower than the validation loss. Which of the following can be used to address the issue (you can select more than one solution)?
  - Reduce the size of hidden layers.
  - Decrease dropout probability.
  - Increase L2 regularization weight.
  - Increase the size of each hidden layer, i.e. add more neurons in each hidden layer.
  - Add more training data.
- Deep Learning is the best solution for every problem regardless whether each problem has enough training data or not. True or False? Why?
- You observe that your model predicts very positive sentiment for the following passage: *Yesterday turned out to be a terrible day. I overslept my alarm clock, and to make matters worse, my dog ate my homework. At least my dog seems happy.* Why might the model misclassify the appropriate sentiment for this sentence?