

ELL409: Assignment 3

Maximum Marks: 8 (+3 Extra credit)

Submission deadline: **11 November, 23:59**

1 Objective

To experiment with the use of Neural Networks for a multiclass classification problem, and try and interpret the high-level or hidden representations learnt by it. Also, to try and understand the effects of various parameter choices such as the number of hidden layers, the number of hidden neurons, and the learning rate.

2 Data

For the previous assignment, you were provided a low-dimensional (PCA) representation of a data set of images. We now provide the corresponding original data: a personalised file for each of you, that contains 3,000 images, each of size 28×28 pixels. You should download your file from <http://web.iitd.ac.in/~sumeet/A3/<EntryID>.csv> (for example, <http://web.iitd.ac.in/~sumeet/A3/2014EE10421.csv>). Each row of this file corresponds to an image, and contains 785 comma-separated values. The first 784 are the grayscale pixel values (ordering is column-major), and the last one gives the class label for the image (there are 10 classes, denoted by the labels 0 to 9).

3 Methodology

Your task is to try and learn a Neural Network classifier for these images, starting with the raw pixels as input features, and thereby also to assess the usefulness of the different representations that your Neural Network constructs. Here is how you should proceed:

(a) Write your own implementation of a basic (fully connected) neural network for multiclass classification of the provided images: you essentially need to implement the backpropagation algorithm to be able to train the network weights. Your implementation should allow for some flexibility (the more the better) in setting the hyperparameters and modelling choices: number of hidden layers (try with at least 1 and 2), number of units in each layer, gradient descent parameters (learning rate, batch size, convergence criterion), choice of activation function in the hidden units, mode of regularisation. (2 marks)

In addition, familiarise yourself with one Neural Network library of your choice. One suggestion is PyBrain for Python, but you can find many others. You may wish to play with a simple toy data set to get a feel for using the library, before you move on to the actual data for this assignment.

(b) *Standard backpropagation neural net:* Use your implementation to train a neural network to recognise the images of handwritten digits given to you. Set aside some of the data for validation, or ideally, use cross-validation. Assess the accuracy and speed (both training and testing) of the neural net for different settings of the various hyperparameters mentioned above. Identify cases of overfitting or underfitting; use regularisation to get better results, if you think it will help. Once you have obtained a good model, try to visualise and interpret the representations being learnt by the hidden neurons. Can you make sense of them? Also, take a look at the images which are being misclassified by the network. Can you see what's going wrong? In addition, try using the standard library, instead of your own implementation, just to train the final model with your optimised hyperparameters. Does this alter the results in any way? If so, why

might that be?

(3 marks)

(c) *Comparison with PCA features*: Now consider the PCA-space representation of your data that you were provided for the previous assignment. This was a way of mapping the images to a lower-dimensional space, something that the neural net is also doing via its hidden units. Try to interpret and compare these two representations. Is the neural net in any sense able to learn a better representation than the PCA one? Train another neural network, using the PCA features from last time as the input features instead of raw pixels. First try with *no hidden layers*, i.e., a simple logistic regression model. Now add a hidden layer. Does it help? Why or why not? And how do these results compare with those obtained using just the raw pixels? (3 marks)

(d) *Advanced neural networks (for extra credit)*: This part is more open-ended. You could experiment with one of the two deep learning approaches discussed in class: convolutional neural nets, or sparse autoencoders (using standard libraries). The network you train need not be very deep; 2 or 3 layers is fine. The objective is primarily to see if these approaches can learn more useful representations than a standard neural net as employed above; and how these representations vary as you change the model hyperparameters. In order to have more training data, you may make use of the full MNIST data set, available at <http://yann.lecun.com/exdb/mnist/>. That page also provides a list of benchmark results with different kinds of techniques; you should see how close you can get to those benchmarks. More importantly, you should attempt to interpret the representations being learnt by your deep net(s). Are these in some sense more natural or intuitive than the representations learnt by the standard neural net you trained earlier? How do they vary with hyperparameter changes, and why? (3 marks)

(Please note that there are standard online tools that can do these tasks for you, and many people will have reported example results on MNIST data. You may use such tools for training your model; but all results, graphs/visualisations, and interpretation in your report must be your own. Any copying of these from elsewhere will constitute plagiarism.)

4 Evaluation

- You should prepare a report, compiling all your results and your interpretation of them, along with your overall conclusions. In particular, you should attempt to answer all of the questions posed in the previous section. Any graphs or other visualisations should also be included therein, as well as your code and any other materials which are relevant. The submission link, as well as precise instructions for how to organise and name your files, will be shared later. The submission deadline is **November 11th, 23:59**. Any late submissions will be penalised.
- The schedule for demos/vivas will be announced by your respective TAs, in advance. If for any reason you cannot attend in your scheduled slot, you must arrange for an alternative slot with your TA well in advance. Last-minute requests for rescheduling will normally not be accepted.