

A1

Deep Learning 2015, Spring

In this first assignment we are introducing Torch the Machine Learning platform used throughout this course.

"Torch7 is a scientific computing framework with wide support for machine learning algorithms. It is easy to use and provides a very efficient implementation, thanks to an easy and fast scripting language, LuaJIT, and an underlying C implementation." - <http://torch.ch/>
It has found widespread use in the industry and is being developed at Facebook amongst others. You can find some of the extensions, that have been developed at Facebook here
<https://github.com/facebook/fbcunn>
<https://github.com/facebook/fbnn>

Feel free to use these, if you want, but please keep in mind, that your deliverable must be executable on CIMS. Next, let us introduce your task.

Task

Your first task is to run and understand the starter code. The starter code is based on Clement Farabet's tutorial on supervised learning, which can be found here http://code.madbits.com/wiki/doku.php?id=tutorial_supervised. It is highly recommended, that you work through the tutorial. Our starter code only differs in that it doesn't make use of gfx. You can already run our code on CIMS like this

```
cp1995@crunchy1[your_folder]$ wget http://cs.nyu.edu/~cp1995/A1.tar.gz
cp1995@crunchy1[your_folder]$ tar -xvf A1.tar.gz
cp1995@crunchy1[your_folder]$ module load LuaJIT
cp1995@crunchy1[your_folder]$ cd A1
cp1995@crunchy1[A1]$ th doall.lua
```

This gives you a trained model and first scores on the test data. Now, you need to figure out how to save out these predictions and submit them to our Kaggle competition. After that, you need to figure out how to improve on our code. The predictions need to be made in order of and corresponding to the test data. Please keep in mind, as with any research in machine learning, training on the test data is equivalent to plagiarism.

Once you have your final trained model, create a stand alone script called "result.lua", that produces a file "predictions.csv" when run like this on any CIMS machine

```
module load LuaJIT
th result.lua
```

We will check, that "predictions.csv" matches your Kaggle submission. Your submission is likely to involve a large Model.net file, so please avoid attaching it to the email, if possible, and expose the file through CIMS. You can find instructions on how to host files on CIMS here <http://cims.nyu.edu/webapps/content/systems/userservices/webhosting>. If your code does not work on CIMS, we cannot award you any points.

Evaluation

- 30% - Kaggle performance
- 40% - Two page paper (brevity preferred)
- 30% - Simple, readable, commented code of final, working algorithm able to execute on binary test data

`test_32x32.t7`

as downloaded by doall.lua

The paper should consist of a

- description of the architecture (number and type of layers, number of neurons, size of input)
- description of the learning techniques applied (which data augmentations?, used dropout?, etc.)
- description of the training procedure (learning rate, momentum, error metrics used, train/validation split, training/validation/test error)

Please send your submission to cpuhrsch@nyu.edu Feb. 09 at 7:45pm.

Please use the following title for your email.

[Deep Learning YOUR_TEAM_NAME] Submission A1

Since this is the first assignment you are mostly likely still forming a team.

Please send me your team name, team members and team leader per email with following subject

[Deep Learning YOUR_TEAM_NAME] Introducing YOUR_TEAM_NAME

with content

YOUR_TEAM_NAME

YOUR_TEAM_LEADER - name, nyu email address

YOUR_TEAM_MEMBER_1 - name, nyu email address

YOUR_TEAM_MEMBER_2 - name, nyu email address

As a team leader you will be responsible for managing your team's submissions on Kaggle, submitting your team's assignment and communicating with the teaching staff.

Resources

To give you some idea about the data, please consult this website <http://ufldl.stanford.edu/housenumbers/>. We will give more directions to look at in the first tutorial.

Also view this

<https://github.com/torch/torch7/wiki/Cheatsheet>

<https://github.com/torch/torch7> - Read the docs on the bottom of the page, especially the docs on Tensor

http://code.cogbits.com/tutorials/1_supervised/README.html

Tips

Use the command line options of the starter code to easily use a small model / small data to explore the data or try different hyperparameters.

Read all the code.