

Caffe:

Convolutional Architecture for Fast Feature Embedding

Created by Yangqing Jia

Developed by BVLC

caffe.berkeleyvision.org

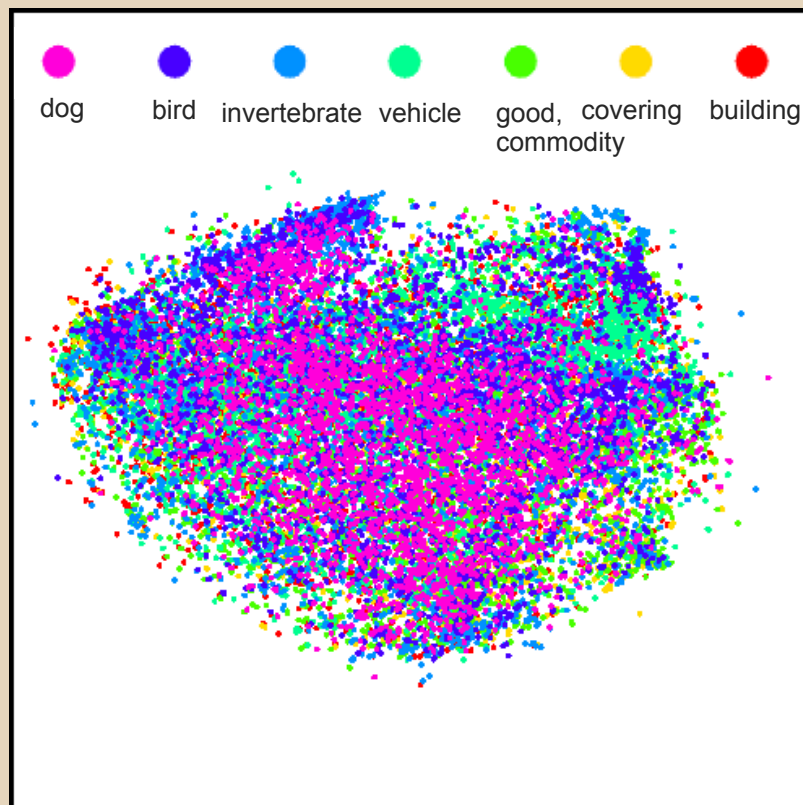
bvlc.eecs.berkeley.edu




Why Caffe?

The Unreasonable Effectiveness of Deep Features

Feature Embedding: learning to map visual features into a discriminative space.



Low-level: Pool₁



High-level: FC₆

Note that classes separate at higher layers of the representation!

So what is Caffe?

Convolutional Architecture for Fast Feature Embedding

- C++/CUDA framework for deep learning and vision
 - library of layers that compose into models
 - fast stochastic gradient descent (SGD) solver
 - tools, demos, and recipes
- Seamless switch between CPU and GPU
 - `Caffe::set_mode(Caffe::CPU);`
- BSD-2 licensed

Train Models



Experiment/Prototype



Inference at Scale



All with essentially the same code!

So what is Caffe?

Convolutional Architecture for Fast Feature Embedding

- Model schemas
 - Define the model and solving strategy and let Caffe take care of the rest
 - Adapt already learned models to new problems in one step

State-of-the-art solving in 14 lines.

```
1 train_net: "imagenet_train.prototxt"
2 test_net: "imagenet_val.prototxt"
3 test_iter: 1000
4 test_interval: 1000
5 base_lr: 0.01
6 lr_policy: "step"
7 gamma: 0.1
8 stepsize: 100000
9 display: 20
10 max_iter: 450000
11 momentum: 0.9
12 weight_decay: 0.0005
13 snapshot: 10000
14 snapshot_prefix: "caffe_imagenet_train"
```

Models are schema, not code.


```
1 name: "CaffeNet"
2 input: "data"
3 input_dim: 10
4 input_dim: 3
5 input_dim: 227
6 input_dim: 227
7 # convolution 1: 96 filters
8 layers {
9   layer {
10     name: "conv1"
11     type: "conv"
12     num_output: 96
13     kernel_size: 11
14     stride: 4
15     weight_filler {
16       type: "gaussian"
17       std: 0.01
18     }
19     bias_filler {
20       type: "constant"
21       value: 0.
22     }
23     blobs_lr: 1.
24     blobs_lr: 2.
25     weight_decay: 1.
26     weight_decay: 0.
27   }
28   bottom: "data"
29   top: "conv1"
30 }
```

...

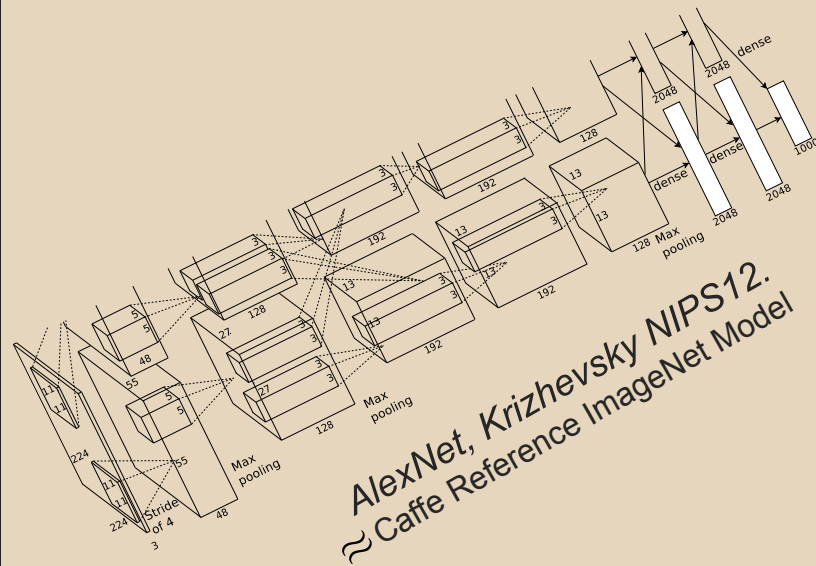
So what is Caffe?

Convolutional Architecture for Fast Feature Embedding

- Research & Engineering
 - Key part of our publication code
 - State-of-the-art models
 - Blazing fast, and it has unit tests!



R-CNN, Girshick CVPR14.



```
Cuda number of devices: 2
Setting to use device 0
Current device id: 0
[=====] Running 10 tests from 2 test cases.
[-----] Global test environment set-up.
[-----] 5 tests from InnerProductLayerTest/0, where TypeParam = float
[ RUN ] InnerProductLayerTest/0.TestSetUp
[ OK ] InnerProductLayerTest/0.TestSetUp (493 ms)
[ RUN ] InnerProductLayerTest/0.TestCPU
[ OK ] InnerProductLayerTest/0.TestCPU (103 ms)
[ RUN ] InnerProductLayerTest/0.TestGPU
[ OK ] InnerProductLayerTest/0.TestGPU (0 ms)
[ RUN ] InnerProductLayerTest/0.TestCPUGradient
[ OK ] InnerProductLayerTest/0.TestCPUGradient (1492 ms)
[ RUN ] InnerProductLayerTest/0.TestGPUGradient
[ OK ] InnerProductLayerTest/0.TestGPUGradient (217 ms)
[-----] 5 tests from InnerProductLayerTest/0 (2305 ms total)
```

So what is Caffe?

Convolutional Architecture for Fast Feature Embedding

- An active research and development community
- See [our contributors](#) and [recent activity](#)

 **527 commits**

★ Star

232

 Fork

113

February 12 2014 - March 12 2014

Period: 1 month ▼

Overview

 **46** Active Pull Requests

 **62** Active Issues

 **23**

Merged Pull Requests

 **23**

Proposed Pull Requests


 **37**

Closed Issues

 **25**

New Issues

23 authors have pushed **511 commits** to all branches, excluding merges. On master, **137 files** have changed and there have been **7,376 additions** and **1,436,167 deletions**.



Why not live caffeine-free?

- It's all about speed.
- cuda-convnet and DeCAF are awesome
 - but cuda-convnet is inflexible
 - and DeCAF is too slow
- Caffe is fast
 - with CPU: 2x speedup over DeCAF
 - with GPU: 10x speedup (under C++)
- Forward pass of a single image takes **2.5ms**
 - Caffe reference ImageNet model with ~60 million parameters
 - (when in batch mode)
 - (~20ms in CPU mode)

Caffe and cuda-convnet


Caffe

- C++/CUDA deep learning and vision
 - library of layers
 - **fast, general-purpose for ILSVRC, PASCAL, your data**
- An active research and development community: public on GitHub
- Seamless GPU/CPU switch
- Model schemas
 - Define the model
 - Configure solver
 - **Finetuning**
- **Wrappers** for Python and MATLAB

cuda-convnet

- C++/CUDA deep learning and vision
 - library of layers
 - **highly-optimized for given data and GPU**
- Static codebase, no community contributions: last update Jul 17, 2012
- GPU only
- Model schemas
 - Define the model
 - Write and run solver command
 - **No finetuning**
- **No wrappers**: monolithic

A Caffe Net

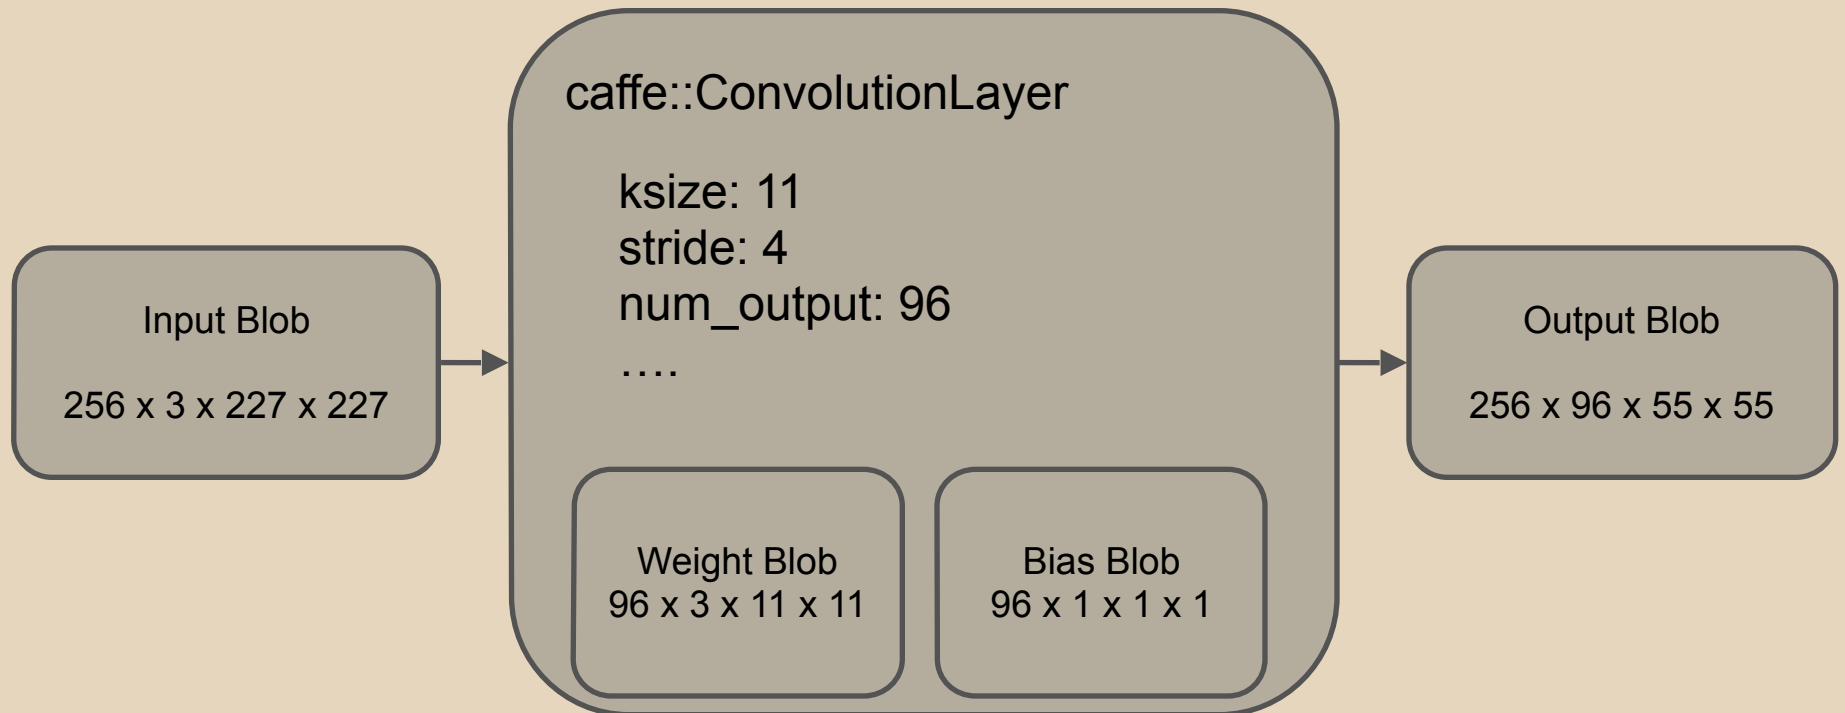


Blob: all your data, derivatives, and parameters.

- example input blob (256 images, RGB, height, width)
 - ImageNet training batches: $256 \times 3 \times 227 \times 227$
- example convolutional parameter blob
 - 96 filters with 3 input channels: $96 \times 3 \times 11 \times 11$


A Layer

- The layer is the fundamental unit of computation.
- Caffe nets are composed of layers as defined in model schema.



A Layer defines...

- Forward: given input, computes the output. \rightarrow
- Backward: given the gradient w.r.t. the output, compute the gradient w.r.t. the input and its internal parameters. \rightarrow
- Setup: how to initialize the layer.



Definition of a Net

Model schema are defined as Protocol Buffers:

```
message NetParameter {  
  optional string name = 1;  
  repeated LayerConnection layers = 2;  
  repeated string input = 3;  
  repeated int32 input_dim = 4;  
}
```

schema definition at /src/caffe/proto/caffe.proto

Protocol Buffer documentation:


<https://developers.google.com/protocol-buffers/docs/overview>

```
name: "linear_regressor"  
input: "data"  
input_dim: 1  
input_dim: 3  
input_dim: 28  
input_dim: 28  
layers {  
  layer {  
    name: "ip"  
    type: "innerproduct"  
    num_output: 10  
  }  
  bottom: "data"  
  top: "prediction"  
}
```

Definition of a Net

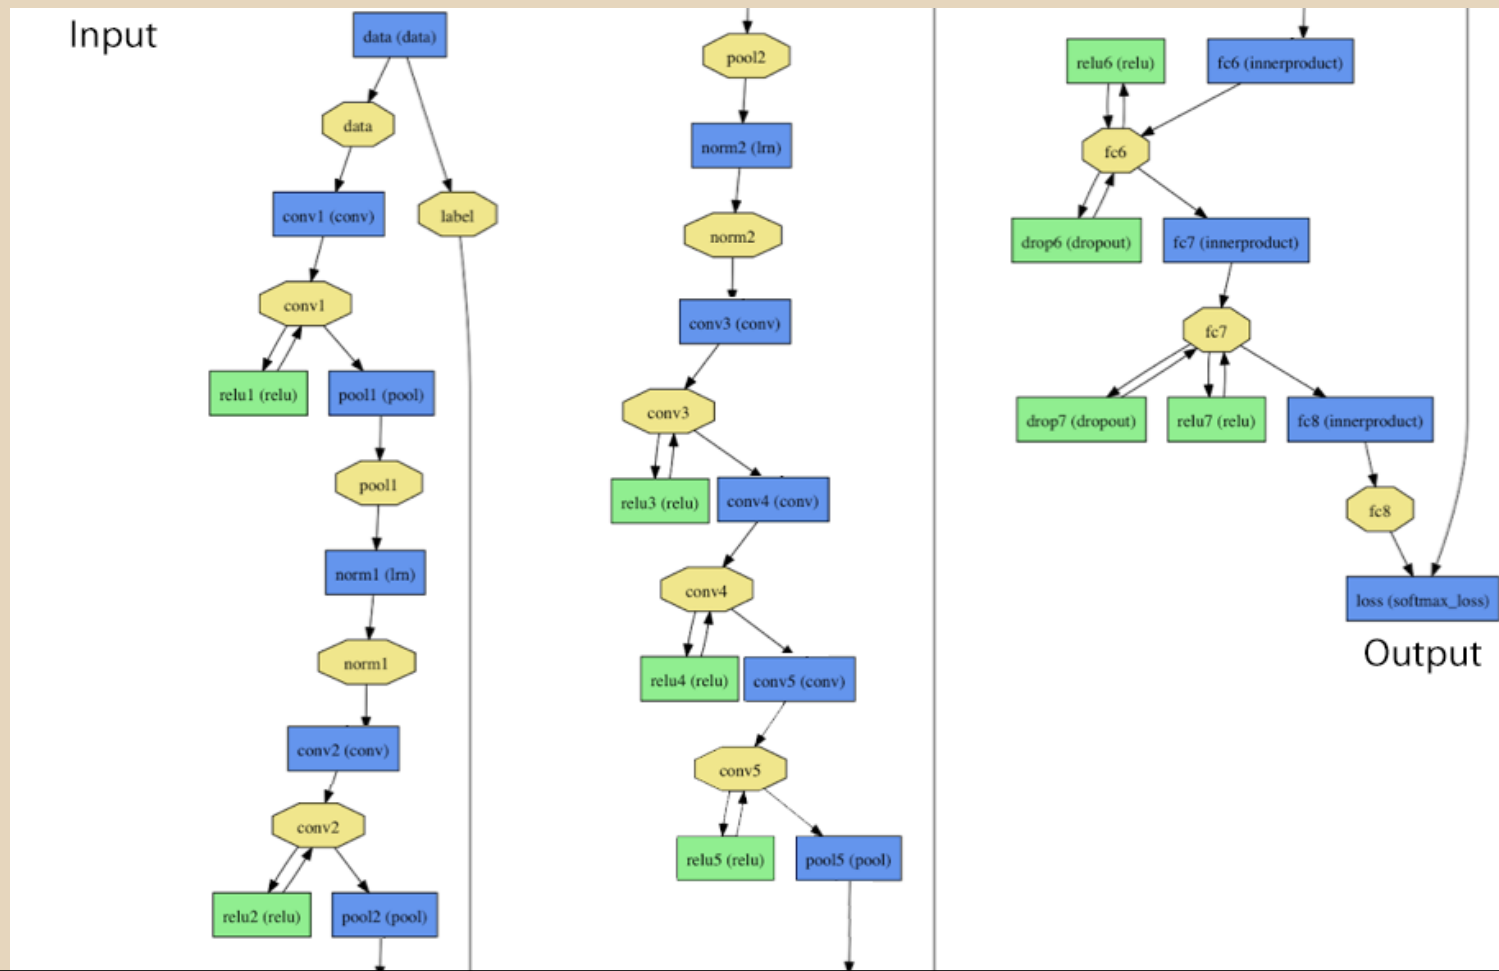
```
name: "mnist-small"
# data layer for input
layers {
  layer {
    name: "mnist"
    type: "data"
    source: "data/mnist-train-leveldb"
    batchsize: 64
    scale: 0.00390625
  }
  top: "data"
  top: "label"
}
# linear classifier by inner product
layers {
  layer {
    name: "ip"
    type: "innerproduct"
    num_output: 10
    weight_filler {
      type: "xavier"
    }
  }
  bottom: "data"
  top: "ip"
}
```

```
# softmax loss for training
# takes classifier output and labels
layers {
  layer {
    name: "prob"
    type: "softmax_loss"
  }
  bottom: "ip"
  bottom: "label"
}
```



How about ImageNet?

- It's another network definition... only this time a state-of-the-art model
- See `caffe/models/imagenet.prototxt`



Finetuning

- Once you have a model—like `caffe_reference_imagenet_model`—you can solve many problems.
- Where training from scratch can fail for lack of sufficient data, finetuning can succeed.
- Rename the layers you need to change...
- ...and continue training.
- No coding needed.

```
layer {  
  name: "fc8"  
  type: "innerproduct"  
  num_output: 1000  
  ...  
}
```

```
layer {  
  name: "fc8-t"  
  type: "innerproduct"  
  num_output: 397  
  ...  
}
```

A Few Practical Questions

- What's the shortest path to features?
 - Swap deep features into your pipeline without tears via the Caffe Reference ImageNet model.
 - Any layer can be extracted.
 - Prototype with Python and MATLAB wrappers.
- Do I have to train from scratch for every problem?
 - Not at all! Finetune learned models to new data and tasks.
 - Define a new model and solver.
 - Call `./finetune_net new_solver old_model #` then get a cup of coffee
- What do I do with my own loss, special operation, or data format?
 - Well, this is trickier but doable.
 - Code the layers needed.
 - Define the model and carry on.

Questions!

Check out caffe.berkeleyvision.org,
the Github repository <https://github.com/BVLC/caffe>,
and our issue tracker <https://github.com/BVLC/caffe/issues> (but search before posting).

Try our examples and tutorials!

References

[SuperVision/AlexNet] A. Krizhevsky, I. Sutskever, and G. Hinton. *Imagenet classification with deep convolutional neural networks*. NIPS, 2012.

[DeCAF] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. *Decaf: A deep convolutional activation feature for generic visual recognition*. ICML, 2014.

[R-CNN] R. Girshick, J. Donahue, T. Darrell, and J. Malik. *Rich feature hierarchies for accurate object detection and semantic segmentation*. CVPR, 2014.