

Deep Learning Tutorial

Introduction to Torch7 and Theano

Current tooling landscape

- 2015: Explosion in frameworks: [Theano](#), [Torch7](#), [MXnet](#), [TensorFlow](#), [Chainer](#), [Nervana/Neon](#), [CNTK](#), [brainstorm](#), [SINGA](#), [Deeplearning4j](#), [Caffe](#), ...
- Most have (multiple) higher-level libraries on top.
- Two currently most popular, two paradigms:
 - Theano: Symbolic graph.
 - Torch7: Direct, numeric layers. (but see nngraph)
- This session: get a quick feeling for paradigms.
 - i.e. Not an in-depth tutorial!




Torch7

- Long history ranging back to 2002, multiple rewrites.
 - Read: a ton of experience flowed into it.
- Big-name users: Google (DeepMind), Facebook, Twitter.
- Written in LuaJIT and C/CUDA.
- Conceptually just like current exercise.
 - Hence, shorter introduction.

Resources

- [Getting started](#).
- Soumith's [CVPR15 notebook](#).
- "Official" [tutorial series](#).
- [Cheat-sheet](#) wiki with many links.



Theano

- Started in 2007 by LISA Lab@Montréal.
- Framework for symbolic computing with tensors.
 - Conceptually very different!
- Python!
- Used by many deep-learning labs.



Resources

- Official [Theano tutorial](#).
- Official [Deep Learning with Theano tutorial series](#).
- [DLSS'15 tutorial series](#).
- Kazuya's [notebook](#).