

Caffe

Caffe is a deep learning framework made with expression, speed, and modularity in mind.

It is developed by the Berkeley Vision and Learning Center (BVLC) and by community contributors.

Caffe is released under the BSD 2-Clause license.

Why Caffe?

- Speed makes Caffe perfect for research experiments and industry deployment. Caffe can process over 60M images per day with a single NVIDIA K40 GPU*. That's 1 ms/image for inference and 4 ms/image for learning. We believe that Caffe is the fastest convnet implementation available.
- Community: Caffe already powers academic research projects, startup prototypes, and even large-scale industrial applications in vision, speech, and multimedia. Join our community of brewers on the [caffe-users](#) group and Github.

Documentation

- DIY Deep Learning for Vision with Caffe
- Tutorial presentation.
- Tutorial Documentation
- Practical guide and framework reference.
- Model Zoo
- BVLC suggests a standard distribution format for Caffe models, and provides trained models.
- Developing & Contributing
- Guidelines for development and contributing to Caffe.
- API Documentation
- Developer documentation automatically generated from code comments.

TUTORIAL ON DEEP LEARNING FOR VISION

- A tutorial in conjunction with the Intl. Conference in Computer Vision (CVPR) 2014.
- Monday June 23, 2014
- Grand Ballroom 2
- Columbus, Ohio
- Schedule
- Morning Session: foundations
- 8.30–9.00 Introduction – Honglak Lee (University of Michigan)
- 9.00–10.00 Supervised learning – Marc'Aurelio Ranzato (Facebook A.I. Research)
- 10.00–10.30 Coffee Break
- 10.30–11.30 Unsupervised learning – Graham Taylor (University of Guelph)
- 11.30–12.30 Practical tools
- Torch7 – Marc'Aurelio Ranzato (Facebook A.I. Research)
- Theano/Pylearn2 – Presented by Ian Goodfellow (Univ. of Montreal)
- Caffe – Presented by Yangqing Jia (Google Research). [additional material]

Examples

- ImageNet Tutorial
 - Train and test "CaffeNet" on ImageNet data.
- LeNet MNIST Tutorial
 - Train and test "LeNet" on the MNIST handwritten digit data.
- CIFAR-10 tutorial
 - Train and test Caffe on CIFAR-10 data.
- Fine-tuning for style recognition
 - Fine-tune the ImageNet-trained CaffeNet on the "Flickr Style" dataset.
- CaffeNet C++ Classification example
 - A simple example performing image classification using the low-level C++ API.
- Feature extraction with Caffe C++ code.
 - Extract CaffeNet / AlexNet features using the Caffe utility.
- Web demo
 - Image classification demo running as a Flask web server.
- Siamese Network Tutorial
 - Train and test a siamese network on MNIST data.
- Finetuning oxford 102