Homework Assignment # 2

DUE: 5:00pm, Sunday March 4th (Electronic turnin required)

In this assignment, you are to experiment with using neural networks for a recognition task. Some freedom will be allowed here. In particular, you can choose which data sets and which neural networks to experiment with.

The important caveat here is that you should limit your ambition and computational scope (unless you have access to powerful GPUs for your experimentation). If you use CSIL, do remember that Tensorflow is available only in the CPU version. As the ECI technical support had so much trouble installing Tensorflow, I did not trouble them to install other packages such as PyTorch or MxNet. You might be able to get some of these packages to work on your own personal computer.

Furthermore, most "entry-level" public datasets have been experimented with extensively. Often times, it is nontrivial for you to beat published benchmark recognition rates that can be higher than 99%. However, hitting such high recognition rates is not the goal of this homework. You might be able to achieve a high recognition rate by adapting some published networks or starting the training from some published parameter sets and fine tuning further. But you will certainly learn more by rolling your own codes and performing training from scratch.

In terms of data sets, we recommend the following two:

1. **Mnist** http://yann.lecun.com/exdb/mnist/: hand written digits, with 60,000 training samples and 10,000 test samples. These images are 28x28 binary images with proper formatting and preprocessing already done.

2. **CiFar10** https://www.cs.toronto.edu/~kriz/cifar.html: 10 classes (airplane, auto, bird, cat, deer, dog, frog, horse, ship, truck) with 60,000 32x32 color images (6,000 per class) separated into 50,000 training and 10,000 test.

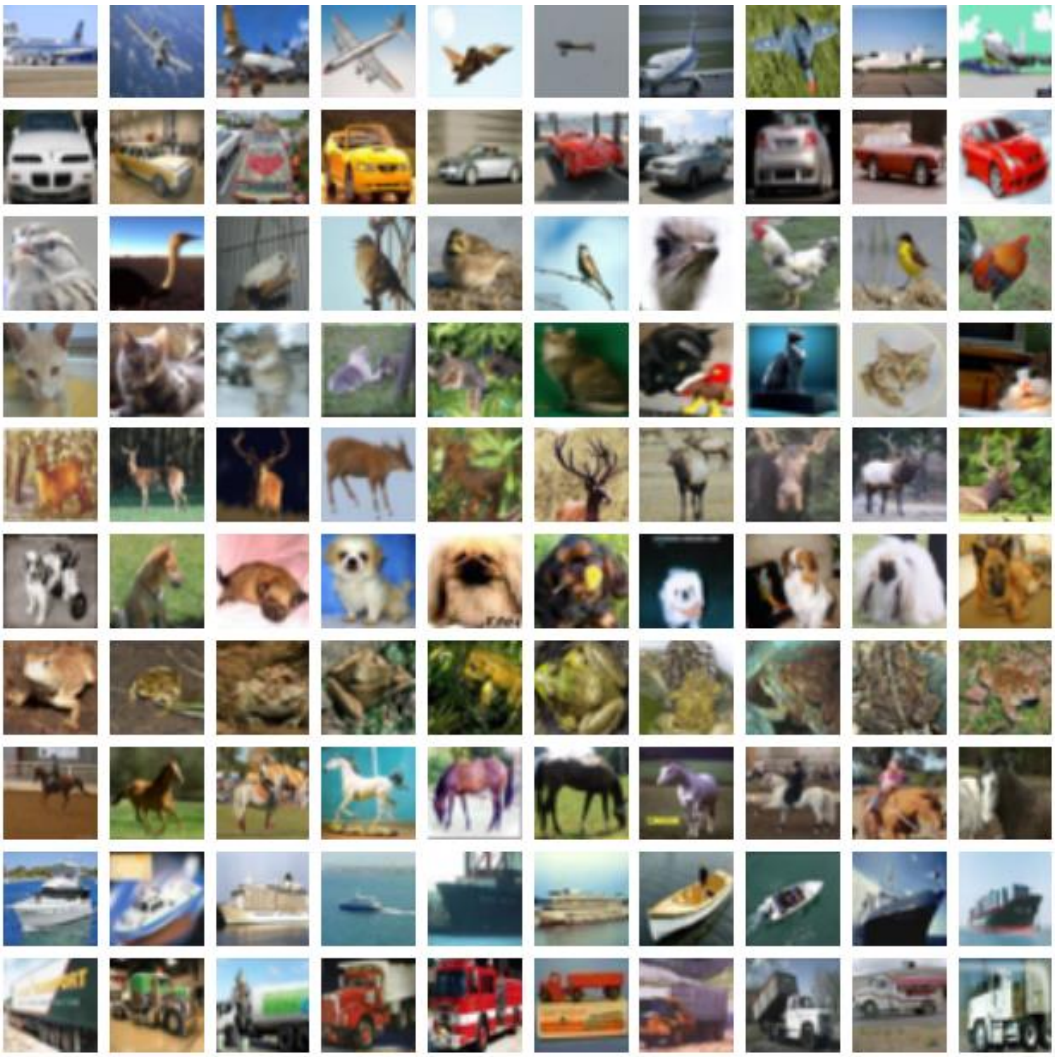

These are “reasonably” small data sets that hopefully won’t overwhelm your CPU, GPU and disk storage. They are normalized, cropped, preprocessed and 1-hot data sets, so you do not need to perform trimming, segmentation and region proposals. If you use other datasets, please make sure that they have such characteristics to simply your task. Processing temporal videos or images requiring elaborate segmentation with region proposals are not recommended for a two-week assignment, but you are welcome to try those for your class project.

As far as networks are concerned, you can use FNN, CNN, Resnet, or Densenet with your choice of number of blocks, number of layers, number of neurons per layer etc. Again, it is important to limit your expectation. Large networks, such as 100-layer Desnset, can consume large storage space and run time.

Your tasks are to write codes that have the following components (1) process images (input, batch, format conversion, etc.), (2) perform training (with proper batching and validating), and (3) estimate performance (e.g., error rate vs. run time and iterations). You should turn in your codes and a short report (PDF) which describes the data sets you use (if not the above two), how you process images and perform training, and test accuracy, runtime, resource usage statistics etc.