Introduction to tensorflow

Jacob Bergstedt

Department of Automatic Control, Lund Institute of Technology, Lund
International Data Analysis Group, Pasteur Institute, Paris

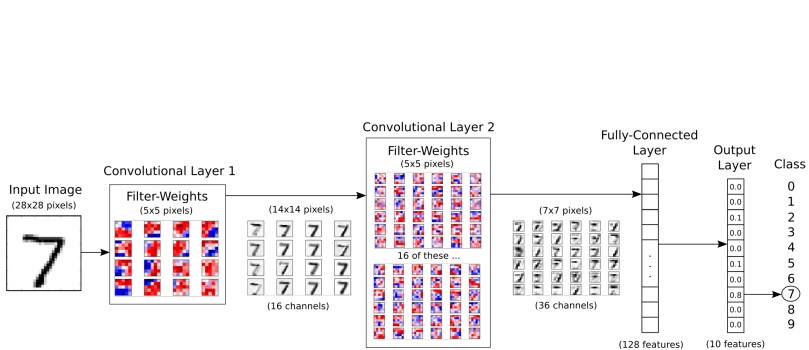*jacoba@control.lth.se*

November 8, 2016

Machine learning in Python

- Data wrangling: Pandas (recommended: R see tidyverse)
- scikit-learn
- XGBoost
- Tensorflow

Tensorflow is

- A modern computation engine
- Designed to be easy to roll out on various platforms
- Designed for large scale applications

- Specify inference network
- Add gradient
- Add summaries
- Run training loop

Input Image
(28x28 pixels)

Convolutional Layer 1

Filter-Weights
(5x5 pixels)

(14x14 pixels)

(16 channels)

Convolutional Layer 2

Filter-Weights
(5x5 pixels)

16 of these ...

(7x7 pixels)

(36 channels)

Fully-Connected
Layer

Output
Layer

Class

0
1
2
3
4
5
6
7
8
9

0.0
0.0
0.1
0.0
0.0
0.1
0.0
0.8
0.0
0.0

(128 features)

(10 features)

Suggestion for today

- Try to improve performance. Look at tensorboard to see what's going on
- Add accuracy summary (you may need to define a new writer)

Discussion

- Tensorflow example of further specialization between computation and application
- Is that development good? Is it reflected in teaching?

Home exercise

- Do the CIFAR-10 tutorial (takes a long time)
- Add summaries, including image summaries for weights and inputs

```
https://www.tensorflow.org/versions/r0.11/tutorials/
deep_cnn/index.html#convolutional-neural-networks
```