

Homework 5

All assignments need to be submitted via github classroom:

<https://classroom.github.com/g/F65AmlrV>

and on gradescope. You can submit in groups of maximum size two.

The homework is due 05/01/19 at 1pm.

All the tasks are to be completed using the [keras Sequential interface](#). It's recommended that you run your code on GPU using google colab:

<https://colab.research.google.com/>

You can enable GPU support at "runtime" -> "change runtime type".

Feel free to experiment with TPU support if you're adventurous.

You can use any other resources at your disposal if you prefer. You should not use k-fold cross-validation for any of these tasks. You can use StratifiedShuffleSplit to create a single train-validation split for use with GridSearchCV.

Use of GridSearchCV might not be the best option for any task but task1, though.

Task 1 [10 Points]

Run a multilayer perceptron (feed forward neural network) with two hidden layers and rectified linear nonlinearities on the iris dataset using the keras [Sequential interface](#). Include code for selecting regularization strength and number of hidden units using GridSearchCV and evaluation on an independent test-set.

Task 2 [40 Points]

Train a multilayer perceptron (fully connected) on the Fashion MNIST dataset using the traditional train/test split as given by `fashion_mnist.load_data` in keras. Use a separate 10000 samples (from the training set) for model selection and to compute learning curves (accuracy vs epochs, not vs `n_samples`). Compare a "vanilla" model with a model using drop-out (potentially a bigger model), and to a model using batch normalization and residual connections (but not dropout). Visualize learning curves for all models.

Task 3 [60 Points]

Train a convolutional neural network on the following dataset:

<https://www.kaggle.com/paultimothymooney/breast-histopathology-images>

3.1 Start with a model without residual connections (using batch normalization is likely to be helpful and you should try it, whether you use dropout is your choice).

3.2 Augment the data using rotations, mirroring and possibly other transformations. How much can you improve your original model by data augmentation?

3.3 Build a deeper model using residual connections. Show that you can build a deep model that would not be able to learn if you remove the residual connections (i.e. compare a deep model with and without residual connections while the rest of the architecture is constant).

Hint: Make sure you are doing the reshape for the training set correctly. A direct reshape might give you garbled images. Display an image after reshaping to make sure they are correct.

Some additional advice to help you along:

- Make sure all your code is running on GPU. Use `"sess = tf.Session(config=tf.ConfigProto(log_device_placement=True))"` to start the tensorflow session to see which device is being used and confirm it is the device you intended.
- Preprocess the images before training a model.
- Test your code on a small part of the data before training the model. You don't want your code to fail on a print statement after waiting for the network to train.