EEL709: Assignment 3

Weightage: 10%

March 23, 2015

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 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 representation of a data set of images. We now provide the corresponding original data: a personalised file for each of you, that contains 2500 images, each of size 28×28 pixels. You should download your file from `http://web.iitd.ac.in/~sumeet/A3/<EntryID>.mat` (for example, `http://web.iitd.ac.in/~sumeet/A3/2010EE50541.mat`). The file format is compatible with both MATLAB and Python; in case you have any difficulties with it and need the data in a different format, let us know. This file will contain two objects: `data_image`, which is a 784×2500 matrix, with each row corresponding to a pixel position (ordering is column-major) and each column corresponding to an image; and `label`, which is a 1×2500 vector, giving the class label for each 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:

1. Familiarise yourself with a Neural Network library of your choice. One suggestion is PyBrain (`http://pybrain.org/`) for Python, but you can find many others. Figure out how you can set various parameters, such as the number of hidden layers, the number of hidden neurons in each layer, the non-linear activation function to be used, and the learning rate for gradient descent. 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.

2. *Standard backpropagation neural net:* Attempt 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 parameters 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?

3. *Comparison with DCT features:* Comment on how your results compare to those obtained in the previous assignment. Is the neural net in any sense able to learn a better representation than the one

that was provided to you last time, which consisted of the top 15 coefficients from a discrete cosine transform of the image? Train another neural network, using the 15 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?

4. *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. 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. 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 respresentations being learnt by your deep net. Are these in some sense more natural or intuitive than the representations learnt by the standard neural net you trained earlier?

You can earn a maximum of 4 extra marks, in addition to the 10 allocated for this assignment, depending on the quantum of extra stuff done and the understanding obtained from it.

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. If you wish, you may also include code or other materials which are relevant, though this is not required. The submission link is `http://web.iitd.ac.in/~sumeet/submit.html`; put everything into a single zip file or tarball, and name it as per the instructions given there. The submission deadline is **April 6th, 23:59**. Any late submissions will be penalised.

- The schedule for demos/vivas will be announced by your respective TAs, in advance (planned dates are 7th and 8th April). 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.