

Appendix C CNN Examples with Caffe, YOLOv3 and PyTorch

 **Note:** This appendix is also available as a colored version of PDF from this book's download website.

This appendix demonstrates a few example CNN implementations with Caffe in C++, YOLOv3 in C and PyTorch in Python. We choose the Caffe, YOLOv3 and PyTorch deep learning frameworks, as they are three of the most popular frameworks for solving computer vision related machine learning tasks. Besides, if you decide to have a career in machine learning, you will have a huge advantage if you have good programming skills in Python and C/C++. However, you don't have to be a C/C++ expert to try out the example CNN models with Caffe, YOLOv3 and PyTorch to be introduced in this appendix. Some basic knowledge about how Python and C/C++ work in general and how Unix shell scripts work would be sufficient.

I have to mention that YOLOv3 perhaps is the state of the art deep learning framework that you may want to focus on if you look for a production-quality DL framework. You can jump to YOLO directly, which starts with §C.2 *The YOLOv3 Framework*. Otherwise, let's start with Caffe first next.

C.1 THE CAFFE FRAMEWORK

Since Caffe is written in C++, you have to build it from the source in order to make it run on your machine. I'll show you how to build the Caffe framework from the source next.

C.1.1 BUILDING THE CAFFE FRAMEWORK FROM THE SOURCE

First, let's see how we can build the Caffe framework from the source. By going through such a process, you will have the following benefits:

- You will understand what other software packages that Caffe depends on.
- You will have access to the C++ source files of Caffe, just in case you want to check out how this popular, production quality framework is implemented in C++.

- As a machine learning engineer, it's important that you can quickly get a framework up and running on your machine and start to get your project going immediately.

Next, I'll share with you how I rebuilt Caffe on my MacBook Pro by following the instructions given at http://caffe.berkeleyvision.org/install_osx.html, especially, what worked and what didn't work, and how I worked around the issues I encountered. If you go along yourself, it may take you several days, but with the help of this appendix, it could be much easier for you, especially if you are not very familiar with C++ and a typical Unix-like environment. Of course, if you are already a C++ professional, it would be easy for you.

The installation of Caffe starts with installing some general dependencies. On macOS, you need to have *homebrew* installed on your machine first. If you do not have homebrew installed already on your machine, search online and get it installed first.

Then, follow the below procedure:

1. Download the latest Caffe source at <https://github.com/BVLC/caffe> and place it in a directory on your machine. For example, I downloaded and placed it on my machine at */Users/henryliu/mspc/devs/ws_cpp/Caffe*. This is my Eclipse C/C++ workspace directory, as I can navigate and view various files easily on such an IDE. Also, add a line in your *.bashrc* file, e.g., *export CAFE_ROOT=/Users/henryliu/mspc/devs/ws_cpp/Caffe*, to set the *CAFFE_ROOT* environment variable. You will need this when you try out some of the CNN models on Caffe later. In case you are not familiar with Unix environment, execute "source ~/.bashrc" on a command terminal to enable all environment variables defined in that file.
2. Execute "cd \$CAFFE_ROOT" on the command terminal and then execute *brew install -vd snappy leveldb gflags glog szip lmdb* by copying this command from that website to your local command terminal. Table C.1 describes what these dependencies are about.
3. The next command to execute is: *brew tap homebrew/science*, which did not work on my machine as it does not exist anymore. It turned out that you can just ignore it.
4. Execute *brew install hdf5 opencv*. Check out what *opencv* is about from Table C.1.
5. I don't use Anaconda since it once messed up my Python environment on my machine. Therefore, I chose the option of no Anaconda for the next part of the installation.
6. Execute *brew install --build-from-source --with-python -vd protobuf*. Check out what *protobuf* is about from Table C.1.
7. Execute *brew install --build-from-source -vd boost boost-python*. Check out what *boost* is about from Table C.1.
8. Execute *brew install protobuf boost*.
9. Next, it mentions that BLAS is already installed as the *Accelerate/vecLib* framework, which is Apple's implementation of BLAS. Check out what BLAS is about from Table C.1.

The dependency installation is completed now. Next, compile Caffe by following the procedure given after Table C.1.

Table C.1 Caffe dependencies

Feature	Semantics
snappy	A fast compressor/decompressor written in C++.

leveldb	A fast key-value storage library written in C++ at Google that provides an ordered mapping from string keys to string values.
gflags	A C++ library that implements command line flags processing.
glog	C++ implementation of the Google logging module.
szip	Provides lossless compression of scientific data from HDF5, which is a unique technology suite that makes possible the management of extremely large and complex data collections.
lmdb	A Btree-based Lightning Memory-Mapped Database Manager (LMDB).
opencv	OpenCV stands for Open Source Computer Vision Library. Written in optimized C/C++, the library can take advantage of multi-core processing. Enabled with OpenCL, it can take advantage of the hardware acceleration of the underlying heterogeneous compute platform.
protobuf	Protocol Buffers - Google's data interchange format.
boost	Over 80 C++ based individual libraries for tasks and data structures such as linear algebra, pseudorandom number generation, multithreading, image processing, regular expressions, and unit testing.
BLAS	The BLAS (Basic Linear Algebra Subprograms) are routines that provide standard building blocks for performing basic vector and matrix operations. The Level 1 BLAS perform scalar, vector and vector-vector operations, the Level 2 BLAS perform matrix-vector operations, and the Level 3 BLAS perform matrix-matrix operations. Because the BLAS are efficient, portable, and widely available, they are commonly used in the development of high quality linear algebra software.

To compile Caffe, it became tricky in my case. I followed the instructions under the section named *Compilation with Make* and it ended up with the following error:

Undefined symbols for architecture x86_64:

```
"cv::imread(cv::String const&, int)", referenced from:
  caffe::WindowDataLayer<float>::load_batch(caffe::Batch<float>*) in window_data_layer.o
  caffe::WindowDataLayer<double>::load_batch(caffe::Batch<double>*) in window_data_layer.o
  caffe::ReadImageToCVMat(std::__1::basic_string<char, std::__1::char_traits<char>, std::__1::allocator<char>>
const&, int, int, bool) in io.o
"cv::imdecode(cv::_InputArray const&, int)", referenced from:
  caffe::DecodeDatumToCVMatNative(caffe::Datum const&) in io.o
  caffe::DecodeDatumToCVMat(caffe::Datum const&, bool) in io.o
"cv::imencode(cv::String const&, cv::_InputArray const&, std::__1::vector<unsigned char,
std::__1::allocator<unsigned char>>&, std::__1::vector<int, std::__1::allocator<int>>> const&)", referenced from:
  caffe::ReadImageToDatum(std::__1::basic_string<char, std::__1::char_traits<char>, std::__1::allocator<char>>
const&, int, int, int, bool, std::__1::basic_string<char, std::__1::char_traits<char>, std::__1::allocator<char>>
const&, caffe::Datum*) in io.o
ld: symbol(s) not found for architecture x86_64
clang: error: linker command failed with exit code 1 (use -v to see invocation)
make: *** [.build_release/lib/libcaffe.so.1.0.0] Error 1
```

I spent a lot of time searching online and nothing helped. Then, it worked when I followed the instructions under the section named *CMake Build*. Therefore, the procedure given below is based on my experience with *CMake Build*:

- `cd $CAFFE_ROOT`
- `cp Makefile.config.example Makefile.config`. Then, in my case, I opened the *Makefile.config* file and made two changes:
 - Uncommented the line of `CPU_ONLY := 1`, since I do not have a GPU on my machine.
 - Uncommented the lines for using Python 3 instead of 2.
- Then, I executed each of the following commands as instructed:

```
$mkdir build
$cd build
$cmake ..
$make all
$make install
$make runtest
```

All of the above commands were successful. However, I tried the command *make distribute* and encountered the error of “*target not defined.*” This was okay as I wanted to run Caffe on my local machine anyway. Figure C.1 shows the code structure of the Caffe framework on my C/C++ Eclipse IDE.

If you have gotten to this step, you are ready to try out a few example CNN models as described in the next few sections.

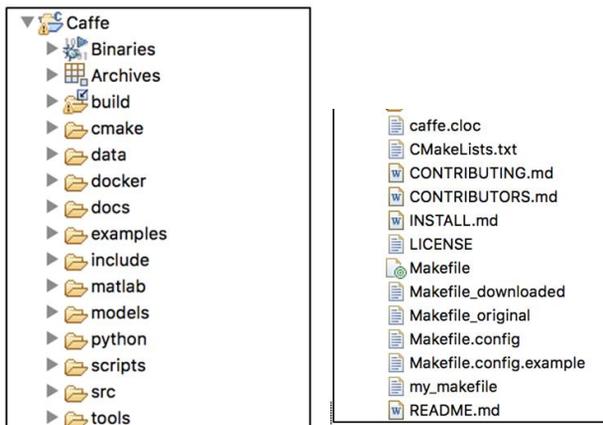


Figure C.1 Code structure of the Caffe framework.

C.1.2 THE LENET CNN MODEL FOR THE MNIST DATASET WITH CAFFE

Caffe has many examples available. However, it's better to start with the MNIST dataset, not because you are already familiar with the MNIST dataset, but because this example contains detailed descriptions about how to define a CNN model to work with Caffe. Therefore, let's get started with this example first. Once again, make sure you have the `CAFFE_ROOT` environment variable set in your environment per instructions given in the previous section.

C.1.2.1 PREPARE THE MNIST DATASET

First, if you don't have `wget` installed on your machine, execute the following command to get it installed:

```
$brew install wget --with-libressl
```

Then, execute the following commands:

```
$cd $CAFFE_ROOT
$./data/mnist/get_mnist.sh
$./examples/mnist/create_mnist.sh
```

After executing the above commands, you should have four files with their names ending with `-ubyte` in the `data/mnist` directory. These are the training and testing dataset we will use.

C.1.2.2 DEFINING THE LENET MODEL

Next, the instruction explains about the LeNet model to be used with the MNIST dataset we have just prepared. I assume that you have studied Chapter 10 of the main text, so I would not repeat about the LeNet here. However, there is a deviation here: The Caffe model here uses the ReLU activation function instead of the sigmoid function as was the case with the original LeNet model, since it has become common knowledge that the ReLU activation function works better than the sigmoid activation function.

Now, let's explain how Caffe defines a CNN model. With Caffe, each model is defined in a text file, e.g., the file `$CAFFE_ROOT/examples/mnist/lenet_train_test.prototxt` in this case for the LeNet model. You can now open this file and examine its contents. It starts with a line of `name: "LeNet"`, followed by 11 segments labeled "layer." To understand this model definition file, perhaps this is a good time for me to help you understand several Caffe jargons as follows:

- **Blobs.** Caffe stores and communicates data in 4D arrays called Blobs.
- **Models.** Caffe models are saved to disk using Google Protocol Buffers.
- **Data.** Caffe stores large scale data in LevelDB databases.
- **Layer.** Defines one or more blobs as input or output to be used in forward and backward passes.
- **Layer Types.** Include: data, convolution, pooling, inner products (ip's), nonlinearities (ReLU, logistic, etc.), local response normalization, element-wise operations, losses (softmax, hinge, etc.), and so on.

Given what we have covered in the main text, you should have no difficulties understanding the above concepts.

Defining a data layer

The data layers define the *data* and *label* blobs for the training and testing datasets, as shown in Listing C.1. Here, every item is obvious except that (1) the `transform_param` segment defines how data should

be transformed, e.g., scaled or normalized by being multiplied with a number 0.00390625, which is just the reciprocal of 256, and (2) the `data_param` segment defines the data source. In addition, this one file defines data blobs with the `include` attribute for both the training phase and the testing phase, and Caffe knows which data blob to choose, based on the phase it is in. These are called layer rules, which are defined in a large file in `$CAFFE_ROOT/src/caffe/proto/caffe.proto`. You can take a quick look at this file to get an idea on how Caffe rules are defined.

Next, we discuss the convolution layer.

Listing C.1. MNIST training and testing data layers with Caffe

```

layer {
  name: "mnist"
  type: "Data"
  top: "data"
  top: "label"
  include {
    phase: TRAIN
  }
  transform_param {
    scale: 0.00390625
  }
  data_param {
    source: "examples/mnist/mnist_train_lmdb"
    batch_size: 64
    backend: LMDB
  }
}
layer {
  name: "mnist"
  type: "Data"
  top: "data"
  top: "label"
  include {
    phase: TEST
  }
  transform_param {
    scale: 0.00390625
  }
  data_param {
    source: "examples/mnist/mnist_test_lmdb"
    batch_size: 100
    backend: LMDB
  }
}

```

Defining a convolution layer with Caffe

Listing C.2 shows how a convolution layer is defined. It can be understood as follows:

- The `bottom` attribute defines the prior layer while the `top` attribute defines the current layer.
- The `param` attributes define the learning rate multipliers for the weights and biases, respectively. In this case, the weights multiplier is 1 and the biases multiplier is 2, which are applied to the learning rate determined by the solver during runtime.
- The `convolution_param` attribute defines the settings for carrying out the convolution. In this case, it specifies 20 output channels with a kernel size of 5 and a stride of 1.
- The `weight_filler` attribute specifies how weights should be randomly initialized. In this case, it specifies to use the Xavier algorithm to automatically determine the scale of initialization based on the number of input and output neurons.
- The `bias_filler` attribute specifies how biases should be initialized. In this case, it specifies that biases should be initialized as constant, with the default filling value of 0.

Next, we discuss how a pooling layer is defined with Caffe.

Listing C.2 A convolution layer defined with Caffe

```
layer {
  name: "conv1"
  type: "Convolution"
  bottom: "data"
  top: "conv1"
  param {
    lr_mult: 1
  }
  param {
    lr_mult: 2
  }
  convolution_param {
    num_output: 20
    kernel_size: 5
    stride: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
    }
  }
}
```

Defining a pooling layer with Caffe

Listing C.3 shows how a pooling layer can be defined with Caffe. In this case, it specifies which convolution layer to follow as defined by the `bottom` attribute, and the pooling settings such as using the max pooling with a kernel size of 2 and a stride of 2. In this case, there are no overlaps between neighboring pooling regions.

Next, we discuss how a fully connected layer is defined with Caffe.

Listing C.3 A pooling layer defined with Caffe

```
layer {
  name: "pool1"
  type: "Pooling"
  bottom: "conv1"
  top: "pool1"
  pooling_param {
    pool: MAX
    kernel_size: 2
    stride: 2
  }
}
```

Defining a fully connected layer with Caffe

Listing C.4 shows how a fully connected layer, which designated as type `InnerProduct`, can be defined with Caffe. In this case, it specifies which layer to follow as defined by the `bottom` attribute, and uses an `inner_product_param` attribute to specify the number of outputs as well as the weight and bias fillers.

Next, we discuss how an ReLU layer is defined with Caffe.

Listing C.4 A fully connected layer defined with Caffe

```
layer {
  name: "ip1"
  type: "InnerProduct"
  bottom: "pool2"
  top: "ip1"
  param {
    lr_mult: 1
  }
  param {
    lr_mult: 2
  }
  inner_product_param {
    num_output: 500
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
    }
  }
}
```

Defining an ReLU layer with Caffe

Listing C.5 shows how an ReLU layer can be defined with Caffe. In this case, both the `bottom` attribute and the `top` attribute are specified to be the same fully connected layer, which makes sense as an ReLU is not necessarily a layer by itself at all – it just performs an element-wise operation, which can be done *in-place* to save memory.

However, note how Listing C.6 defines another fully connected layer, following the ReLU layer described in Listing C.5. In particular, the `ip1` layer, not the `relu1` layer, is assigned to the `bottom` attribute, as an ReLU layer is more of an element-wise operation than an actual layer.

Next, we discuss how an accuracy layer is defined with Caffe.

Listing C.5 An ReLU layer defined with Caffe

```
layer {
  name: "relu1"
  type: "ReLU"
  bottom: "ip1"
  top: "ip1"
}
```

Listing C.6 A fully connected layer following an ReLU layer defined with Caffe

```
layer {
  name: "ip2"
  type: "InnerProduct"
  bottom: "ip1"
  top: "ip2"
  param {
    lr_mult: 1
  }
  param {
    lr_mult: 2
  }
  inner_product_param {
    num_output: 10
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
    }
  }
}
```

Defining an accuracy layer with Caffe

Listing C.7 shows how an accuracy layer can be defined with Caffe. In this case, two `bottom` attributes are specified as inputs, the `ip2` layer and the `label` “layer.” It also specifies that this layer should be used in the `TEST` phase.

Next, we discuss how a loss layer is defined with Caffe.

Listing C.7 An accuracy layer defined with Caffe

```
layer {
  name: "accuracy"
  type: "Accuracy"
  bottom: "ip2"
  bottom: "label"
  top: "accuracy"
  include {
    phase: TEST
  }
}
```

Defining a loss layer with Caffe

Listing C.8 shows how a loss layer can be defined with Caffe, which should be the final layer of a CNN model with Caffe. In this case, two `bottom` attributes are specified as inputs, the `ip2` layer and the `label` “layer.” The `ip2` layer provides predictions while the `label` layer provides target values, both of which are used for computing the loss, which is the basis for the back-propagation algorithm to work.

Next, we discuss how the solver is defined with Caffe for the LetNet model with the MNIST dataset.

Listing C.8 A loss layer defined with Caffe

```
layer {
  name: "loss"
  type: "SoftmaxWithLoss"
  bottom: "ip2"
  bottom: "label"
  top: "loss"
}
```

C.1.2.3 DEFINING THE SOLVER FOR THE MNIST DATASET WITH CAFFE

The file `$CAFFE_ROOT/examples/mnist/lenet_solver.prototxt` defines the solver, which specifies the end-to-end process for running the entire job. Listing C.9 shows the entire contents of this file. Since we have basic concepts covered in the main text and every line is clearly annotated, we would not spend time to explain every line, except that the `solver_mode` specified at the end of the file should be changed to `CPU` if you do not have a GPU equipped with your machine.

Listing C.9 lenet_solver.prototxt

```

# The train/test net protocol buffer definition
net: "examples/mnist/lenet_train_test.prototxt"
# test_iter specifies how many forward passes the test should carry out.
# In the case of MNIST, we have test batch size 100 and 100 test iterations,
# covering the full 10,000 testing images.
test_iter: 100
# Carry out testing every 500 training iterations.
test_interval: 500
# The base learning rate, momentum and the weight decay of the network.
base_lr: 0.01
momentum: 0.9
weight_decay: 0.0005
# The learning rate policy
lr_policy: "inv"
gamma: 0.0001
power: 0.75
# Display every 100 iterations
display: 100
# The maximum number of iterations
max_iter: 10000
# snapshot intermediate results
snapshot: 5000
snapshot_prefix: "examples/mnist/lenet"
# solver mode: CPU or GPU
solver_mode: CPU

```

C.1.2.4 KICKING OFF TRAINING AND TESTING WITH CAFFE

The `examples/mnist/lenet_train_test.prototxt` and `examples/mnist/lenet_solver.prototxt` files are called Caffe *protobuf* files. Once they are prepared, just run the following two commands to kick off training and testing:

```

cd $CAFFE_ROOT
./examples/mnist/train_lenet.sh

```

The command specified in the script `train_lenet.sh` is as follows:

```

./build/tools/caffe train --solver=examples/mnist/lenet_solver.prototxt

```

As you see, use Caffe to solve an applicable machine learning problem consists of the following three steps:

1. Compose a network model definition file similar to the `lenet_train_test.prototxt` file.
2. Compose a job process definition file similar to the `lenet_solver.prototxt` file.
3. Compose a script similar to the script `train_lenet.sh` and run it.

Listing C.10 shows running the above MNIST LeNet model with Caffe on my machine. Note that I just picked a few segments for illustrative purposes. As you can see, the test started at 21:57:27 and ended at 22:03:19 for a total duration of 4m36s, with an accuracy of 99.09% achieved after 10000 iterations! This is outstanding performance by any means.

Listing C.10 Sample output of running the MNIST LeNet model with Caffe

```

henryliu:Caffe henryliu$ ./examples/mnist/train_lenet.sh
I0310 21:57:27.226054 2508161984 caffe.cpp:197] Use CPU.
I0310 21:57:27.227905 2508161984 solver.cpp:45] Initializing solver from parameters:
.....
O310 21:57:27.230612 2508161984 layer_factory.hpp:77] Creating layer mnist
I0310 21:57:27.231889 2508161984 db_lmdb.cpp:35] Opened lmdb examples/mnist/mnist_train_lmdb
I0310 21:57:27.232677 2508161984 net.cpp:84] Creating Layer mnist
I0310 21:57:27.232699 2508161984 net.cpp:380] mnist -> data
I0310 21:57:27.232717 2508161984 net.cpp:380] mnist -> label
I0310 21:57:27.232748 2508161984 data_layer.cpp:45] output data size: 64,1,28,28
I0310 21:57:27.237839 2508161984 net.cpp:122] Setting up mnist
I0310 21:57:27.237856 2508161984 net.cpp:129] Top shape: 64 1 28 28 (50176)
I0310 21:57:27.237865 2508161984 net.cpp:129] Top shape: 64 (64)
.....
I0310 21:58:13.941082 2508161984 solver.cpp:239] Iteration 1300 (33.0033 iter/s, 3.03s/100 iters), loss =
0.0233421
I0310 21:58:13.941115 2508161984 solver.cpp:258]   Train net output #0: loss = 0.0233422 (* 1 = 0.0233422 loss)
I0310 21:58:13.941123 2508161984 sgd_solver.cpp:112] Iteration 1300, lr = 0.00912412
I0310 21:58:16.960737 2508161984 solver.cpp:239] Iteration 1400 (33.1236 iter/s, 3.019s/100 iters), loss =
0.00798987
I0310 21:58:16.960772 2508161984 solver.cpp:258]   Train net output #0: loss = 0.00798988 (* 1 = 0.00798988
loss)
I0310 21:58:16.960778 2508161984 sgd_solver.cpp:112] Iteration 1400, lr = 0.00906403
I0310 21:58:19.946302 2508161984 solver.cpp:351] Iteration 1500, Testing net (#0)
.....
I0310 22:03:13.821404 2508161984 sgd_solver.cpp:112] Iteration 9900, lr = 0.00596843
I0310 22:03:16.879815 2508161984 solver.cpp:468] Snapshotting to binary proto file
examples/mnist/lenet_iter_10000.caffemodel
I0310 22:03:16.900782 2508161984 sgd_solver.cpp:280] Snapshotting solver state to binary proto file
examples/mnist/lenet_iter_10000.solverstate
I0310 22:03:16.918725 2508161984 solver.cpp:331] Iteration 10000, loss = 0.00294297
I0310 22:03:16.918767 2508161984 solver.cpp:351] Iteration 10000, Testing net (#0)
I0310 22:03:19.175561 131223552 data_layer.cpp:73] Restarting data prefetching from start.
I0310 22:03:19.274293 2508161984 solver.cpp:418]   Test net output #0: accuracy = 0.9909
I0310 22:03:19.274327 2508161984 solver.cpp:418]   Test net output #1: loss = 0.0286514 (* 1 = 0.0286514 loss)
I0310 22:03:19.274333 2508161984 solver.cpp:336] Optimization Done.
I0310 22:03:19.274336 2508161984 caffe.cpp:250] Optimization Done.

```

C.1.3 ALEX'S CIFAR-10 WITH CAFFE

If you have successfully completed the previous exercise, then you have learnt how Caffe works! You can verify your learning with this second Caffe deep learning CNN example.

First, let's learn a bit about the CIFAR-10 dataset. This is a dataset created by Alex Krizhevsky at the Canadian Institute for Advanced Research (CIFAR), with 10 classes of images of 32x32 pixels. It has 6000 images per class, for a total of 60,000 images. Out of 60,000 images, 50,000 are used as training images and 10,000 are used as test images. The 50,000 training images are split into 5 batches, with each

batch containing 10,000 images. Figure C.2 shows 10 sample images for each of the 10 classes. You can find more about this dataset at <https://www.cs.toronto.edu/~kriz/cifar.html>.

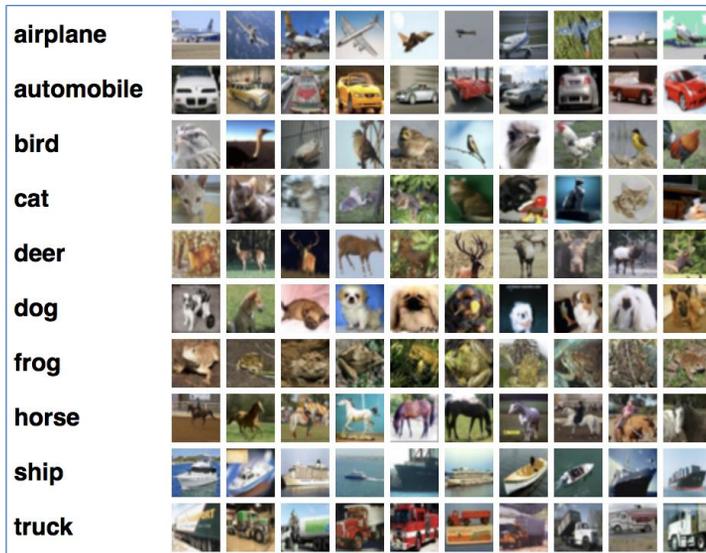


Figure C.2 Samples for Alex's CIFAR-10 dataset.

Now, in terms of trying out this dataset with Caffe, I'd like to take a different approach. In the previous section with the MNIST dataset, we first examined the model description file, then the job description file, and finally the script for kicking off the training process. For this example, I'd like to reverse the process, namely, we first look at the script for kicking off the training process, then the job description file, and finally the model definition file. I feel this may help you understand how Caffe framework works better.

C.1.3.1 THE SCRIPT FOR KICKING OFF THE TRAINING PROCESS

Listing C.11 shows the script `$CAFFE_ROOT/examples/cifar10/train_quick.sh`. It requires to have two job description files to feed to the solver: `cifar10_quick_solver.prototxt` and `cifar10_quick_solver_lr1.prototxt`, which will be discussed in the next section. The trained model is saved to a snapshot file named `cifar10_quick_iter_4000.solverstate`.

Listing C.11 CIFAR-10 train-quick.sh script

```
#!/usr/bin/env sh
set -e

TOOLS=./build/tools

$TOOLS/caffe train \
  --solver=examples/cifar10/cifar10_quick_solver.prototxt $@
```

```
# reduce learning rate by factor of 10 after 8 epochs
$TOOLS/caffe train \
  --solver=examples/cifar10/cifar10_quick_solver_lr1.prototxt \
  --snapshot=examples/cifar10/cifar10_quick_iter_4000.solverstate $@
```

C.1.3.2 THE JOB DESCRIPTION FILES

Listings C.12 and C.13 show the two job description files: `cifar10_quick_solver.prototxt` and `cifar10_quick_solver_lr1.prototxt`, respectively. This can be considered a two-phase training, with the learning rate reduced to 10x smaller in the second training phase. Note also that by default, the `solver_mode` was set to GPU, but I have changed it to CPU for the same reason explained in the previous section. You should do the same if you do not have a GPU installed on your machine.

Next, we check out the model definition file for this example.

Listing C.12 The `cifar10_quick_solver.prototxt` file

```
# reduce the learning rate after 8 epochs (4000 iters) by a factor of 10

# The train/test net protocol buffer definition
net: "examples/cifar10/cifar10_quick_train_test.prototxt"
# test_iter specifies how many forward passes the test should carry out.
# In the case of MNIST, we have test batch size 100 and 100 test iterations,
# covering the full 10,000 testing images.
test_iter: 100
# Carry out testing every 500 training iterations.
test_interval: 500
# The base learning rate, momentum and the weight decay of the network.
base_lr: 0.001
momentum: 0.9
weight_decay: 0.004
# The learning rate policy
lr_policy: "fixed"
# Display every 100 iterations
display: 100
# The maximum number of iterations
max_iter: 4000
# snapshot intermediate results
snapshot: 4000
snapshot_prefix: "examples/cifar10/cifar10_quick"
# solver mode: CPU or GPU
solver_mode: CPU
```

Listing C.13 The `cifar10_quick_solver_lr1.prototxt` file

```
# reduce the learning rate after 8 epochs (4000 iters) by a factor of 10
```

```

# The train/test net protocol buffer definition
net: "examples/cifar10/cifar10_quick_train_test.prototxt"
# test_iter specifies how many forward passes the test should carry out.
# In the case of MNIST, we have test batch size 100 and 100 test iterations,
# covering the full 10,000 testing images.
test_iter: 100
# Carry out testing every 500 training iterations.
test_interval: 500
# The base learning rate, momentum and the weight decay of the network.
base_lr: 0.0001
momentum: 0.9
weight_decay: 0.004
# The learning rate policy
lr_policy: "fixed"
# Display every 100 iterations
display: 100
# The maximum number of iterations
max_iter: 5000
# snapshot intermediate results
snapshot: 5000
snapshot_format: HDF5
snapshot_prefix: "examples/cifar10/cifar10_quick"
# solver mode: CPU or GPU
solver_mode: CPU

```

C.1.3.3 THE MODEL DEFINITION FILE

Listing C.14 shows the model definition file for this example. It's kind of lengthy, but not very different from the model file we discussed in the previous section for the MNIST dataset, except that it has more layers. Please take your time and go through it end to end to make sure that you understand it, or even better, make a sketch drawing by going through all layers from bottom to top.

Listing C.14 The model definition file for the Caffe CIFAR-10 example (partial)

```

name: "CIFAR10_quick"
layer {
  name: "cifar"
  type: "Data"
  top: "data"
  top: "label"
  include {
    phase: TRAIN
  }
  transform_param {
    mean_file:
"examples/cifar10/mean.binaryproto"
  }
  data_param {
    source:
"examples/cifar10/cifar10_train_lmdb"
    batch_size: 100
    backend: LMDB
  }
}
layer {
  name: "cifar"
  type: "Data"
  top: "data"
  top: "label"
  include {

```

```

    phase: TEST
  }
  transform_param {
    mean_file:
"examples/cifar10/mean.binaryproto"
  }
  data_param {
    source:
"examples/cifar10/cifar10_test_lmdb"
    batch_size: 100
    backend: LMDB
  }
}
.....
layer {
  name: "accuracy"
  type: "Accuracy"
  bottom: "ip2"
  bottom: "label"
  top: "accuracy"
  include {
    phase: TEST
  }
}
layer {
  name: "loss"
  type: "SoftmaxWithLoss"
  bottom: "ip2"
  bottom: "label"
  top: "loss"
}

```

C.1.3.4 RUNNING THE CIFAR-10 EXAMPLE

I ran this example successfully on my machine, except that it took close to an hour to download the CIFAR-10 dataset of ~170MB, due to my slow wifi connection. Listing C.15 shows the final accuracy of 75.68%.

If you want to try it out, make necessary changes such as the `solver_mode`, and then run the following commands on your machine to get it going:

```

$cd $CAFFE_ROOT
$./examples/cifar10/train_quick.sh

```

If you encounter any issues, check out <http://caffe.berkeleyvision.org/gathered/examples/cifar10.html> for more detailed instructions.

Listing C.15 Sample output of running the CIFAR-10 example

```

.....
I0310 14:35:48.097317 2508161984 sgd_solver.cpp:112] Iteration 4800, lr = 0.0001
I0310 14:36:09.057718 2508161984 solver.cpp:239] Iteration 4900 (4.77099 iter/s, 20.96s/100 iters), loss =
0.465986
I0310 14:36:09.057766 2508161984 solver.cpp:258] Train net output #0: loss = 0.465986 (* 1 = 0.465986 loss)
I0310 14:36:09.057773 2508161984 sgd_solver.cpp:112] Iteration 4900, lr = 0.0001
I0310 14:36:29.173247 73412608 data_layer.cpp:73] Restarting data prefetching from start.
I0310 14:36:30.006633 2508161984 solver.cpp:478] Snapshotting to HDF5 file
examples/cifar10/cifar10_quick_iter_5000.caffemodel.h5
I0310 14:36:30.015507 2508161984 sgd_solver.cpp:290] Snapshotting solver state to HDF5 file
examples/cifar10/cifar10_quick_iter_5000.solverstate.h5
I0310 14:36:30.114841 2508161984 solver.cpp:331] Iteration 5000, loss = 0.525545
I0310 14:36:30.114869 2508161984 solver.cpp:351] Iteration 5000, Testing net (#0)
I0310 14:36:39.559231 73949184 data_layer.cpp:73] Restarting data prefetching from start.
I0310 14:36:39.941176 2508161984 solver.cpp:418] Test net output #0: accuracy = 0.7568
I0310 14:36:39.941207 2508161984 solver.cpp:418] Test net output #1: loss = 0.735389 (* 1 = 0.735389 loss)

```

I0310 14:36:39.941213 2508161984 solver.cpp:336] Optimization Done.
 I0310 14:36:39.941217 2508161984 caffe.cpp:250] Optimization Done.

C.1.4 THE IMAGENET EXAMPLE WITH CAFFE

Given the two examples we covered in the previous sections, you should be able to follow the instructions at <http://caffe.berkeleyvision.org/gathered/examples/imagenet.html> to try out the ImageNet example with Caffe. If you decide to try it out, download the ImageNet data from the website at <http://www.image-net.org/challenges/LSVRC/2012/nonpub-downloads>. The entire data amounts to ~160 GB, which could be challenging to download if you do not have a fast Internet connection. In my case, I downloaded the following three files at home with a cable connected to a Windows PC:

- *ILSVRC2012_img_train.tar* of 32.96GB with 258,434 images (~22%) instead of the full set of ~1.2M images of ~138GB.
- *ILSVRC2012_img_val.tar* of 6.74GB with all 50,000 images.
- *ILSVRC2012_img_test.tar* of 13.69GB with all 100,000 images.

Then I double-clicked on the file *ILSVRC2012_img_train.tar*, renamed the directory to *train*, created the following shell script, and executed it to untar all JPEG files from each tar file.

```
#!/bin/bash
for name in ./*.tar; do
    tar_name=$(basename "$name")
    dir_name="${tar_name%.*}"
    #echo $dir_name
    mkdir -p $dir_name
    tar -xvf $name -C $dir_name
done
```

Then I followed the instructions given in the *readme.md* file located in the directory of *examples/imagenet* as follows:

1. **Data Preparation.** I executed the script *./data/ilsvrc12/get_ilsvrc_aux.sh* and downloaded the required auxiliary data from http://dl.caffe.berkeleyvision.org/caffe_ilsvrc12.tar.gz, which is not ImageNet data. After this step, the files placed in the *data/ilsvrc12* directory include: *det_synset_words.txt*, *imagenet_mean.binaryproto*, *imagenet_bet.pickle*, *synset_words.txt*, *synsets.txt*, *test.txt*, *train.txt*, and *val.txt*. The *imagenet_mean.binaryproto* and *imagenet_bet.pickle* are binary files, while all others ending with *.txt* are text files. The text files describe what each of the images is, either with a number from 0 to 999 or an actual name. This kind of information had already been prepared for us, so we just use it as is.
2. **Resize Image.** Now open the *examples/imagenet/create_imagenet.sh* file, and make two changes: (1) set *RESIZE* to true if you have not resized the images, and (2) set the path for *TRAIN_DATA_ROOT* and *VAL_DATA_ROOT* so that Caffe would know where the ImageNet data resides. After executing this step, training and validation datasets would be inserted into the LevelDB database.
3. **Compute Image Mean.** Caffe requires that all image data be centered around the mean, so this step accomplishes that. Execute the command *./examples/imagenet/make_imagenet_mean.sh* and a file named *data/ilsvrc12/imagenet_mean.binaryproto* will be created.

4. **Model Definition.** This example attempts to mimic the work by Krizhevsky et al. as we introduced in Chapter 10. The file *models/bvlc_reference_caffenet/train_val.prototxt* describes the model, as shown in Listing C.16. Although it's quite lengthy, all layers should be familiar to you, so we would not repeat explaining them.
5. **Job Definition.** The file *models/bvlc_reference_caffenet/solver.prototxt* specifies how the training job should be carried out. Once again, remember to change `solver_mode` to CPU if you do not have a GPU installed on your machine.
6. **Kick off the training job.** When you are ready, simply kick off the training job by executing the command `./build/tools/caffe train --solver=models/bvlc_reference_caffenet/solver.prototxt`. However, for your reference, without a GPU, it would be slow. For example, on my MacBook Pro with an Intel i7 quad-core processor, it took ~4 minutes per 20 iterations, which is roughly 10x slower than on a K40 GPU. Listing C.18 shows a partial output of running this example on my machine. It is seen that at the end of the 50,000 iterations, training loss and test loss reached 1.4091 and 8.42039, respectively, while the test accuracy reached 0.10892 only, after running for 8684 minutes or about 6 days. This means that we do need GPUs for training deep learning models.

If you decide to develop your skills in applying CNN models to computer vision, delve into the internal implementations of Caffe or Caffe2. Your investment in your time will be paid off nicely.

Listing C.16 ImageNet AlexNet model definition file (*train_val.prototxt*, partial)

```

name: "CaffeNet"
layer {
  name: "data"
  type: "Data"
  top: "data"
  top: "label"
  include {
    phase: TRAIN
  }
  transform_param {
    mirror: true
    crop_size: 227
    mean_file:
"data/ilsvrc12/imagenet_mean.binarypro
to"
  }
  # mean pixel / channel-wise mean
  instead of mean image
  # transform_param {
  #   crop_size: 227
  #   mean_value: 104
  #   mean_value: 117
  #   mean_value: 123
  #   mirror: true
  # }

  data_param {
    source:
"examples/imagenet/ilsvrc12_train_lmdb
"
    batch_size: 256
    backend: LMDB
  }
  layer {
    name: "data"
    type: "Data"
    top: "data"
    top: "label"
    include {
      phase: TEST
    }
    transform_param {
      mirror: false
      crop_size: 227
      mean_file:
"data/ilsvrc12/imagenet_mean.binarypro
to"
    }
    # mean pixel / channel-wise mean
    instead of mean image
  }
}

```



```

gamma: 0.1
stepsize: 100000
display: 20
max_iter: 45000
momentum: 0.9
weight_decay: 0.0005
snapshot: 10000
snapshot_prefix: "models/bvlc_reference_caffenet/caffenet_train"
solver_mode: CPU

```

Listing C.18 Output of running the Caffe AlexNet job

```

I0324 20:06:34.952026 2506531648 layer_factory.hpp:77] Creating layer data
I0324 20:06:34.952397 2506531648 db_lmdb.cpp:35] Opened lmdb examples/imagenet/ilsvrc12_val_lmdb
.....
I0324 20:06:35.674441 2506531648 net.cpp:255] Network initialization done.
I0324 20:06:35.674546 2506531648 solver.cpp:57] Solver scaffolding done.
I0324 20:06:35.674772 2506531648 caffe.cpp:239] Starting Optimization
I0324 20:06:35.674787 2506531648 solver.cpp:293] Solving CaffeNet
I0324 20:06:35.674794 2506531648 solver.cpp:294] Learning Rate Policy: step
I0324 20:06:35.785261 2506531648 solver.cpp:351] Iteration 0, Testing net (#0)
I0324 20:23:23.643776 97710080 data_layer.cpp:73] Restarting data prefetching from start.
I0324 20:23:27.673923 2506531648 solver.cpp:418] Test net output #0: accuracy = 0.001
I0324 20:23:27.673967 2506531648 solver.cpp:418] Test net output #1: loss = 7.15056 (* 1 = 7.15056 loss)
I0324 20:23:41.583783 2506531648 solver.cpp:239] Iteration 0 (0 iter/s, 1025.91s/20 iters), loss = 7.60255
I0324 20:23:41.583819 2506531648 solver.cpp:258] Train net output #0: loss = 7.60255 (* 1 = 7.60255 loss)
I0324 20:23:41.583847 2506531648 sgd_solver.cpp:112] Iteration 0, lr = 0.01
I0324 20:27:48.385308 2506531648 solver.cpp:239] Iteration 20 (0.0810369 iter/s, 246.801s/20 iters), loss =
5.78311
I0324 20:27:48.385622 2506531648 solver.cpp:258] Train net output #0: loss = 5.78311 (* 1 = 5.78311 loss)
I0324 20:27:48.385632 2506531648 sgd_solver.cpp:112] Iteration 20, lr = 0.01
I0324 20:31:46.621083 2506531648 solver.cpp:239] Iteration 40 (0.0839507 iter/s, 238.235s/20 iters), loss =
5.56459
.....
I0324 22:53:46.022282 2506531648 solver.cpp:258] Train net output #0: loss = 4.17744 (* 1 = 4.17744 loss)
I0324 22:53:46.022291 2506531648 sgd_solver.cpp:112] Iteration 740, lr = 0.01
I0324 22:57:44.885599 2506531648 solver.cpp:239] Iteration 760 (0.08373 iter/s, 238.863s/20 iters), loss =
4.13973
I0324 22:57:44.885979 2506531648 solver.cpp:258] Train net output #0: loss = 4.13973 (* 1 = 4.13973 loss)
I0324 22:57:44.885989 2506531648 sgd_solver.cpp:112] Iteration 760, lr = 0.01
.....
I0330 20:31:20.443524 2506531648 solver.cpp:239] Iteration 49960 (0.101471 iter/s, 197.1s/20 iters), loss =
1.25496
I0330 20:31:20.445861 2506531648 solver.cpp:258] Train net output #0: loss = 1.25496 (* 1 = 1.25496 loss)
I0330 20:31:20.445873 2506531648 sgd_solver.cpp:112] Iteration 49960, lr = 0.01
I0330 20:34:37.205741 2506531648 solver.cpp:239] Iteration 49980 (0.101647 iter/s, 196.759s/20 iters), loss =
1.4091
I0330 20:34:37.206394 2506531648 solver.cpp:258] Train net output #0: loss = 1.4091 (* 1 = 1.4091 loss)

```

```
I0330 20:34:37.206403 2506531648 sgd_solver.cpp:112] Iteration 49980, lr = 0.01
I0330 20:37:44.142716 2506531648 solver.cpp:468] Snapshotting to binary proto file
models/bvlc_reference_caffenet/caffenet_train_iter_50000.caffemodel
I0330 20:37:45.423261 2506531648 sgd_solver.cpp:280] Snapshotting solver state to binary proto file
models/bvlc_reference_caffenet/caffenet_train_iter_50000.solverstate
I0330 20:37:50.101991 2506531648 solver.cpp:331] Iteration 50000, loss = 1.15807
I0330 20:37:50.102022 2506531648 solver.cpp:351] Iteration 50000, Testing net (#0)
I0330 20:51:07.974370 97710080 data_layer.cpp:73] Restarting data prefetching from start.
I0330 20:51:11.204300 2506531648 solver.cpp:418] Test net output #0: accuracy = 0.10892
I0330 20:51:11.204347 2506531648 solver.cpp:418] Test net output #1: loss = 8.42039 (* 1 = 8.42039 loss)
I0330 20:51:11.204352 2506531648 solver.cpp:336] Optimization Done.
I0330 20:51:11.207332 2506531648 caffe.cpp:250] Optimization Done.
```

```
real8684m38.099s
user 28888m26.225s
sys 416m16.676s
```

C.2 THE YOLOV3 FRAMEWORK

In this section, we focus on the YOLOv3 framework. This is my favorite, as it is written in C, highly-performing, and most of all, it offers many ways to build and run it, which is ideal for individual machine learning researchers.

Once again, since it is written in C, I'll show you how to build it from the source next to get started.

C.2.1 BUILDING THE YOLOV3 FRAMEWORK FROM THE SOURCE

You may want to watch the YouTube video at <https://www.youtube.com/watch?v=Cgxsv1riJhI> about how amazing YOLOv3 is. This perhaps can motivate you a bit on getting deep with the YOLOv3 framework. If so, let's begin with how to build YOLOv3 from the source next. You can check out YOLOv3 at <https://pjreddie.com/darknet/yolo/> for more information about this framework now or later.

YOLOv3 runs on an engine named *darknet*. The link at <https://pjreddie.com/darknet/install/> gives information on how to install *darknet*. In the section of *Compiling with OpenCV*, it mentions that by default, *darknet* uses *stb_image.h* for image loading, but *stb_image* does not support all image formats. Besides, OpenCV is a production-quality computer vision library, so I was interested in re-compiling *darknet* with OpenCV. However, when I followed the simple instructions given there for re-compiling *darknet* with OpenCV on my MacBook Pro, it did not work! It took me some substantial amount of time to rebuild YOLOv3 from the source on my macOS, which motivated me to summarize my experience here so that you don't have to go through all the difficulties I once had.

The steps to recompile *darknet* with OpenCV3 include:

1. Installing XCode
2. Installing Homebrew
3. Installing Python 3
4. Installing OpenCV 3 with Python bindings
5. Recompiling *darknet* with OpenCV3

If you already have 1-3 on your MacBook, you can skip to step 4. Otherwise, install 1-3 first, as instructed below.

C.2.1.1 INSTALLING XCODE

Get the latest version of XCode from the App Store and install it on your macOS machine. Then apply the developer license by executing the below command:

```
$sudo xcodebuild -license
```

Install the Command Line Tools by executing the below command:

```
$ sudo xcode-select --install
```

C.2.1.2 INSTALLING HOMEBREW

If you do not have Homebrew installed on your macOS machine, install it by executing the below command (all in one line):

```
$ ruby -e "$(curl -fsSL  
https://raw.githubusercontent.com/Homebrew/install/master/install)"
```

C.2.1.3 INSTALLING PYTHON3

If you do not have Python 3 installed on your macOS machine, install it with the below command:

```
$ brew install python3
```

To check the python version, execute the following command:

```
$python3 --version
```

I have Python 3.6.5 installed on my machine.

You can build YOLOv3 with STB (<https://github.com/nothings/stb>) just for simple tasks such as loading and saving images, but OpenCV gives you additional capability for displaying images. Performance-wise, they are about the same, but I like to install OpenCV on my MacBook to run YOLOv3 with that extra capability for displaying images. However, if you plan to upload your YOLOv3 bundle (binary + data) to a particular hosted env with no OpenCV support, then, building YOLOv3 with STB is the only option by setting `OPENCV = 0` in the `Makefile` that comes with the YOLOv3 download.

Next, I share with you how I installed OpenCV 3 on my MacBook, just in case you are interested as well.

C.2.1.4 INSTALLING OPENCV 3 WITH PYTHON BINDINGS

This is where you may run into difficulties. First of all, if you run the following command as instructed by many online blogs:

```
$brew tap homebrew/science
```

you may get the following:

```
$Error: homebrew/science was deprecated. This tap is now empty as all its
formulae were migrated.
```

So what do you do? Just ignore it.

Next, if you install `opencv3` as follows:

```
$brew install opencv3
```

you will get the latest OpenCV 3.4.1_2 installed. Then, when you change to the `darknet` directory and re-compile `darknet` by typing `make`, you will get the following error:

```
$In file included from ./src/gemm.c:2:
In file included from src/utils.h:5:
In file included from include/darknet.h:25:
In file included from /usr/local/Cellar/opencv/3.4.1_2/include/opencv2/highgui/highgui_c.h:45:
In file included from /usr/local/Cellar/opencv/3.4.1_2/include/opencv2/core/core_c.h:48:1_2
In file included from /usr/local/Cellar/opencv/3.4.1_2/include/opencv2/core/types_c.h:59:
/usr/local/Cellar/opencv/3.4.1_2/include/opencv2/core/cvdef.h:485:1: fatal error: unknown type name
'namespace'
namespace cv {
^
1 error generated.
make: *** [obj/gemm.o] Error 1
```

So what's wrong here? It only turned out that `opencv3.4.1` does not work with YOLOv3. We have to fall back to `opencv3.4.0`. I got YOLOv3 compiled successfully with `opencv3.4.0` by following the below procedure:

1. Install prerequisites for `opencv` by executing the following commands:

```
$brew install cmake pkg-config
$brew install jpeg libpng libtiff openexr
$brew install eigen tbb
```
2. Download `opencv 3.4.0` and `opencv_contrib 3.4.0` to a directory:

OpenCV3.4.0: <https://github.com/opencv/opencv/releases/tag/3.4.0>. Click on *Source code* (*tar.gz*).

OpenCV3.4.0 contrib: https://github.com/opencv/opencv_contrib/releases/tag/3.4.0.
3. Change to your `opencv-3.4.0` directory and execute the following three commands:

```
$mkdir build
$cd build
$cmake -D CMAKE_BUILD_TYPE=RELEASE -D CMAKE_INSTALL_PREFIX=/usr/local -D
OPENCV_EXTRA_MODULES_PATH=/Users/henryliu/mspc/devs/opencv_contrib-
3.4.0/modules -D
PYTHON3_LIBRARY=/usr/local/Cellar/python3/3.6.5/Frameworks/Python.framew
ork/Versions/3.6/lib/python3.6/config-3.6m/libpython3.6.dylib -D
PYTHON3_INCLUDE_DIR=/usr/local/Cellar/python3/3.6.5/Frameworks/Python.fr
amework/Versions/3.6/include/python3.6m/ -D BUILD_opencv_python2=OFF -D
BUILD_opencv_python3=ON -D INSTALL_PYTHON_EXAMPLES=ON -D
INSTALL_C_EXAMPLES=OFF -D BUILD_EXAMPLES=ON ..
```

Note that with the above *cmake* command, make sure you set `OPENCV_EXTRA_MODULES_PATH`, `PYTHON3_LIBRARY` and `PYTHON3_INCLUDE_DIR` to your own corresponding paths, respectively. At the end, you should see something similar to the following as I got on my machine:

```
-- Configuring done
-- Generating done
-- Build files have been written to: /Users/henryliu/mspc/devs/opencv-3.4.0/build
```

Next, execute the following two commands:

```
$ sudo make -j4
$ sudo make install
```

To verify that you have installed `opencv3.4.0` successfully, start up `python3` and issue the `import cv2` statement as shown below and you should get `'3.4.0'` as I got on my machine:

```
henryliu:build henryliu$ python3
Python 3.6.5 (default, Mar 30 2018, 06:42:10)
[GCC 4.2.1 Compatible Apple LLVM 9.0.0 (clang-900.0.39.2)] on darwin
Type "help", "copyright", "credits" or "license" for more information.
>>> import cv2
>>> cv2.__version__
'3.4.0'
>>>
```

C.2.1.5 BUILD YOLOv3

Now download the latest YOLOv3 source code from <https://github.com/pjreddie/darknet> and save it to a directory on your machine. Then, change to the *darknet* directory, edit the *Makefile* file to enable OPENCV by setting

```
OPENCV=1
```

Now type *make* and it should start re-compiling YOLOv3. After completion, execute the following command:

```
$ ./darknet imtest data/eagle.jpg
```

You should see images as shown below. This is an indication that you have successfully recompiled YOLOv3 on your macOS machine, as these images are supposed to be loaded by OpenCV.

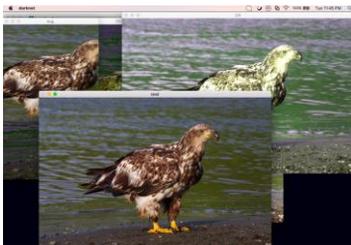


Figure C.3 Testing YOLOv3 recompiled with OpenCV3.4.0.

C.2.2 ALEX'S CIFAR-10 WITH YOLOv3

If you did not skip §C.1.3, you should already know Alex's work with CIFAR-10. Now. Let's follow <https://pjreddie.com/darknet/train-cifar/> to train a CNN model with YOLOv3 using the CIFAR-10 dataset. The steps include:

1. Getting the CIFAR dataset
2. Making a data file to define the job
3. Making a network config file to define the net
4. Training the model

Let's follow these steps to train a classifier.

C.2.2.1 GETTING THE CIFAR DATASET

To get the CIFAR dataset, change to the *darknet* directory and run the following commands:

```
$cd data
$wget https://pjreddie.com/media/files/cifar.tgz
$tar xzf cifar.tgz
```

After the above step, you should have the directories of *train* and *test* as well as a file named *labels.txt*. You can check them out by executing the following commands:

```
henryliu:cifar henryliu$ ls train | head -5
0_frog.png
10000_automobile.png
10001_frog.png
10002_frog.png
10003_ship.png
henryliu:cifar henryliu$ ls train | wc -l
50000
henryliu:cifar henryliu$ ls test | wc -l
10000
henryliu:cifar henryliu$ cat labels.txt
airplane
automobile
bird
cat
deer
dog
frog
horse
ship
truck
```

The *train* directory contains 50k image files in PNG format, while the *test* directory contains 10k images for testing. The *labels.txt* file contains the 10 classes as shown above that those images belong to.

Next, execute the following commands in the *cifar* directory to create the path files for the training and testing datasets, respectively:

```

$find `pwd`/train -name \*.png > train.list
$find `pwd`/test -name \*.png > test.list
henryliu:data henryliu$ head -5 cifar/train.list
/Users/henryliu/mspc/devs/ws_cpp/darknet/data/cifar/train/0_frog.png
/Users/henryliu/mspc/devs/ws_cpp/darknet/data/cifar/train/10000_automobile.png
/Users/henryliu/mspc/devs/ws_cpp/darknet/data/cifar/train/10001_frog.png
/Users/henryliu/mspc/devs/ws_cpp/darknet/data/cifar/train/10002_frog.png
/Users/henryliu/mspc/devs/ws_cpp/darknet/data/cifar/train/10003_ship.png
henryliu:data henryliu$ head -5 cifar/test.list
/Users/henryliu/mspc/devs/ws_cpp/darknet/data/cifar/test/0_cat.png
/Users/henryliu/mspc/devs/ws_cpp/darknet/data/cifar/test/1000_dog.png
/Users/henryliu/mspc/devs/ws_cpp/darknet/data/cifar/test/1001_airplane.png
/Users/henryliu/mspc/devs/ws_cpp/darknet/data/cifar/test/1002_ship.png
/Users/henryliu/mspc/devs/ws_cpp/darknet/data/cifar/test/1003_deer.png

```

 **Note:** If you need to run the CIFAR job somewhere else, use the relative path in place of the absolute path for each image file by removing the part preceding `data/cifar/...`, e.g., `data/cifar/test/1003_deer.png`, etc., in the `train.list` and `test.list` files. Otherwise, you will get “*file not found*” errors.

Next, create the job definition file.

C.2.2.2 MAKING A DATA FILE TO DEFINE THE JOB

Now, check out or create a `cifar.data` file in the `darknet/cfg` directory with the following contents:

```

classes=10
train = data/cifar/train.list
valid = data/cifar/test.list
labels = data/cifar/labels.txt
backup = backup
top=2

```

Since the `backup` directory does not exist yet, you need to create it yourself now. Then, familiarize yourself with the meaning of each line as follows:

- **classes=10:** the number of unique classes that all images belong to
- **train:** The file that contains the absolute path of each training image file, including the file name
- **valid:** The file that contains the absolute path of each validation/test image file, including the file name
- **labels:** The file containing a list of all possible classes by name
- **backup:** The directory for saving backup weights during training
- **top = 2:** # of top-*n* classes to classify at test time (in addition to top-1)

Next, define the network configuration for training the model.

C.2.2.3 MAKING A NETWORK CONFIG FILE TO DEFINE THE NET

To define the model for training, the file *cfg/cifar_small.cfg* is composed with the following contents:

```
[net]
batch=128
subdivisions=1
height=28
width=28
channels=3
max_crop=32
min_crop=32

hue=.1
saturation=.75
exposure=.75

learning_rate=0.1
policy=poly
power=4
max_batches = 5000
momentum=0.9
decay=0.0005

[convolutional]
batch_normalize=1
filters=64
size=3
stride=1
pad=1
activation=leaky

[maxpool]
size=2
stride=2

[convolutional]
batch_normalize=1
filters=128
size=3
stride=1
pad=1
activation=leaky

[convolutional]
filters=10
size=1
stride=1
pad=1
activation=leaky

[avgpool]

[softmax]
groups=1

[convolutional]
type=sse
```

Note that YOLO uses the `max_batches` to define the # of maximum iterations. You can change it to a smaller number, say, 500, just to make sure that it runs. The other parameters should be obvious, given what you have learnt from the main text. The link at <https://pjreddie.com/darknet/train-cifar/> has a brief description about the model.

Next, let's see how we can train this model with YOLOv3.

C.2.2.4 TRAINING THE MODEL

To train the model, just launch it with the following command:

```
./darknet classifier train cfg/cifar.data cfg/cifar_small.cfg
```

The command to restart the training with a backup file located in the `backup` directory is:

```
./darknet classifier train cfg/cifar.data cfg/cifar_small.cfg backup/cifar_small.backup
```

The instructions at <https://pjreddie.com/darknet/train-cifar/> stopped here, so I would take over and share with you what I got on my macOS machine.

I first set `max_batches` to 512 and ended up with the following output, which took about six minutes:

```
henryliu:darknet henryliu$ ./darknet classifier train cfg/cifar.data cfg/cifar_small.cfg
cifar_small
1
layer  filters  size      input       output
 0 conv   32 3x3 / 1  28x 28x 3  -> 28x 28x 32 0.001 BFLOPs
 1 max     2x2 / 2  28x 28x 32  -> 14x 14x 32
 2 conv   64 3x3 / 1  14x 14x 32  -> 14x 14x 64 0.007 BFLOPs
 3 max     2x2 / 2  14x 14x 64  -> 7x 7x 64
 4 conv  128 3x3 / 1   7x 7x 64  -> 7x 7x 128 0.007 BFLOPs
 5 conv   10 1x1 / 1   7x 7x 128  -> 7x 7x 10 0.000 BFLOPs
 6 avg                7x 7x 10  -> 10
 7 softmax                10
 8 cost                  10
Learning Rate: 0.1, Momentum: 0.9, Decay: 0.0005
50000
32 32
1, 0.003: 1.628222, 1.628222 avg, 0.099221 rate, 0.789039 seconds, 128 images, 08-28-2018 19:52:09.000
2, 0.005: 1.607311, 1.626131 avg, 0.098447 rate, 0.692393 seconds, 256 images, 08-28-2018 19:52:09.000
3, 0.008: 1.580675, 1.621585 avg, 0.097677 rate, 0.790754 seconds, 384 images, 08-28-2018 19:52:10.000
...
510, 1.306: 0.966246, 1.016550 avg, 0.000000 rate, 0.732217 seconds, 65280 images, 08-28-2018 19:58:28.000
511, 1.308: 1.038416, 1.018736 avg, 0.000000 rate, 0.762059 seconds, 65408 images, 08-28-2018 19:58:28.000
512, 1.311: 0.986644, 1.015527 avg, 0.000000 rate, 0.711117 seconds, 65536 images, 08-28-2018 19:58:29.000
Saving weights to backup/cifar_small.weights
```

The last line of **512, 1.311: 0.986644, 1.015527 avg, 0.000000 rate, 0.711117 seconds, 65536 images, 08-28-2018 19:58:29.000** represents the iteration, total loss: current loss, average loss so far, current learning rate, time taken for this iteration, the number of images processed so far and timestamp that was added by me in the source code. Notice the following:

- The learning rate started with 0.099221 at iteration 1 and ended up with 0.000000 at iteration 512.
- The number of images started with 128 at iteration 1 and ended up with 65536 at iteration 512. This is because each iteration uses 128 images, so $128 \times 512 = 65536$. This also means that after $50000/128 = 390$ iterations, each image would have been used at least once.
- The avg loss started with 1.628222 and ended with 1.015527 after 512 iterations over a duration of 6m20s. Besides, each iteration took about 0.7 seconds, or $128 \text{ images}/0.7\text{s} = 183 \text{ images per second}$, which may look great, but actually not, as this is a fairly simple and small example.

As you may have realized, this is an image classification ML example, which means that a single-object is given to the model, which predicts what the image might be. Next, I'll show you how good this simple model is by giving it an image from the YOLOv3 download as shown in Figure C.4.



Figure C.4 A picture containing multiple objects.

Here is what the model predicted when the following command was executed on my MacBook:

```
henryliu:cifar henryliu$ ./darknet_mac_no_accel classifier predict cfg/cifar.data cfg/cifar_small.cfg  
backup/cifar_small_512.weights data/dog.jpg
```

.....

Loading weights from backup/cifar_small_512.weights...Done!

data/dog.jpg: Predicted in 0.003164 seconds.

17.42%: automobile

16.94%: frog

The model trained predicted that this picture might be an automobile with a probability of 17.42% or a frog with a probability of 16.94%. Neither is true, though. However, partially it's my fault that I gave a multi-object image to a single-image classification model to predict what the object was. To be fair, I made the following three pictures with one object per picture and fed them to the model one by one and see what the model would predict.

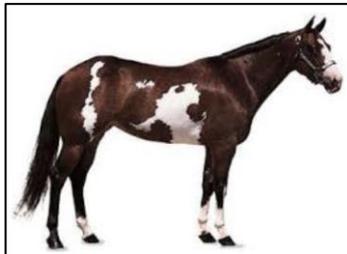


Figure C.5 Three separate images with one object per picture.

And these were what the model predicted now:

```
henryliu:cifar henryliu$ ./darknet_mac_no_accel classifier predict cfg/cifar.data cfg/cifar_small.cfg
backup/cifar_small_512.weights data/dog-1.png
```

```
.....
data/dog-1.png: Predicted in 0.002243 seconds.
```

```
33.32%: dog
```

```
23.08%: cat
```

```
henryliu:cifar henryliu$ ./darknet_mac_no_accel classifier predict cfg/cifar.data cfg/cifar_small.cfg
backup/cifar_small_512.weights data/horse-1.png
```

```
.....
data/horse-1.png: Predicted in 0.002402 seconds.
```

```
36.22%: airplane
```

```
21.94%: automobile
```

```
henryliu:cifar henryliu$ ./darknet_mac_no_accel classifier predict cfg/cifar.data cfg/cifar_small.cfg
backup/cifar_small_512.weights data/car-1.png
```

```
.....
data/car-1.png: Predicted in 0.002592 seconds.
```

```
42.72%: automobile
```

```
33.49%: airplane
```

Namely, the model predicted that the dog was 33.32% likely to be a dog and 23.08% likely to be a cat, the horse was 36.22% likely to be an airplane and 21.94% likely to be an automobile, and the car was 42.72% likely to be an automobile and 33.49% likely to be an airplane. The model was not sufficiently accurate, but still not too bad, given that it was trained for about 6 minutes only.

Next I changed `max_batches` from 512 back to 5000 and got the loss down from 1.01 to ~0.58 after ~63 minutes. I ran the same tests and got:

```
henryliu:cifar henryliu$ ./darknet_mac_accel classifier predict cfg/cifar.data cfg/cifar_small.cfg
backup/cifar_small_5000.weights data/dog.jpg
```

```
.....
data/dog.jpg: Predicted in 0.002281 seconds.
```

```
81.79%: ship
```

```
10.94%: airplane
```

```
henryliu:cifar henryliu$ ./darknet_mac_accel classifier predict cfg/cifar.data cfg/cifar_small.cfg
backup/cifar_small_5000.weights data/dog-1.png
```

```
.....
data/dog-1.png: Predicted in 0.001843 seconds.
```

```
53.68%: horse
```

```
24.24%: cat
```

```
henryliu:cifar henryliu$ ./darknet_mac_accel classifier predict cfg/cifar.data cfg/cifar_small.cfg
backup/cifar_small_5000.weights data/horse-1.png
```

```
.....
data/horse-1.png: Predicted in 0.002292 seconds.
```

```
71.48%: horse
```

```
28.38%: airplane
```

```
henryliu:cifar henryliu$ ./darknet_mac_accel classifier predict cfg/cifar.data cfg/cifar_small.cfg
backup/cifar_small_5000.weights data/car-1.png
```

```
.....
data/car-1.png: Predicted in 0.002093 seconds.
```

```
67.99%: automobile
```

15.19%: ship

Finally I changed `max_batches` from 5000 to 105000 and got the loss down from ~0.58 to ~0.24, using a version of YOLOv3 I optimized on my MacBook, which is about 3x faster for this particular small CIFAR-10 model. I'll share the details of the optimization later so that you can run a much larger dataset named COCO on your Mac machine by roughly 20x faster as I did. But for now, here are the same tests and the results with this much better trained model:

```
henryliu:cifar henryliu$ ./darknet_mac_accel classifier predict cfg/cifar.data cfg/cifar_small.cfg
backup/cifar_small_105000.weights data/dog.jpg
```

.....

data/dog.jpg: Predicted in 0.002316 seconds.

79.54%: airplane

16.67%: cat

```
henryliu:cifar henryliu$ ./darknet_mac_accel classifier predict cfg/cifar.data cfg/cifar_small.cfg
backup/cifar_small_105000.weights data/dog-1.png
```

.....

data/dog-1.png: Predicted in 0.002554 seconds.

99.91%: dog

0.05%: horse

```
henryliu:cifar henryliu$ ./darknet_mac_accel classifier predict cfg/cifar.data cfg/cifar_small.cfg
backup/cifar_small_105000.weights data/horse-1.png
```

.....

data/horse-1.png: Predicted in 0.002598 seconds.

74.49%: horse

14.32%: airplane

```
henryliu:cifar henryliu$ ./darknet_mac_accel classifier predict cfg/cifar.data cfg/cifar_small.cfg
backup/cifar_small_105000.weights data/car-1.png
```

data/car-1.png: Predicted in 0.002157 seconds.

95.72%: automobile

2.68%: truck

Now, except the first “*all-in-one*” picture test, all single-object pictures have been predicted correctly, with a dog being 99.91% a dog, a horse 74.49% a horse, and a car 95.72% an automobile, respectively. In fact, I made another run with 205000 max-batches and the loss was reduced from 0.24 to 0.20 only, which would not help much further.

Finally, I tried to test/validate the model trained above using the validation/test dataset, and got the results as shown in Listing C.19. Since the CIFAR data falls into the 10 classes of *airplane* – 0, *automobile* – 1, *bird* – 2, *cat* – 3, *deer* – 4, *dog* – 5, *frog* – 6, *horse* – 7, *ship* – 8, and *truck* 9, for a given test image, the model attempts to predict the probabilities of being in the top 2 classes of *airplane* – 0 and *automobile* – 1. For example, the line of `data/cifar/test/1001_airplane.png, 0, 0.959015, 0.002087` means that the airplane picture is 95.9% an airplane and 0.2087% an automobile, and so forth. Overall, this model trained with 105000 batches seems to be very accurate.

Next we explore YOLO’s another use for detecting multiple objects in an image or bounding boxes.

Listing C.19 Validation runs with the CIFAR small configuration model trained with 105000 iterations

```
henryliu:cifar henryliu$ ./darknet_mac_accel classifier valid cfg/cifar.data cfg/cifar_small.cfg
backup/cifar_small_105000.weights
```

layer	filters	size	input	output
0 conv	32	3 x 3 / 1	28 x 28 x 3	-> 28 x 28 x 32 0.001 BFLOPs
1 max	2 x 2 / 2		28 x 28 x 32	-> 14 x 14 x 32
2 conv	64	3 x 3 / 1	14 x 14 x 32	-> 14 x 14 x 64 0.007 BFLOPs
3 max	2 x 2 / 2		14 x 14 x 64	-> 7 x 7 x 64
4 conv	128	3 x 3 / 1	7 x 7 x 64	-> 7 x 7 x 128 0.007 BFLOPs
5 conv	10	1 x 1 / 1	7 x 7 x 128	-> 7 x 7 x 10 0.000 BFLOPs
6 avg			7 x 7 x 10	-> 10
7 softmax				10
8 cost				10

```
Loading weights from backup/cifar_small_105000.weights...Done!
```

```
data/cifar/test/0_cat.png, 3, 0.002164, 0.003230,
```

```
0: top 1: 1.000000, top 2: 1.000000
```

```
data/cifar/test/1000_dog.png, 5, 0.000000, 0.000000,
```

```
1: top 1: 1.000000, top 2: 1.000000
```

```
data/cifar/test/1001_airplane.png, 0, 0.959015, 0.002087,
```

```
2: top 1: 1.000000, top 2: 1.000000
```

```
data/cifar/test/1002_ship.png, 8, 0.007687, 0.162669,
```

```
3: top 1: 1.000000, top 2: 1.000000
```

```
data/cifar/test/1003_deer.png, 4, 0.000001, 0.000000,
```

```
4: top 1: 1.000000, top 2: 1.000000
```

```
data/cifar/test/1004_ship.png, 8, 0.010149, 0.000043,
```

```
5: top 1: 1.000000, top 2: 1.000000
```

```
data/cifar/test/1005_automobile.png, 1, 0.000031, 0.922122,
```

```
6: top 1: 1.000000, top 2: 1.000000
```

```
data/cifar/test/1006_automobile.png, 1, 0.000000, 0.999992,
```

```
7: top 1: 1.000000, top 2: 1.000000
```

```
data/cifar/test/1007_ship.png, 8, 0.066883, 0.002106,
```

```
8: top 1: 1.000000, top 2: 1.000000
```

```
data/cifar/test/1008_truck.png, 9, 0.236106, 0.022653,
```

```
9: top 1: 0.900000, top 2: 0.900000
```

```
data/cifar/test/1009_frog.png, 6, 0.000022, 0.000192,
```

```
10: top 1: 0.909091, top 2: 0.909091
```

```
data/cifar/test/100_deer.png, 4, 0.000003, 0.000001,
```

```
11: top 1: 0.916667, top 2: 0.916667
```

```
data/cifar/test/1010_airplane.png, 0, 0.999245, 0.000001,
```

```
12: top 1: 0.923077, top 2: 0.923077
```

```
.....
```

C.2.3 USE YOLOv3 TO DETECT MULTI-OBJECTS IN AN IMAGE OR BOUNDING BOXES

From YOLO's main web page at <https://pjreddie.com/darknet/yolo/>, you can find a section about detecting bounding boxes using YOLO's pre-trained model. It starts with getting *darknet*, but you already have it if you followed the previous instructions and re-compiled YOLOv3 with OpenCV3. Then, you can directly go to the next step of retrieving the pre-trained YOLO weights as follows:

```
$ wget https://pjreddie.com/media/files/yolov3.weights
```

Then, execute the following command to detect objects in the picture:

```
$ ./darknet detect cfg/yolov3.cfg yolov3.weights data/dog.jpg
```

The output from the above command should look similar to the following (the list is a bit lengthy, but it gives all details layer by layer about how this model is composed exactly):

```
henryliu:darknet henryliu$ ./darknet detect cfg/yolov3.cfg yolov3.weights data/dog.jpg
```

layer	filters	size	input	output	
0	conv	32 3 x 3 / 1	416 x 416 x 3	-> 416 x 416 x 32	0.299 BFLOPs
1	conv	64 3 x 3 / 2	416 x 416 x 32	-> 208 x 208 x 64	1.595 BFLOPs
2	conv	32 1 x 1 / 1	208 x 208 x 64	-> 208 x 208 x 32	0.177 BFLOPs
3	conv	64 3 x 3 / 1	208 x 208 x 32	-> 208 x 208 x 64	1.595 BFLOPs
4	res	1	208 x 208 x 64	-> 208 x 208 x 64	
5	conv	128 3 x 3 / 2	208 x 208 x 64	-> 104 x 104 x 128	1.595 BFLOPs
6	conv	64 1 x 1 / 1	104 x 104 x 128	-> 104 x 104 x 64	0.177 BFLOPs
7	conv	128 3 x 3 / 1	104 x 104 x 64	-> 104 x 104 x 128	1.595 BFLOPs
8	res	5	104 x 104 x 128	-> 104 x 104 x 128	
9	conv	64 1 x 1 / 1	104 x 104 x 128	-> 104 x 104 x 64	0.177 BFLOPs
10	conv	128 3 x 3 / 1	104 x 104 x 64	-> 104 x 104 x 128	1.595 BFLOPs
11	res	8	104 x 104 x 128	-> 104 x 104 x 128	
12	conv	256 3 x 3 / 2	104 x 104 x 128	-> 52 x 52 x 256	1.595 BFLOPs
13	conv	128 1 x 1 / 1	52 x 52 x 256	-> 52 x 52 x 128	0.177 BFLOPs
14	conv	256 3 x 3 / 1	52 x 52 x 128	-> 52 x 52 x 256	1.595 BFLOPs
15	res	12	52 x 52 x 256	-> 52 x 52 x 256	
16	conv	128 1 x 1 / 1	52 x 52 x 256	-> 52 x 52 x 128	0.177 BFLOPs
17	conv	256 3 x 3 / 1	52 x 52 x 128	-> 52 x 52 x 256	1.595 BFLOPs
18	res	15	52 x 52 x 256	-> 52 x 52 x 256	
19	conv	128 1 x 1 / 1	52 x 52 x 256	-> 52 x 52 x 128	0.177 BFLOPs
20	conv	256 3 x 3 / 1	52 x 52 x 128	-> 52 x 52 x 256	1.595 BFLOPs
21	res	18	52 x 52 x 256	-> 52 x 52 x 256	
22	conv	128 1 x 1 / 1	52 x 52 x 256	-> 52 x 52 x 128	0.177 BFLOPs
23	conv	256 3 x 3 / 1	52 x 52 x 128	-> 52 x 52 x 256	1.595 BFLOPs
24	res	21	52 x 52 x 256	-> 52 x 52 x 256	
25	conv	128 1 x 1 / 1	52 x 52 x 256	-> 52 x 52 x 128	0.177 BFLOPs
26	conv	256 3 x 3 / 1	52 x 52 x 128	-> 52 x 52 x 256	1.595 BFLOPs
27	res	24	52 x 52 x 256	-> 52 x 52 x 256	
28	conv	128 1 x 1 / 1	52 x 52 x 256	-> 52 x 52 x 128	0.177 BFLOPs
29	conv	256 3 x 3 / 1	52 x 52 x 128	-> 52 x 52 x 256	1.595 BFLOPs
30	res	27	52 x 52 x 256	-> 52 x 52 x 256	
31	conv	128 1 x 1 / 1	52 x 52 x 256	-> 52 x 52 x 128	0.177 BFLOPs
32	conv	256 3 x 3 / 1	52 x 52 x 128	-> 52 x 52 x 256	1.595 BFLOPs
33	res	30	52 x 52 x 256	-> 52 x 52 x 256	
34	conv	128 1 x 1 / 1	52 x 52 x 256	-> 52 x 52 x 128	0.177 BFLOPs
35	conv	256 3 x 3 / 1	52 x 52 x 128	-> 52 x 52 x 256	1.595 BFLOPs
36	res	33	52 x 52 x 256	-> 52 x 52 x 256	
37	conv	512 3 x 3 / 2	52 x 52 x 256	-> 26 x 26 x 512	1.595 BFLOPs
38	conv	256 1 x 1 / 1	26 x 26 x 512	-> 26 x 26 x 256	0.177 BFLOPs
39	conv	512 3 x 3 / 1	26 x 26 x 256	-> 26 x 26 x 512	1.595 BFLOPs
40	res	37	26 x 26 x 512	-> 26 x 26 x 512	

41 conv 256 1 x 1 / 1 26 x 26 x 512 -> 26 x 26 x 256 0.177 BFLOPs
 42 conv 512 3 x 3 / 1 26 x 26 x 256 -> 26 x 26 x 512 1.595 BFLOPs
 43 res 40 26 x 26 x 512 -> 26 x 26 x 512
 44 conv 256 1 x 1 / 1 26 x 26 x 512 -> 26 x 26 x 256 0.177 BFLOPs
 45 conv 512 3 x 3 / 1 26 x 26 x 256 -> 26 x 26 x 512 1.595 BFLOPs
 46 res 43 26 x 26 x 512 -> 26 x 26 x 512
 47 conv 256 1 x 1 / 1 26 x 26 x 512 -> 26 x 26 x 256 0.177 BFLOPs
 48 conv 512 3 x 3 / 1 26 x 26 x 256 -> 26 x 26 x 512 1.595 BFLOPs
 49 res 46 26 x 26 x 512 -> 26 x 26 x 512
 50 conv 256 1 x 1 / 1 26 x 26 x 512 -> 26 x 26 x 256 0.177 BFLOPs
 51 conv 512 3 x 3 / 1 26 x 26 x 256 -> 26 x 26 x 512 1.595 BFLOPs
 52 res 49 26 x 26 x 512 -> 26 x 26 x 512
 53 conv 256 1 x 1 / 1 26 x 26 x 512 -> 26 x 26 x 256 0.177 BFLOPs
 54 conv 512 3 x 3 / 1 26 x 26 x 256 -> 26 x 26 x 512 1.595 BFLOPs
 55 res 52 26 x 26 x 512 -> 26 x 26 x 512
 56 conv 256 1 x 1 / 1 26 x 26 x 512 -> 26 x 26 x 256 0.177 BFLOPs
 57 conv 512 3 x 3 / 1 26 x 26 x 256 -> 26 x 26 x 512 1.595 BFLOPs
 58 res 55 26 x 26 x 512 -> 26 x 26 x 512
 59 conv 256 1 x 1 / 1 26 x 26 x 512 -> 26 x 26 x 256 0.177 BFLOPs
 60 conv 512 3 x 3 / 1 26 x 26 x 256 -> 26 x 26 x 512 1.595 BFLOPs
 61 res 58 26 x 26 x 512 -> 26 x 26 x 512
 62 conv 1024 3 x 3 / 2 26 x 26 x 512 -> 13 x 13 x 1024 1.595 BFLOPs
 63 conv 512 1 x 1 / 1 13 x 13 x 1024 -> 13 x 13 x 512 0.177 BFLOPs
 64 conv 1024 3 x 3 / 1 13 x 13 x 512 -> 13 x 13 x 1024 1.595 BFLOPs
 65 res 62 13 x 13 x 1024 -> 13 x 13 x 1024
 66 conv 512 1 x 1 / 1 13 x 13 x 1024 -> 13 x 13 x 512 0.177 BFLOPs
 67 conv 1024 3 x 3 / 1 13 x 13 x 512 -> 13 x 13 x 1024 1.595 BFLOPs
 68 res 65 13 x 13 x 1024 -> 13 x 13 x 1024
 69 conv 512 1 x 1 / 1 13 x 13 x 1024 -> 13 x 13 x 512 0.177 BFLOPs
 70 conv 1024 3 x 3 / 1 13 x 13 x 512 -> 13 x 13 x 1024 1.595 BFLOPs
 71 res 68 13 x 13 x 1024 -> 13 x 13 x 1024
 72 conv 512 1 x 1 / 1 13 x 13 x 1024 -> 13 x 13 x 512 0.177 BFLOPs
 73 conv 1024 3 x 3 / 1 13 x 13 x 512 -> 13 x 13 x 1024 1.595 BFLOPs
 74 res 71 13 x 13 x 1024 -> 13 x 13 x 1024
 75 conv 512 1 x 1 / 1 13 x 13 x 1024 -> 13 x 13 x 512 0.177 BFLOPs
 76 conv 1024 3 x 3 / 1 13 x 13 x 512 -> 13 x 13 x 1024 1.595 BFLOPs
 77 conv 512 1 x 1 / 1 13 x 13 x 1024 -> 13 x 13 x 512 0.177 BFLOPs
 78 conv 1024 3 x 3 / 1 13 x 13 x 512 -> 13 x 13 x 1024 1.595 BFLOPs
 79 conv 512 1 x 1 / 1 13 x 13 x 1024 -> 13 x 13 x 512 0.177 BFLOPs
 80 conv 1024 3 x 3 / 1 13 x 13 x 512 -> 13 x 13 x 1024 1.595 BFLOPs
 81 conv 255 1 x 1 / 1 13 x 13 x 1024 -> 13 x 13 x 255 0.088 BFLOPs
 82 detection
 83 route 79
 84 conv 256 1 x 1 / 1 13 x 13 x 512 -> 13 x 13 x 256 0.044 BFLOPs
 85 upsample 2x 13 x 13 x 256 -> 26 x 26 x 256
 86 route 85 61
 87 conv 256 1 x 1 / 1 26 x 26 x 768 -> 26 x 26 x 256 0.266 BFLOPs
 88 conv 512 3 x 3 / 1 26 x 26 x 256 -> 26 x 26 x 512 1.595 BFLOPs
 89 conv 256 1 x 1 / 1 26 x 26 x 512 -> 26 x 26 x 256 0.177 BFLOPs

```

90 conv  512 3 x 3 / 1  26 x 26 x 256 -> 26 x 26 x 512 1.595 BFLOPs
91 conv  256 1 x 1 / 1  26 x 26 x 512 -> 26 x 26 x 256 0.177 BFLOPs
92 conv  512 3 x 3 / 1  26 x 26 x 256 -> 26 x 26 x 512 1.595 BFLOPs
93 conv  255 1 x 1 / 1  26 x 26 x 512 -> 26 x 26 x 255 0.177 BFLOPs
94 detection
95 route 91
96 conv  128 1 x 1 / 1  26 x 26 x 256 -> 26 x 26 x 128 0.044 BFLOPs
97 upsample      2x 26 x 26 x 128 -> 52 x 52 x 128
98 route 97 36
99 conv  128 1 x 1 / 1  52 x 52 x 384 -> 52 x 52 x 128 0.266 BFLOPs
100 conv  256 3 x 3 / 1  52 x 52 x 128 -> 52 x 52 x 256 1.595 BFLOPs
101 conv  128 1 x 1 / 1  52 x 52 x 256 -> 52 x 52 x 128 0.177 BFLOPs
102 conv  256 3 x 3 / 1  52 x 52 x 128 -> 52 x 52 x 256 1.595 BFLOPs
103 conv  128 1 x 1 / 1  52 x 52 x 256 -> 52 x 52 x 128 0.177 BFLOPs
104 conv  256 3 x 3 / 1  52 x 52 x 128 -> 52 x 52 x 256 1.595 BFLOPs
105 conv  255 1 x 1 / 1  52 x 52 x 256 -> 52 x 52 x 255 0.353 BFLOPs
106 detection

```

Loading weights from yolov3.weights...Done!

data/dog.jpg: Predicted in 7.581085 seconds.

truck: 93%

bicycle: 99%

dog: 99%

The above output indicates that it took 7.581 seconds and predicted a truck, a bicycle and a dog with the probabilities of 93%, 99% and 99%, respectively. If you open the *projections.png* file in the *darknet* directory, you should see those bounding boxes predicted as shown in Fig. C.6. You can also locate this image on the Dock by clicking on the Terminal icon labeled *darknet* as shown in Fig. C.6.

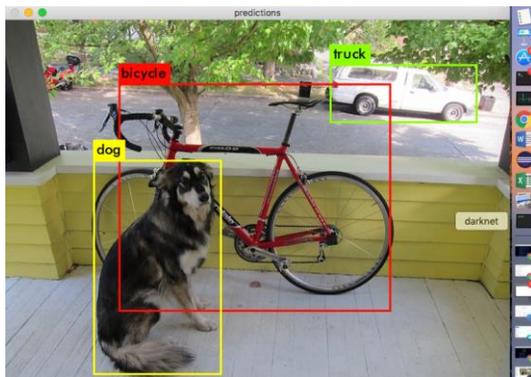


Figure C.6 Bounding boxes predicted by YOLOv3 with its own pre-trained weights.

Now you can press Ctrl-C to end the session. You can try another image, e.g., *data/horses.jpg*, and you would get the output as shown below, showing that four horses have been detected, as shown in Fig. C.7:

data/horses.jpg: Predicted in 7.799737 seconds.

horse: 98%

horse: 97%
horse: 91%
horse: 89%

It is amazing that YOLO can easily tell what objects and how many are in a picture!

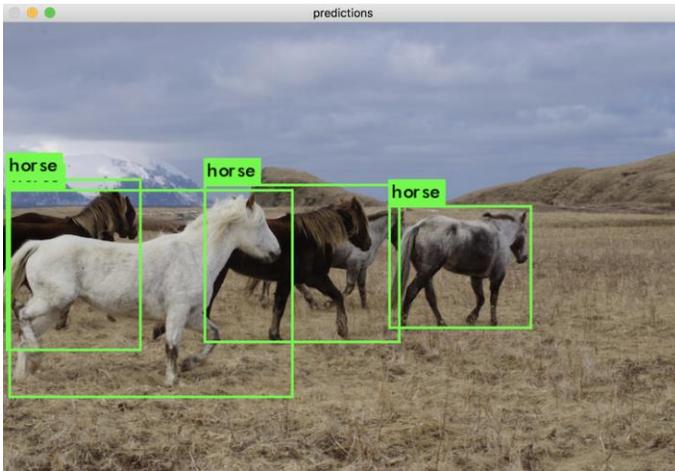


Figure C.7 Four horses detected by YOLO with its own pre-trained weights.

You can also try YOLO with live videos from a webcam as described there. I'll leave this to you, though.

C.2.4 TRAINING YOLO ON THE COCO DATASET

To train YOLO on the 2014 COCO dataset, check out this paper <https://arxiv.org/pdf/1405.0312.pdf> to learn a bit more about the COCO dataset first. Then, download the 2014 COCO dataset directly from COCO's download site at <http://cocodataset.org/#download>, which is much faster than from YOLO's website. After downloading 2014 COCO dataset zip files, create a `data/coco/images` sub-directory in the `darknet` directory, and place the 2014 COCO zip files there. Then, follow the instructions at YOLO's main page at <https://pjreddie.com/darknet/yolo/> under the section titled *Training YOLO on COCO*, except that you need to execute the following command:

```
$ cp scripts/get_coco_dataset.sh data
```

Then, make some changes in the `get_coco_dataset.sh` file as shown in Listing C.19. The commands I executed next were:

```
$ cd data  
$ bash get_coco_dataset.sh
```

The `get_coco_dataset.sh` script, shown in Listing C.19, explains what this script does, as a good example for how to retrieve dataset and prepare the data. Note that downloading the COCO dataset may take many hours, depending on the Internet speed you have with your machine. In my case, downloading the

train/val/test data concurrently took me about two hours at home with a download speed of up to 14 MB/s with a direct Ethernet cable connection, as shown in Fig. C.8.

Listing C.19 `get_coco_dataset.sh` (those marked red were modified)

```
#!/bin/bash

# Clone COCO API
#git clone https://github.com/pdollar/coco
cd coco

#mkdir images
#cd images

# Download Images
#wget -c https://pjreddie.com/media/files/train2014.zip
#wget -c https://pjreddie.com/media/files/val2014.zip

# Unzip
#unzip -q train2014.zip
#unzip -q val2014.zip
#unzip -q test2014.zip

#cd ..

# Download COCO Metadata
wget -c https://pjreddie.com/media/files/instances_train-val2014.zip
wget -c https://pjreddie.com/media/files/coco/5k.part
wget -c https://pjreddie.com/media/files/coco/trainvalno5k.part
wget -c https://pjreddie.com/media/files/coco/labels.tgz
tar xzf labels.tgz
unzip -q instances_train-val2014.zip

# Set Up Image Lists
paste <(awk "{print \"\$PWD\"}" <5k.part) 5k.part | tr -d '\t' > 5k.txt
paste <(awk "{print \"\$PWD\"}" <trainvalno5k.part) trainvalno5k.part | tr -d '\t' > trainvalno5k.txt
```

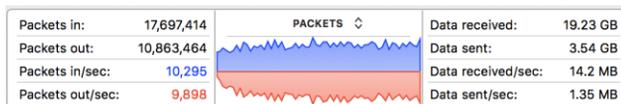


Figure C.8 COCO dataset download speed at home, directly from COCO’s download site.

Once again, you should use relative path for each image file in the `5k.txt` and `trainvalno5k.txt` files if you plan to copy and run COCO on a different environment.

Then, I modified the `cfg/coco.data` file to have the following contents:

```

classes= 80
train = data/coco/trainvalno5k.txt
valid = data/coco/5k.txt
names = data/coco.names
backup = backup

```

The `cfg/yolov3.cfg` was modified for training as follows, as marked in red (`max_batches` was changed from 500200 to 5000 just to try it out for the first time):

```

[net]
# Testing
#batch=1
#subdivisions=1
# Training
batch=64
subdivisions=16
width=416
height=416
channels=3
momentum=0.9
decay=0.0005
angle=0
saturation = 1.5
exposure = 1.5
hue=.1

learning_rate=0.001
burn_in=1000
#max_batches = 500200
max_batches = 5000
policy=steps
steps=400000,450000
scales=.1,.1

[convolutional]
batch_normalize=1
filters=32
size=3
stride=1
pad=1
activation=leaky

# Downsample

[convolutional]
batch_normalize=1
filters=64
size=3
stride=2
pad=1
activation=leaky
...
[yolo]
mask = 0,1,2
anchors = 10,13, 16,30, 33,23, 30,61, 62,45,
59,119, 116,90, 156,198, 373,326
classes=80
num=9
jitter=.3
ignore_thresh = .5
truth_thresh = 1
random=1

```

Now you need to download the pre-trained convolutional weights that have been pre-trained on Imagenet using the darknet53 model. You can just download the weights for the convolutional layers by executing the following command:

```
$wget https://pjreddie.com/media/files/darknet53.conv.74
```

Then I trained the model with COCO dataset by executing the following command:

```
$/darknet detector train cfg/coco.data cfg/yolov3.cfg darknet53.conv.74
```

The output you get should be similar to what I got as shown below:

```

./darknet detector train cfg/coco.data cfg/yolov3.cfg darknet53.conv.74
yolov3
layer  filters  size      input           output

```

```

0 conv  32 3 x 3 / 1  416 x 416 x 3  ->  416 x 416 x 32  0.299 BFLOPs
.....
105 conv  255 1 x 1 / 1  52 x 52 x 256  ->  52 x 52 x 255  0.353 BFLOPs
106 detection
Loading weights from darknet53.conv.74...Done!
Learning Rate: 0.001, Momentum: 0.9, Decay: 0.0005
Resizing
352
Loaded: 0.142245 seconds
Region 82 Avg IOU: 0.306305, Class: 0.459760, Obj: 0.475054, No Obj: 0.443880, .5R: 0.200000, .75R: 0.000000, count: 10
Region 94 Avg IOU: 0.191476, Class: 0.515562, Obj: 0.416233, No Obj: 0.460233, .5R: 0.000000, .75R: 0.000000, count: 7
.....
Region 82 Avg IOU: 0.303231, Class: 0.362318, Obj: 0.471073, No Obj: 0.442284, .5R: 0.142857, .75R: 0.000000, count: 7
.....
Region 106 Avg IOU: 0.260889, Class: 0.538027, Obj: 0.485092, No Obj: 0.472694, .5R: 0.133333, .75R: 0.000000, count: 15
.....
1: 728.814819, 728.814819 avg, 0.000000 rate, 1017.449199 seconds, 64 images
Loaded: 0.000042 seconds
Region 82 Avg IOU: 0.157086, Class: 0.629885, Obj: 0.393714, No Obj: 0.445066, .5R: 0.000000, .75R: 0.000000, count: 5
.....
Region 106 Avg IOU: 0.158112, Class: 0.301957, Obj: 0.455966, No Obj: 0.464749, .5R: 0.000000, .75R: 0.000000, count: 4
2: 716.763916, 727.609741 avg, 0.000000 rate, 1039.281667 seconds, 128 images
.....

```

So what does each output line shown above mean? I found that the above output is generated by line 239 from function `forward_yolo_layer` in the `yolo_layer.c` file as shown below:

```

printf("Region %d Avg IOU: %f, Class: %f, Obj: %f, No Obj: %f, .5R: %f, .75R:
%f, count: %d\n", net.index, avg_iou/count, avg_cat/class_count,
avg_obj/count, avg_anyobj/(1.w*1.h*1.n*1.batch), recall/count,
recall75/count, count);

```

So what is IOU? It stands for *Intersection over Union*, which measures how well the ground-truth bounding box overlaps with the predicted bounding box, as shown in Figure C.9. If IOU = 1, it means that the ground-truth bounding box and the predicted bounding box overlap exactly.

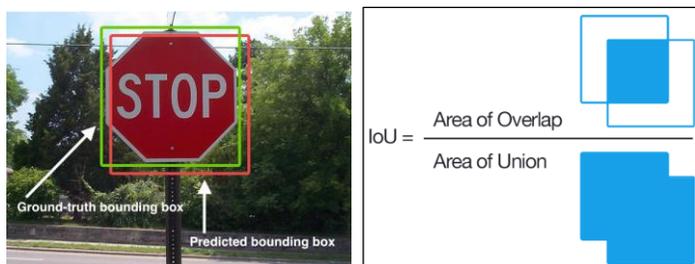


Figure C.9 IOU: Intersection over Union (Source: Courtesy of <https://www.pyimagesearch.com/>).

The portion of the code from function `forward_yolo_layer` is listed in Listing C.20 below, which shows how quantities on each of the output lines starting with “Region ...” are computed:

Listing C.20 How a Region is computed with YOLOv3

```

int mask_n = int_index(l.mask, best_n, l.n);
if(mask_n >= 0){
    int box_index = entry_index(l, b, mask_n*l.w*l.h + j*l.w + i, 0);
    float iou = delta_yolo_box(truth, l.output, l.biases, best_n, box_index, i, j, l.w, l.h, net.w, net.h, l.delta,
        (2-truth.w*truth.h), l.w*l.h);

    int obj_index = entry_index(l, b, mask_n*l.w*l.h + j*l.w + i, 4);
    avg_obj += l.output[obj_index];
    l.delta[obj_index] = 1 - l.output[obj_index];

    int class = net.truth[t*(4 + 1) + b*l.truths + 4];
    if (l.map) class = l.map[class];
    int class_index = entry_index(l, b, mask_n*l.w*l.h + j*l.w + i, 4 + 1);
    delta_yolo_class(l.output, l.delta, class_index, class, l.classes, l.w*l.h, &avg_cat);

    ++count;
    ++class_count;
    if(iou > .5) recall += 1;
    if(iou > .75) recall75 += 1;
    avg_iou += iou;
}
}

```

For your reference, I installed *darkent* on an Eclipse IDE for C/C++, which helps manage files and search better. For example, Figure C.10 shows how to search a string pattern from all C source files under the *Remote Search* tab.

 **Note:** Some tips about how to navigate through the YOLOv3 source code or any C/C++ programs on Eclipse for C/C++ projects:

- **Viewing the definition of a function:** Place your cursor to a function name called in the current function, and a pop-up window will show the definition of that function, which can be fully viewed by scrolling up and down.
 - **Jumping forward to where a function is defined:** Place your cursor to a function name called in the current function, right-click and select *Open Declaration*.
 - **Jumping backward to a caller:** Place your cursor to a function name called in the current function, right-click and select *Open Call Hierarchy*. Then in the *Call Hierarchy* window, right-click on the caller and select *Open*.
-

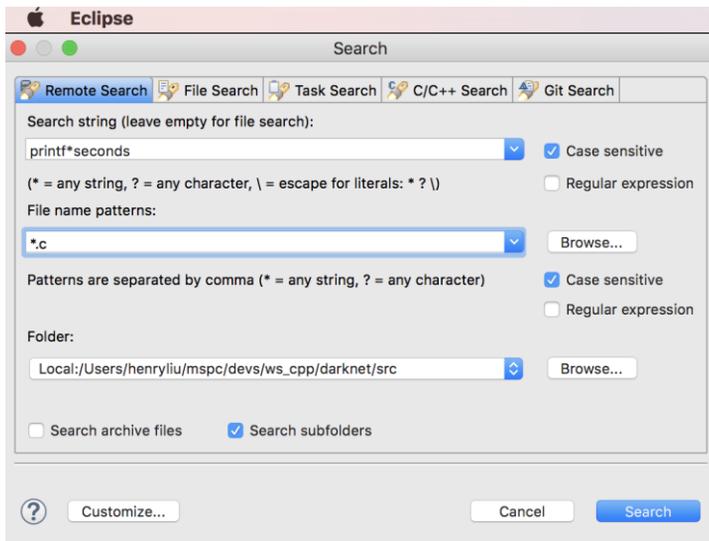


Figure C.10 Search capability from the Eclipse IDE for C/C++.

It's amazing that the entire darknet library is written in C with just 45 files with a total size of ~455 kB, as shown below:

12746	yolo_layer.c	3243	upsample_layer.c
14476	utils.c	3730	tree.c
3552	softmax_layer.c	8188	gemm.c
2940	shortcut_layer.c	1606	dropout_layer.c
3937	route_layer.c	10201	detection_layer.c
10093	rnn_layer.c	10568	demo.c
5037	reorg_layer.c	9787	deconvolutional_layer.c
19388	region_layer.c	44905	data.c
44504	parser.c	4095	cuda.c
3121	option_list.c	2759	crop_layer.c
5532	normalization_layer.c	9388	crnn_layer.c
30214	network.c	5174	cost_layer.c
3940	maxpool_layer.c	18620	convolutional_layer.c
4262	matrix.c	11056	connected_layer.c
24438	lstm_layer.c	10819	compare.c
2095	logistic_layer.c	1340	col2im.c
8929	local_layer.c	8435	box.c
1370	list.c	9397	blas.c
4471	layer.c	10366	batchnorm_layer.c
1794	l2norm_layer.c	1877	avgpool_layer.c
42105	image.c	3560	activations.c
1337	im2col.c	1707	activation_layer.c
13715	gru_layer.c	454817	total_size_in_byt

The *examples* directory contains the driver code for specific examples, including *darknet.c*, *detector.c*, *network.c*, etc., which call darknet library functions defined in the *src* directory. The code execution path with the COCO training example will be illustrated in the next section.

Once again, from the previous output lines, it shows that the first batch of 64 images took about $1017/60 = 17$ minutes or each image took about $1017/64 = 16$ seconds, and the second batch of 64 images took about $1039/60 = 17$ minutes as well or each image took about $1039/64 = 16$ seconds as well. I started training on April 21, and as of May 27, I got:

```
1056: 10.239875, 7.783431 avg, 0.001000 rate, 1969.657717 seconds, 67584 images
```

That is, the training reached batch # 1056 with an average loss of 7.783431. Compared with the loss of 728 at batch #1, the loss is about 100x smaller, but still a long way to reach the end of 500200 batches or iterations as specified in the original *yolov3.cfg* file! This motivated me to speed it up on my MacBook, as described next.

C.2.5 PROFILING YOLO

YOLO is written in C. You might want to know how YOLO is coded exactly, in which case a call graph will help. Or, you might want to analyze YOLO's performance as a software program, in which case, you need to profile YOLO while it is running. I had similar interests and figured out how we can do this easily. It turned out that using *Instruments* - the profiling feature of the XCode IDE on macOS - is the easiest way out of several options. In this section, I share my experience with you on how to obtain call graphs and CPU usage profiles with the Instruments tool on macOS.

To use Instruments, you need to have XCode and its Command Line Tools installed on your macOS, which you should already have if you did not skip §C.2.1. Then, just fire up *darknet* with a training task such as the one with COCO we demonstrated earlier. The next step is to find the process id (PID) of *darknet*, as shown from the Activity Monitor in Figure C.11, which was 13557 in my case. Finally, execute the following command with the *darknet*'s PID as shown below as in my case:

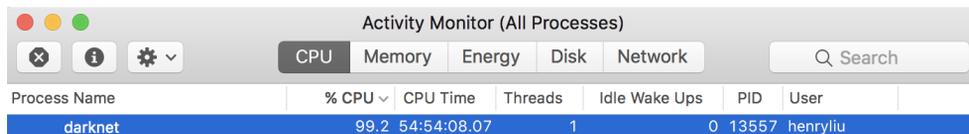
```
$ instruments -l 60000 -t Time\ Profiler -p 13557
```

The above command instructs the Instruments tool to instrument the darknet process for 60000 milliseconds or 60 seconds while it is running. If you get an error like

```
xcodeselect: error: tool 'instruments' requires Xcode, but active developer directory  
'/Library/Developer/CommandLineTools' is a command line tools instance
```

then, executing the following command should fix it as in my case:

```
$ sudo xcode-select -s /Applications/Xcode.app/Contents/Developer
```



The screenshot shows the Activity Monitor window titled "Activity Monitor (All Processes)". The "CPU" tab is selected. The table below shows the details for the "darknet" process.

Process Name	% CPU	CPU Time	Threads	Idle Wake Ups	PID	User
darknet	99.2	54:54:08.07	1	0	13557	henryliu

Figure C.11 The PID of the *darknet* process displayed on the Activity Monitor.

After the specified amount of time has passed, look for the Instruments icon on the Dock as shown in Figure C.12 below:

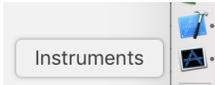


Figure C.12 The Instruments icon on the Dock.

Clicking on the above icon should bring up the Instruments panel as shown in Figure C.13. As you see, I got both the call graph and CPU stats in one shot. It is seen that *darknet*'s main function calls the `train_detector` function, which calls the `train_network` function, which in turn calls the `train_network_datum` function and the `get_next_batch` function.

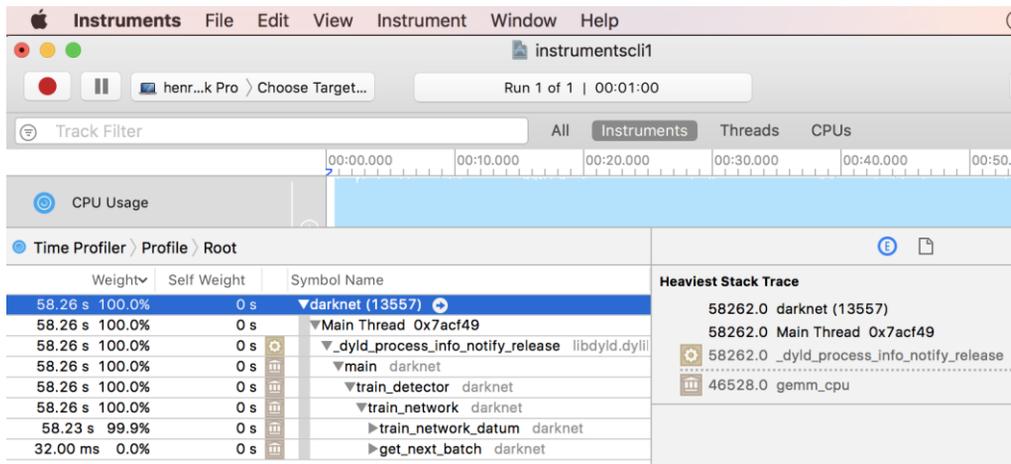


Figure C.13 The *darknet* CPU stats profiled with the Instruments tool

I drilled down further by expanding the `train_network_datum` function, as shown in Figure C.14. It is seen that the `train_network_datum` function called two more functions: `forward_network` and `backward_convolutional_layer`, which took ~87% and ~13% of the total CPU time, respectively. Figure C.15 shows the entire call graph in more details.

58.26 s	100.0%	0 s	darknet (13557)
58.26 s	100.0%	0 s	Main Thread Ox7acf49
58.26 s	100.0%	0 s	_dyld_process_info_notify_release libdyld.dylib
58.26 s	100.0%	0 s	main darknet
58.26 s	100.0%	0 s	train_detector darknet
58.26 s	100.0%	0 s	train_network darknet
58.23 s	99.9%	0 s	train_network_datum darknet
50.52 s	86.7%	0 s	forward_network darknet
7.71 s	13.2%	0 s	backward_convolutional_layer darknet
2.00 ms	0.0%	2.00 ms	axy_cpu darknet
32.00 ms	0.0%	0 s	get_next_batch darknet

Figure C.14 CPU time breakdown between the *darknet*'s two functions of `forward_network` and `backward_convolutional_layer` as revealed by the *Instruments* tool.

Weight	Self Weight	Symbol Name
58.26 s	100.0%	0 s
58.26 s	100.0%	0 s
58.26 s	100.0%	0 s
58.26 s	100.0%	0 s
58.26 s	100.0%	0 s
58.26 s	100.0%	0 s
58.26 s	100.0%	0 s
58.23 s	99.9%	0 s
50.52 s	86.7%	0 s
50.05 s	85.9%	2.00 ms
46.53 s	79.8%	46.53 s
1.74 s	2.9%	0 s
965.00 ms	1.6%	965.00 ms
282.00 ms	0.4%	282.00 ms
271.00 ms	0.4%	271.00 ms
150.00 ms	0.2%	150.00 ms
73.00 ms	0.1%	73.00 ms
863.00 ms	1.4%	862.00 ms
1.00 ms	0.0%	1.00 ms
778.00 ms	1.3%	121.00 ms
657.00 ms	1.1%	657.00 ms
78.00 ms	0.1%	78.00 ms
64.00 ms	0.1%	64.00 ms
157.00 ms	0.2%	31.00 ms
126.00 ms	0.2%	126.00 ms
121.00 ms	0.2%	121.00 ms
114.00 ms	0.1%	0 s
67.00 ms	0.1%	67.00 ms
47.00 ms	0.0%	47.00 ms
61.00 ms	0.1%	2.00 ms
49.00 ms	0.0%	4.00 ms
29.00 ms	0.0%	28.00 ms
1.00 ms	0.0%	1.00 ms
16.00 ms	0.0%	16.00 ms
6.00 ms	0.0%	6.00 ms
2.00 ms	0.0%	2.00 ms
1.00 ms	0.0%	1.00 ms
1.00 ms	0.0%	1.00 ms
8.00 ms	0.0%	0 s
7.00 ms	0.0%	7.00 ms
1.00 ms	0.0%	1.00 ms
5.00 ms	0.0%	0 s
5.00 ms	0.0%	5.00 ms
1.00 ms	0.0%	1.00 ms
7.71 s	13.2%	0 s
2.00 ms	0.0%	2.00 ms

Figure C.15 The *darknet* call graph as revealed with the Instruments tool.

As shown in Fig. C.15, the program execution begins with the main function in *darknet.c*. The “detector” argument initiates calling the function `train_detector` in *detector.c*. Then, the “train” argument initiates calling the function `train_network` in *network.c*, which calls the `train_network_datum` in *network.c* in turn. Refer to Listing C.21 for how this function is coded. As you see, this is where how `forward_network` and `backward_network` functions are called, how error is computed, and how the network is updated by calling the `update_network` function, which was not recorded during this profiling, as it was called only every subdivision or 16 images, which would take about ~50 minutes – much longer than the profiling duration of one minute or so. However, this is the part that involves the input model parameters of `momentum = 0.9` and `decay = 0.0005` as shown in the `yolov3.cfg` file given in §C.2.4, which are explained further next.

Listing C.21 Function `train_network_datum(network *net)` (in `network.c`)

```

289 float train_network_datum(network *net)
290 {
291     *net->seen += net->batch;
292     net->train = 1;
293     forward_network(net);
294     backward_network(net);
295     float error = *net->cost;
296     if ((*net->seen)/net->batch)%net->subdivisions == 0)
297         update_network(net);
297     return error;
298 }

```

To understand how the `forward_network` function calls the `forward_convolutional_layer` function, we show the `forward_network` function in Listing C.22. This function essentially loops through all layers defined for a given network with the `for`-loop defined from line 198 to 209. The line 204 initiates the call to the `forward_convolutional_layer` function, defined at line 221 with the function `make_convolutional_layer` in `convolutional_layer.c`, shown in Listing C.23. This function demonstrates how a convolutional layer is made. Listing C.24 shows how YOLOV3 implemented the key CNN functions of `forward_convolutional_layer`, `backward_convolutional_layer`, and `update_convolutional_layer`. These functions explain how these common CNN layers work, as is explained further next.

Listing C.22 Function `forward_network(network *net)` (in `network.c`)

```

188 void forward_network(network *netp)
189 {
190     #ifdef GPU
191         if (netp->gpu_index >= 0) {
192             forward_network_gpu(netp);
193             return;
194         }
195     #endif
196     network net = *netp;
197     int i;
198     for (i = 0; i < net.n; ++i) {
199         net.index = i;
200         layer l = net.layers[i];
201         if (l.delta) {
202             fill_cpu(l.outputs * l.batch, 0, l.delta, 1);
203         }
204         l.forward(l, net);
205         net.input = l.output;
206         if (l.truth) {
207             net.truth = l.output;
208         }
209     }
210     calc_network_cost(netp);
211 }

```

Listing C.23 Function `make_convolutional_layer` (in `src/convolutional_layer.c`)

```

176 convolutional_layer make_convolutional_layer(int batch, int h, int w, int
    c, int n, int groups, int size, int stride, int padding, ACTIVATION
    activation, int batch_normalize, int binary, int xnor, int adam)
177 {
178     int i;
179     convolutional_layer l = {0};
180     l.type = CONVOLUTIONAL;
181
182     l.groups = groups;
183     l.h = h;
184     l.w = w;
185     l.c = c;
186     l.n = n;
187     l.binary = binary;
188     l.xnor = xnor;
189     l.batch = batch;
190     l.stride = stride;
191     l.size = size;
192     l.pad = padding;
193     l.batch_normalize = batch_normalize;
194
195     l.weights = calloc(c/groups*n*size*size, sizeof(float));
196     l.weight_updates = calloc(c/groups*n*size*size, sizeof(float));
197
198     l.biases = calloc(n, sizeof(float));
199     l.bias_updates = calloc(n, sizeof(float));
200
201     l.nweights = c/groups*n*size*size;
202     l.nbiases = n;
203
204     // float scale = 1./sqrt(size*size*c);
205     float scale = sqrt(2./(size*size*c/l.groups));
206     //printf("convscale %f\n", scale);
207     //scale = .02;
208     //for(i = 0; i < c*n*size*size; ++i) l.weights[i] =
        scale*rand_uniform(-1, 1);
209     for(i = 0; i < l.nweights; ++i) l.weights[i] = scale*rand_normal();
210     int out_w = convolutional_out_width(l);
211     int out_h = convolutional_out_height(l);
212     l.out_h = out_h;
213     l.out_w = out_w;
214     l.out_c = n;
215     l.outputs = l.out_h * l.out_w * l.out_c;
216     l.inputs = l.w * l.h * l.c;
217
218     l.output = calloc(l.batch*l.outputs, sizeof(float));
219     l.delta = calloc(l.batch*l.outputs, sizeof(float));
220
221     l.forward = forward_convolutional_layer;
222     l.backward = backward_convolutional_layer;

```

```

223     l.update = update_convolutional_layer;
224     if(binary){
225         l.binary_weights = calloc(l.nweights, sizeof(float));
226         l.cweights = calloc(l.nweights, sizeof(char));
227         l.scales = calloc(n, sizeof(float));
228     }
229     if(xnor){
230         l.binary_weights = calloc(l.nweights, sizeof(float));
231         l.binary_input = calloc(l.inputs*l.batch, sizeof(float));
232     }
233
234     if(batch_normalize){
235         l.scales = calloc(n, sizeof(float));
236         l.scale_updates = calloc(n, sizeof(float));
237         for(i = 0; i < n; ++i){
238             l.scales[i] = 1;
239         }
240
241         l.mean = calloc(n, sizeof(float));
242         l.variance = calloc(n, sizeof(float));
243
244         l.mean_delta = calloc(n, sizeof(float));
245         l.variance_delta = calloc(n, sizeof(float));
246
247         l.rolling_mean = calloc(n, sizeof(float));
248         l.rolling_variance = calloc(n, sizeof(float));
249         l.x = calloc(l.batch*l.outputs, sizeof(float));
250         l.x_norm = calloc(l.batch*l.outputs, sizeof(float));
251     }
252     if(adam){
253         l.m = calloc(l.nweights, sizeof(float));
254         l.v = calloc(l.nweights, sizeof(float));
255         l.bias_m = calloc(n, sizeof(float));
256         l.scale_m = calloc(n, sizeof(float));
257         l.bias_v = calloc(n, sizeof(float));
258         l.scale_v = calloc(n, sizeof(float));
259     }
260
261     .....
322     l.workspace_size = get_workspace_size(l);
323     l.activation = activation;
324
325     fprintf(stderr, "conv %5d %2d x%2d /%2d %4d x%4d x%4d -> %4d x%4d
x%4d %5.3f BFLOPs\n", n, size, size, stride, w, h, c, l.out_w, l.out_h,
l.out_c, (2.0 * l.n * l.size*l.size*l.c/l.groups *
l.out_h*l.out_w)/1000000000.);
326
327     return l;
328 }

```

Listing C.24 Function `forward_convolutional_layer` (in `src/convolutional_layer.c`)

```

445 void forward_convolutional_layer(convolutional_layer l, network net)
446 {
447     int i, j;
448
449     fill_cpu(l.outputs*l.batch, 0, l.output, 1);
450
451     if(l.xnor){
452         binarize_weights(l.weights, l.n, l.c/l.groups*l.size*l.size,
453             l.binary_weights);
454         swap_binary(&l);
455         binarize_cpu(net.input, l.c*l.h*l.w*l.batch, l.binary_input);
456         net.input = l.binary_input;
457     }
458     int m = l.n/l.groups;
459     int k = l.size*l.size*l.c/l.groups;
460     int n = l.out_w*l.out_h;
461     for(i = 0; i < l.batch; ++i){
462         for(j = 0; j < l.groups; ++j){
463             float *a = l.weights + j*l.nweights/l.groups;
464             float *b = net.workspace;
465             float *c = l.output + (i*l.groups + j)*n*m;
466
467             im2col_cpu(net.input + (i*l.groups + j)*l.c/l.groups*l.h*l.w,
468                 l.c/l.groups, l.h, l.w, l.size, l.stride, l.pad, b);
469             gemm(0,0,m,n,k,1,a,k,b,n,1,c,n);
470         }
471     }
472
473     if(l.batch_normalize){
474         forward_batchnorm_layer(l, net);
475     } else {
476         add_bias(l.output, l.biases, l.batch, l.n, l.out_h*l.out_w);
477     }
478
479     activate_array(l.output, l.outputs*l.batch, l.activation);
480     if(l.binary || l.xnor) swap_binary(&l);
481 }
482
483 void backward_convolutional_layer(convolutional_layer l, network net)
484 {
485     int i, j;
486     int m = l.n/l.groups;
487     int n = l.size*l.size*l.c/l.groups;
488     int k = l.out_w*l.out_h;
489
490     gradient_array(l.output, l.outputs*l.batch, l.activation, l.delta);
491
492     if(l.batch_normalize){
493         backward_batchnorm_layer(l, net);
494     } else {
495         backward_bias(l.bias_updates, l.delta, l.batch, l.n, k);
496     }

```

```

497
498     for(i = 0; i < l.batch; ++i){
499         for(j = 0; j < l.groups; ++j){
500             float *a = l.delta + (i*l.groups + j)*m*k;
501             float *b = net.workspace;
502             float *c = l.weight_updates + j*l.nweights/l.groups;
503
504             float *im = net.input+(i*l.groups + j)*l.c/l.groups*l.h*l.w;
505
506             im2col_cpu(im, l.c/l.groups, l.h, l.w,
507                 l.size, l.stride, l.pad, b);
508             gemm(0, l, m, n, k, l, a, k, b, k, l, c, n);
509
510             if(net.delta){
511                 a = l.weights + j*l.nweights/l.groups;
512                 b = l.delta + (i*l.groups + j)*m*k;
513                 c = net.workspace;
514
515                 gemm(1, 0, n, k, m, l, a, n, b, k, 0, c, k);
516
517                 col2im_cpu(net.workspace, l.c/l.groups, l.h, l.w, l.size,
518                     l.stride,
519                     l.pad, net.delta + (i*l.groups +
520                         j)*l.c/l.groups*l.h*l.w);
521             }
522         }
523     }
524 void update_convolutional_layer(convolutional_layer l, update_args a)
525 {
526     float learning_rate = a.learning_rate*l.learning_rate_scale;
527     float momentum = a.momentum;
528     float decay = a.decay;
529     int batch = a.batch;
530
531     axpy_cpu(l.n, learning_rate/batch, l.bias_updates, 1, l.biases, 1);
532     scal_cpu(l.n, momentum, l.bias_updates, 1);
533
534     if(l.scales){
535         axpy_cpu(l.n, learning_rate/batch, l.scale_updates, 1, l.scales,
536             1);
537         scal_cpu(l.n, momentum, l.scale_updates, 1);
538     }
539
540     axpy_cpu(l.nweights, -decay*batch, l.weights, 1, l.weight_updates, 1);
541     axpy_cpu(l.nweights, learning_rate/batch, l.weight_updates, 1,
542         l.weights, 1);
543     scal_cpu(l.nweights, momentum, l.weight_updates, 1);
544 }

```

In general, without taking into account of the momentum and decay, the weights would be updated as follows:

$$w_{i+1} = w_i - \epsilon \left\langle \frac{\partial L}{\partial w} \Big|_{w_i} \right\rangle_{D_i} \quad (\text{C.1})$$

, where L is the loss, ϵ the learning rate, and D_i the current batch. The gradient term is computed based on the back propagation we described in Chapter 9. YOLOv3 implementation of this term can be found in the function of `backward_convolutional_layer` in `convolutional_layer.c`, as shown in Listing C.24. It would be an interesting exercise to further drill down how the `gradient_array` function is defined, etc.

The PDF file http://vision.stanford.edu/teaching/cs231b_spring1415/slides/alexnet_tugce_kyunghee.pdf shows Alex's original work, demonstrating how weights are updated in two steps as follows, when the momentum and decay terms are taken into account:

$$v_{i+1} = \textit{momentum} * v_i - \textit{decay} * \epsilon * w_i - \epsilon \left\langle \frac{\partial L}{\partial w} \Big|_{w_i} \right\rangle_{D_i} \quad (\text{C.2a})$$

$$w_{i+1} = w_i + v_{i+1} \quad (\text{C.2b})$$

, where v_i is the preceding correction. It's now clear that the gradient term in (C.2a) is computed in the function of `backward_convolutional_layer` in `convolutional_layer.c`, as shown in Listing C.24, while all other terms are computed and combined in the function of `update_convolutional_layer` in `convolutional_layer.c`. The `forward_convolutional_layer` function performs the matrix multiplications between the input and weight matrices to generate the outputs.

Figure C.15 shows that both the forward and backward convolutional layers call the `gemm` function, which took about 80% of the total CPU time. The update convolutional layer calls the `axpy` function in place of the `gemm` function, though. These two functions of `gemm` and `axpy` are explained further in the next section.

C.2.6 THE GEMM AND AXPY FUNCTIONS

So what are the `gemm` and `axpy` functions after all? Listing C.25 shows the CPU-version of the `gemm` function from lines 145-166, together with two of its variants of `gemm_nn` shown from line 74 to 89 and `gemm_tt` shown from line 126 to 142, respectively. Listing C.26 shows the CPU version of the `axpy` function, together with the `scale` and `fill` functions as well.

As is explained in wiki https://en.wikipedia.org/wiki/Basic_Linear_Algebra_Subprograms, the BLAS (Basic Linear Algebra Subprograms) spec specifies three levels of vector-matrix computations as follows:

- Level 1 (**axpy**): $\mathbf{y} \leftarrow \alpha \mathbf{x} + \mathbf{y}$ where \mathbf{x} and \mathbf{y} are vectors and α is a coefficient.
- Level 2 (**gemv** – **generalized matrix-vector multiplication**): $\mathbf{y} \leftarrow \alpha \mathbf{A} \mathbf{x} + \beta \mathbf{y}$ where matrix \mathbf{A} and coefficient β are added.
- Level 3 (**gemm** – **general matrix multiplication**): $\mathbf{C} \leftarrow \alpha \mathbf{A} \mathbf{B} + \beta \mathbf{C}$ where vectors \mathbf{x} and \mathbf{y} in level 2 (`gemv`) are replaced with matrices \mathbf{B} and \mathbf{C} , respectively.

As is seen, these functions are essentially multiplication functions involving constant coefficients, vectors and matrices. They are in general optimized and tuned on most particular platforms. Not surprisingly, they are at the core of machine learning in general and deep learning in particular.

To help you understand *gemm* a bit deeper, I added the next section with a standalone C program to illustrate exactly how general matrix multiplications are carried out with the *gemm* function implemented in YOLO.

Listing C.25 Function *gemm* (in *src/gemm.c*)

```

74 void gemm_nn(int M, int N, int K, float ALPHA,
75             float *A, int lda,
76             float *B, int ldb,
77             float *C, int ldc)
78 {
79     int i,j,k;
80     #pragma omp parallel for
81     for(i = 0; i < M; ++i){
82         for(k = 0; k < K; ++k){
83             register float A_PART = ALPHA*A[i*lda+k];
84             for(j = 0; j < N; ++j){
85                 C[i*ldc+j] += A_PART*B[k*ldb+j];
86             }
87         }
88     }
89 }

126 void gemm_tt(int M, int N, int K, float ALPHA,
127              float *A, int lda,
128              float *B, int ldb,
129              float *C, int ldc)
130 {
131     int i,j,k;
132     #pragma omp parallel for
133     for(i = 0; i < M; ++i){
134         for(j = 0; j < N; ++j){
135             register float sum = 0;
136             for(k = 0; k < K; ++k){
137                 sum += ALPHA*A[i+k*lda]*B[k+j*ldb];
138             }
139             C[i*ldc+j] += sum;
140         }
141     }
142 }
143
144
145 void gemm_cpu(int TA, int TB, int M, int N, int K, float ALPHA,
146              float *A, int lda,
147              float *B, int ldb,
148              float BETA,
149              float *C, int ldc)

```

```

150 {
151     //printf("cpu: %d %d %d %d %d %f %d %d %f %d\n",TA, TB, M, N, K,
        ALPHA, lda, ldb, BETA, ldc);
152     int i, j;
153     for(i = 0; i < M; ++i){
154         for(j = 0; j < N; ++j){
155             C[i*ldc + j] *= BETA;
156         }
157     }
158     if(!TA && !TB)
159         gemm_nn(M, N, K, ALPHA,A,lda, B, ldb,C,ldc);
160     else if(TA && !TB)
161         gemm_tn(M, N, K, ALPHA,A,lda, B, ldb,C,ldc);
162     else if(!TA && TB)
163         gemm_nt(M, N, K, ALPHA,A,lda, B, ldb,C,ldc);
164     else
165         gemm_tt(M, N, K, ALPHA,A,lda, B, ldb,C,ldc);
166 }

```

Listing C.26 CPU-version of the axpy, scale, and fill functions (in src/blas.c)

```

178 void axpy_cpu(int N, float ALPHA, float *X, int INCX, float *Y, int INCY)
179 {
180     int i;
181     for(i = 0; i < N; ++i) Y[i*INCY] += ALPHA*X[i*INCX];
182 }
183
184 void scal_cpu(int N, float ALPHA, float *X, int INCX)
185 {
186     int i;
187     for(i = 0; i < N; ++i) X[i*INCX] *= ALPHA;
188 }
189
190 void fill_cpu(int N, float ALPHA, float *X, int INCX)
191 {
192     int i;
193     for(i = 0; i < N; ++i) X[i*INCX] = ALPHA;
194 }

```

C.2.7 GENERAL MATRIX MULTIPLICATION (GEMM) EXAMPLES

From the computational point of view, machine learning is mostly about matrix multiplications, which is why the *gemm* function turned out to be responsible for ~80% of the CPU utilizations for a YOLOv3 training example, as illustrated in Figure C.15. Therefore, it is not too exaggerating to say that if you know how to make *gemm* run super-fast from either hardware or software perspective, you can have your own business.

Next, I'd like to share an example I worked out to illustrate how *gemm* is implemented in YOLOv3 and how it works in general. It is a standalone C project named *ml01* that I created on the Eclipse IDE for C/C++ Developers, as shown in Figure C.16. There are essentially four files: *ml01.c*, *gemm.h*, *gemm.c*

and makefile. I created the `ml10.c` and `makefile` myself, with `gemm.h` and `gemm.c` moved over from the original `darknet` project directory to make it “standalone.”

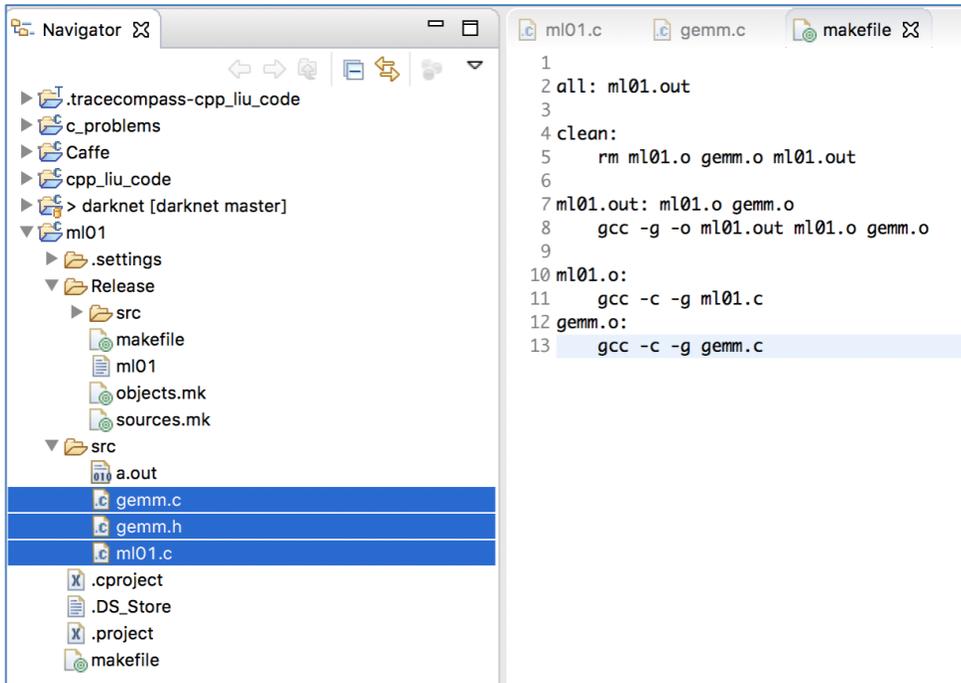


Figure C.16 The standalone C project illustrating how `gemm` works.

There are two ways to build this project. The easiest is to open up a Terminal and change to the `src` directory, and then issue the following command:

```
$gcc -o a.out *.c
```

An alternative is to change to the `Release` directory and issue the following command:

```
$make -k all
```

With the above command, “-k” means “*Keep going when some targets can't be made.*”

Note that the second approach is preferred if your project is large and more complicated. If you take this approach, you need to understand at least the following:

- The `makefile`, `objects.mk` and `sources.mk` in the `Release` directory are automatically generated from the `makefile` in the main directory. Therefore, make changes to the `makefile` in the main directory only if needed.
- You may want to spend a few minutes to get familiar with those `gcc` flags as shown below (or type `gcc -help` on the command prompt to look up yourself):
 - `-c`: Compile or assemble the source files but do not link.

- `-o file`: Write output as an executable to `file`.
- `-g`: Generate source level debug information.

Now, revisit Listing C.25 to see how `gemm_nn` and `gemm_cpu` are implemented. You can now correlate those functions with the `m101.c` program as shown in Listing C.27. You can start with the `main` function, which calls the `demo` function and the `gemm_test` function. The `demo` function illustrates a simple example with known result, whereas the `gemm_test` function allows us to experiment with larger matrices and more iterations with command line arguments. I suggest that you uncomment line 124 in Listing C.27, build and run this program first as follows, and make sure you get the following output from the `demo` function:

```
henryliu:src henryliu$ gcc -framework ACCELERATE -o a.out m101.c gemm.c
henryliu:Release henryliu$ ./m101
```

```
.....
array c in matrix format:
[ 1007.76, 1008.12
  1014.06, 1014.72 ]
.....
```

In the above command, the program is compiled with Apple's Accelerate framework, which has Apple's C-library for computationally intensive calculations. We will discuss more about this in the next section.

Next, let's dive a bit deeper into this simple program and see what we can learn from it about `gemm`, following the end of Listing C.27.

Listing C.27 `m101.c`

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <time.h>
4  #include <math.h>
5  #include "gemm.h"
6  #include <Accelerate/Accelerate.h>
7  #include <assert.h>
8
9  /*
10 [ 0.11 0.12 0.13 ] [ 11 12 ] [ 7.76 8.12 ]
11 [ 0.21 0.22 0.23 ] [ 21 22 ] = [ 4.06 4.72 ]
12 [ 31 32 ]
13 */
14 void demo () {
15     puts("Hello, gemm!!!");
16     int m = 2, k = 3, lda = 3;
17     float a[] = { 0.11, 0.12, 0.13,
18                 0.21, 0.22, 0.23 };
19
20     int n = 2, ldb = 2;
21     float b[] = { 11, 12,
22                 21, 22,
23                 31, 32 };
24     int ldc = 2;
```

```

25 float c[] = { 1000.00, 1000.00,
26              1000.00, 1000.00 };
27
28 puts("print array a ...");
29 for (int i = 0; i < m; i++)
30     for (int j = 0; j < k; j++)
31         printf ("%i, %i, %g\n", i, j, a[i*lda + j]);
32
33 puts("\nprint array b ...");
34 for (int i = 0; i < k; i++)
35     for (int j = 0; j < n; j++)
36         printf ("%i, %i, %g\n", i, j, b[i*ldb + j]);
37
38 // "0, 0" means no transpose
39 gemm(0, 0, 2, 2, 3, 1.0, a, 3, b, 2, 1, c, 2);
40
41 puts("\nprint array c ...");
42 for (int i = 0; i < m; i++)
43     for (int j = 0; j < n; j++)
44         printf ("%i, %i, %g\n", i, j, c[i*ldc + j]);
45 printf ("\narray c in matrix format:\n[ %g, %g\n", c[0], c[1]);
46 printf (" %g, %g ]\n", c[2], c[3]);
47 }
48
49 void init_array (float *c, int N) {
50     for (int i = 0; i < N; i++) {
51         c[i] = 0.0;
52     }
53 }
54
55 void gemm_test (int m, int k, int n, int lda, int ldb, int ldc, float
alpha, float beta, int iter) {
56     puts("\ngemm_test ...");
57     float *a = random_matrix (m, k);
58     float *b = random_matrix (k, n);
59     float *c = random_matrix (m, n);
60     float *c1 = malloc (m*n*sizeof(float));
61     memcpy(c1, c, m*n*sizeof(float));
62
63     for (int i = 0; i < m * n; i++)
64         assert (c[i] == c1[i]);
65     printf("c == c1 after memcpy\n");
66
67     clock_t start = clock(), end;
68     for (int i = 0; i < iter; i++) {
69         //init_array(c, m*n);
70         gemm(0, 0, m, n, k, alpha, a, m, b, k, beta, c, n);
71     }
72
73     double flop = ((double)m)*n*(2.*k + 2.)*iter;
74     double gflop = flop/pow(10., 9);
75     end = clock();

```

```

76     double seconds = ((double) (end - start)) / CLOCKS_PER_SEC;
77     printf("gemm: Matrix Multiplication %dx%d * %dx%d: %lf s, %lf
          GFLOPS\n",m,k,k,n, seconds, gflop/seconds);
78
79     puts("\nprint array c (partial) from gemm_test ...");
80     for (int i = 0; i < 2; i++)
81         for (int j = 0; j < 2; j++)
82             printf ("%i, %i, %lf\n", i, j, c[i*ldc + j]);
83
84
85     start = clock();
86     for (int i = 0; i < iter; i++) {
87         //init_array(c1, m*n);
88         cblas_sgemm(101, 111,111, m, n, k, alpha, a, lda, b, ldb, beta, c1,
          ldc);
89     }
90
91     end = clock();
92     seconds = ((double) (end - start)) / CLOCKS_PER_SEC;
93     printf("cblas_gemm: Matrix Multiplication %dx%d * %dx%d: %lf s, %lf
          GFLOPS\n",m,k,k,n, seconds, gflop/seconds);
94
95     puts("\nprint array c (partial) from gemm_test ...");
96     for (int i = 0; i < 2; i++)
97         for (int j = 0; j < 2; j++)
98             printf ("%i, %i, %lf\n", i, j, c1[i*ldc + j]);
99
100    /* c and c1 are not exactly the same
101    for (int i = 0; i < m * n; i++) {
102        printf ("%i, %lf, %lf\n", i, c[i], c1[i]);
103        assert (c[i] == c1[i]);
104    }
105    */
106    //printf("Passed\n");
107    free(a), free(b), free(c), free(c1);
108}
109int main(int argc, char *argv[]) {
110    int m = 2, k = 3, n = 2;
111    int lda = k, ldb = n, ldc = n;
112    int iter = 10;
113    float alpha = 1.0, beta = 1.0;
114    //demo();
115
116    for (int i = 0; i < argc; ++i)
117        printf("%i: %s ", i, argv[i]);
118
119    if (argc == 5) {
120        m = atoi(argv[1]);
121        k = atoi(argv[2]);
122        n = atoi(argv[3]);
123        iter = atoi(argv[4]);
124    } else {
125        m = 64, k = 64, n = 64, iter = 10;

```

```

126 }
127
128 lda = k, ldb = n, ldc = n;
129 gemm_test(m, k, n, lda, ldb, ldc, alpha, beta, iter);
130 return EXIT_SUCCESS;
131 }

```

First of all, we need to understand that matrices are 2D data structures, but represented in C as 1D arrays under the assumption of “row-major,” which means that the arrays start with row 1, then row 2, ..., and so on, sequentially. Secondly, the `gemm` function can specify whether matrix **A** or **B** or both are to be transposed. For example, the function `gemm_nn` means matrices **A** and **B** are taken with no transposition, while the function `gemm_tt` means both **A** and **B** are to be transposed. For simplicity, we assume that we deal with not-to-be-transposed matrices only, or the function `gemm_nn` only, or $TA = TB = 0$ at line 145 or line 158 in Listing C.25.

Then, it’s important to get the dimensions of the matrices right. There is a simple verification rule that a (m rows) (k columns) matrix **A** multiplied by a (k rows) (n columns) matrix **B** would yield a (m rows) (n columns) matrix **C**, or $(m \times k) \bullet (k \times n) \Rightarrow (m \times n)$. The other parameters of (lda , ldb , ldc) are simply the number of columns for matrices **a**, **b** and **c**, respectively.

Once you understand the above points, it’s easy to follow through the `main` function in Listing C.27. When you run the `ml01` program with no additional parameters, it will call `demo()` function first if not commented out and then call `gemm_test` with a (64×64) matrix for **a** and (64×64) matrix for **b** for 10 iterations, if no additional command line arguments are given. You can compile the program `ml01` and run the program by adding additional parameters of (m , k , n , $iter$) with the intended matrix dimensions and iteration as follows:

```
henryliu:src henryliu$ gcc -framework ACCELERATE -o a.out ml01.c gemm.c
```

```
henryliu:Release henryliu$ ./a.out 64 64 64 100000
```

```
0: ./a.out 1: 64 2: 64 3: 64 4: 10000
```

```
gemm_test ...
```

```
c == c1 after memcpy
```

```
gemm: Matrix Multiplication 64x64 * 64x64: 7.443356 s, 0.715376 GFLOPS
```

```
print array c (partial) from gemm_test ...
```

```
0, 0, 156576.281250
```

```
0, 1, 142320.734375
```

```
1, 0, 163024.203125
```

```
1, 1, 155483.562500
```

```
cblas_gemm: Matrix Multiplication 64x64 * 64x64: 0.105836 s, 50.311803 GFLOPS
```

```
print array c (partial) from gemm_test ...
```

```
0, 0, 156578.265625
```

```
0, 1, 142277.062500
```

```
1, 0, 163063.640625
```

```
1, 1, 155520.109375
```

As you see from the `gemm_test` function in Listing C.27, with given arrays of **a**, **b** and **c**, the program calls the `gemm` function implemented in YOLOv3 first and then the `cblas_sgemm` function from Apple’s

Accelerate framework. For example, the above run took ~7.4 seconds and achieved ~0.7 GFLOPS (Giga floating point operations per second) or 0.7 billion flops with the `gemm` function versus ~0.1 seconds and achieved ~50 GFLOPS with the `cblas_sgemm` function, i.e., Apple’s `cblas_gemm` function is 74x faster than the `gemm` function implemented in YOLOv3 in C. If you wonder how the number of floating point operations is estimated, it is given at line 73 in Listing C.27 as follows:

```
double flop = ((double)m)*n*(2.*k + 2.)*iter;
```

This is how it is estimated in the original `gemm.c` function implemented in YOLOv3.

If you run the same program as above, you may get a different GFLOPS number on your machine, but the partial result listed at the end should remain the same, since the same `random_matrix` function implemented in the same `gemm` function is used every time. I also made a `random_matrix2` function in `gemm.c` for this project with all randomly generated floating point data centered around “0” but the results were similar except that some elements in the array `c` had negative numbers. In addition, I’d like to mention that memory allocated to each array needs to be de-allocated properly as well.

Since this is a standalone program, you can experiment as much as you can. For example, I removed the “register” modifier in the following statement in `gemm.c`:

```
//register float A_PART = ALPHA*A[i*lda+k];
float A_PART = ALPHA*A[i*lda+k];
```

and it made very little difference in terms of performance.

C.2.8 ACCELERATING YOLOV3’S GEMM ON MAC OS WITH APPLE’S ACCELERATE FRAMEWORK

The test results illustrated in the preceding section with 64×64 matrices show that Apple’s `cblas_sgemm` function from its *Accelerate* framework is about $50/0.7 = 71$ times faster than the `gemm` function implemented in C in YOLOv3. Since the machine I used to run the test is a latest MacBook Pro, I wanted to try out more to see how much faster the same `gemm` test runs could go if I used Apple’s *Accelerate* framework on MacOS.

To achieve the above objective, I added the following header in `m101.c`:

```
#include <Accelerate/Accelerate.h>
```

Then, I made the `gemm_test` function to call the `cblas_sgemm` function at line 88 as follows:

```
cblas_sgemm(101, 111, 111, m, n, k, alpha, a, lda, b, ldb, beta, c1,
ldc)
```

where the item “101” means `CblasRowMajor` and the term “111” means `CblasNoTrans`, respectively. The document at https://developer.apple.com/documentation/accelerate/blas/cblas_transpose shows more options for the first three parameters of the `cblas_sgemm` function as follows:

```
enum CBLAS_TRANSPOSE {
    CblasNoTrans=111,
    CblasTrans=112,
    CblasConjTrans=113,
    AtlasConj=114
```

```
};
typedef enum CBLAS_TRANSPOSE CBLAS_TRANSPOSE
```

And here is the document https://developer.apple.com/documentation/accelerate/1513264-cblas_sgemm?language=objc showing the signature of the `cblas_sgemm` function:

```
void cblas_sgemm(const enum CBLAS_ORDER __Order, const enum CBLAS_TRANSPOSE __TransA, const enum
CBLAS_TRANSPOSE __TransB, const int __M, const int __N, const int __K, const float __alpha, const float *__A,
const int __lda, const float *__B, const int __ldb, const float __beta, float *__C, const int __ldc);
```

Once again, I re-compiled the `ml01.c` program with the following command:

```
$gcc -framework Accelerate ml01.c gemm.c
```

Then I ran more tests with 128×128 and 256×256 matrices, with all results listed below::

```
henryliu:src henryliu$ ./a.out 128 128 128 10000
```

```
0: ./a.out 1: 128 2: 128 3: 128 4: 10000
```

```
gemm_test ...
```

```
c == c1 after memcpy
```

```
gemm: Matrix Multiplication 128x128 * 128x128: 50.783187 s, 0.832376 GFLOPS
```

```
print array c (partial) from gemm_test ...
```

```
0, 0, 282547.093750
```

```
0, 1, 312655.281250
```

```
1, 0, 321351.843750
```

```
1, 1, 333862.968750
```

```
cblas_gemm: Matrix Multiplication 128x128 * 128x128: 1.039302 s, 40.672220 GFLOPS
```

```
print array c (partial) from gemm_test ...
```

```
0, 0, 282617.656250
```

```
0, 1, 312713.250000
```

```
1, 0, 321226.312500
```

```
1, 1, 333843.156250
```

```
henryliu:src henryliu$ ./a.out 256 256 256 10000
```

```
0: ./a.out 1: 256 2: 256 3: 256 4: 10000
```

```
gemm_test ...
```

```
c == c1 after memcpy
```

```
gemm: Matrix Multiplication 256x256 * 256x256: 390.243553 s, 0.863192 GFLOPS
```

```
print array c (partial) from gemm_test ...
```

```
0, 0, 657736.125000
```

```
0, 1, 603507.000000
```

```
1, 0, 607466.937500
```

```
1, 1, 564777.000000
```

```
cblas_gemm: Matrix Multiplication 256x256 * 256x256: 8.805666 s, 38.254351 GFLOPS
```

```
print array c (partial) from gemm_test ...
```

```
0, 0, 658317.812500
```

```
0, 1, 603902.687500
```

```
1, 0, 607938.625000
1, 1, 564726.500000
```

As is seen, `cblas_sgemm` is $6.27/0.08 = 78$, $50.78/1.04 = 49$, and $390/8.8 = 44$ times faster than `gemm` implemented in YOLOv3 for 64×64 , 128×128 and 256×256 matrices, respectively, which is very impressive.

C.2.9 COCO TRAINING ON MAC OS BY TAKING ADVANTAGE OF APPLE'S ACCELERATE FRAMEWORK

Given the huge benefit of using Apple's Accelerate framework, I was motivated to replace `gemm` with `cblas_sgemm` in YOLOv3. I made the following changes to `gemm.c` in YOLOv3's `darknet` directory:

- Added

```
#ifndef ACCEL
#include <Accelerate/Accelerate.h>
#endif
```

- The `gemm_cpu` function now looks like this:

```
void gemm_cpu(int TA, int TB, int M, int N, int K, float ALPHA,
             float *A, int lda,
             float *B, int ldb,
             float BETA,
             float *C, int ldc)
{
    //printf("cpu: %d %d %d %d %d %f %d %d %f %d\n",TA, TB, M, N, K, ALPHA, lda, ldb, BETA, ldc);
    // commented out for calling cblas_sgemm
```

```
#ifndef ACCEL
    if(!TA && !TB) {
        cblas_sgemm(101, 111,111, M, N, K, ALPHA, A, lda, B, ldb, BETA, C, ldc);
    } else if(TA && !TB) {
        cblas_sgemm(101, 112,111, M, N, K, ALPHA, A, lda, B, ldb, BETA, C, ldc);
    } else if(!TA && TB) {
        cblas_sgemm(101, 111,112, M, N, K, ALPHA, A, lda, B, ldb, BETA, C, ldc);
    } else {
        cblas_sgemm(101, 112,112, M, N, K, ALPHA, A, lda, B, ldb, BETA, C, ldc);
    }
#else
    int i, j;
    for(i = 0; i < M; ++i) {
        for(j = 0; j < N; ++j) {
            C[i*ldc + j] *= BETA;
        }
    }
    if(!TA && !TB) {
        gemm_nn(M, N, K, ALPHA,A,lda, B, ldb,C,ldc);
    } else if(TA && !TB) {
```

```

    gemm_tn(M, N, K, ALPHA,A,lda, B, ldb,C,ldc);
} else if (!TA && TB) {
    gemm_nt(M, N, K, ALPHA,A,lda, B, ldb,C,ldc);
} else {
    gemm_tt(M, N, K, ALPHA,A,lda, B, ldb,C,ldc);
}
#endif
}

```

Then I added the following at the beginning of the Makefile in the darknet directory:

```
ACCEL=1
```

and after the CFLAG line:

```

ifeq ($(ACCEL), 1)
COMMON+= -DACCEL
CFLAGS+= -DACCEL
CFLAGS+= -framework ACCELERATE
endif

```

Then, I rebuilt YOLOv3 with the commands of “make clean” and “make.” These changes have made YOLOv3 10 - 20 times faster than before. For example, these are the output lines before replacing `gemm` with `cblas_sgemm`:

```

2801: 6.420076, 6.379949 avg, 0.001000 rate, 3134.186031 seconds, 179264 images
.....

```

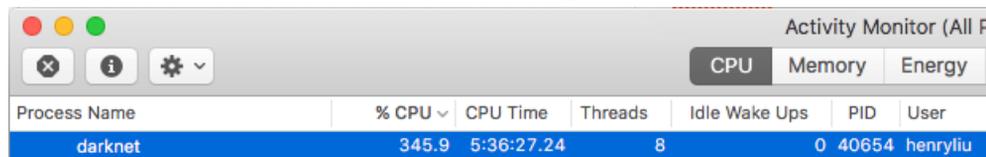
And these are the same iterations after replacing `gemm` with `cblas_sgemm`:

```

2801: 5.855477, 5.855477 avg, 0.001000 rate, 152.833456 seconds, 179264 images
.....

```

In addition, the YOLOv3 is now able to use all 8 CPU cores on my machine, as indicated by the number “8” under Threads shown in Figure C.17 (a) below. Figure C.17 (b) shows the new profile with YOLOv3 optimized with Apple’s Accelerate framework. Compare with Figure C.15 to see how this optimization has reduced the CPU time spent on the `gemm` functions significantly.



The screenshot shows the Activity Monitor window with the 'CPU' tab selected. The 'darknet' process is highlighted in blue. The table below shows the process details:

Process Name	% CPU	CPU Time	Threads	Idle Wake Ups	PID	User
darknet	345.9	5:36:27.24	8	0	40654	henryliu

Figure C.17 (a) YOLOv3 optimized with Apple’s Accelerate framework.

Time Profiler > Profile > Root > main > train_network

Weight	Self Weight	Symbol Name
50.86 s 100.0%	0 s	▼train_network darknet_mac_accel
50.85 s 99.9%	0 s	▼train_network_datum darknet_mac_accel
26.53 s 52.1%	0 s	▼forward_network darknet_mac_accel
339.00 ms 0.6%	16.00 ms	▶forward_yolo_layer darknet_mac_accel
49.00 ms 0.0%	0 s	▶forward_upsample_layer darknet_mac_accel
663.00 ms 1.3%	1.00 ms	▶forward_shortcut_layer darknet_mac_accel
36.00 ms 0.0%	0 s	▶forward_route_layer darknet_mac_accel
23.93 s 47.0%	13.00 ms	▼forward_convolutional_layer darknet_mac_accel
4.57 s 8.9%	4.55 s	▼im2col_cpu darknet_mac_accel
27.00 ms 0.0%	27.00 ms	▶_platform_bzero\$VARIANT\$Haswell libsystem_platform.dylib
5.96 s 11.7%	0 s	▶gemm_cpu darknet_mac_accel
9.09 s 17.8%	0 s	▼forward_batchnorm_layer darknet_mac_accel
5.06 s 9.9%	5.06 s	variance_cpu darknet_mac_accel
400.00 ms 0.7%	400.00 ms	scale_bias darknet_mac_accel
1.00 ms 0.0%	1.00 ms	scal_cpu darknet_mac_accel
782.00 ms 1.5%	782.00 ms	normalize_cpu darknet_mac_accel
1.48 s 2.9%	1.48 s	mean_cpu darknet_mac_accel
1.37 s 2.6%	1.37 s	copy_cpu darknet_mac_accel
400.00 ms 0.7%	400.00 ms	fill_cpu darknet_mac_accel
391.00 ms 0.7%	391.00 ms	add_bias darknet_mac_accel
3.51 s 6.8%	673.00 ms	▼activate_array darknet_mac_accel
2.83 s 5.5%	2.83 s	activate darknet_mac_accel
2.00 ms 0.0%	2.00 ms	DYLD-STUB\$\$_bzero darknet_mac_accel
2.00 ms 0.0%	0 s	▶<Unknown Address>
610.00 ms 1.1%	610.00 ms	fill_cpu darknet_mac_accel
5.00 ms 0.0%	5.00 ms	exp libsystem_m.dylib
894.00 ms 1.7%	204.00 ms	▼activate_array darknet_mac_accel
690.00 ms 1.3%	690.00 ms	activate darknet_mac_accel
1.00 ms 0.0%	1.00 ms	DYLD-STUB\$\$exp darknet_mac_accel
4.00 ms 0.0%	0 s	▶<Unknown Address>
33.00 ms 0.0%	0 s	▶backward_upsample_layer darknet_mac_accel
469.00 ms 0.9%	0 s	▶backward_shortcut_layer darknet_mac_accel
28.00 ms 0.0%	0 s	▶backward_route_layer darknet_mac_accel
23.77 s 46.7%	0 s	▼backward_convolutional_layer darknet_mac_accel
7.00 ms 0.0%	7.00 ms	sum_array darknet_mac_accel
3.66 s 7.1%	3.64 s	▼im2col_cpu darknet_mac_accel
16.00 ms 0.0%	16.00 ms	▶_platform_bzero\$VARIANT\$Haswell libsystem_platform.dylib
1.77 s 3.4%	1.77 s	gradient_array darknet_mac_accel
10.47 s 20.5%	0 s	▶gemm_cpu darknet_mac_accel
3.27 s 6.4%	3.27 s	col2im_cpu darknet_mac_accel
4.60 s 9.0%	4.07 s	▼backward_batchnorm_layer darknet_mac_accel
317.00 ms 0.6%	317.00 ms	scale_bias darknet_mac_accel
220.00 ms 0.4%	0 s	▼backward_bias darknet_mac_accel
220.00 ms 0.4%	220.00 ms	sum_array darknet_mac_accel
4.00 ms 0.0%	4.00 ms	▶0x7ffdfd47600 libsystem_m.dylib
12.00 ms 0.0%	12.00 ms	axpy_cpu darknet_mac_accel
1.00 ms 0.0%	1.00 ms	▶_platform_memmove\$VARIANT\$Haswell libsystem_platform.dylib
8.00 ms 0.0%	0 s	▶get_next_batch darknet_mac_accel

Figure C.17 (b) Profile with YOLOv3 optimized with Apple's Accelerate framework.

If you are interested in repeating what I did to speed up YOLOv3 on macOS, you can follow the steps I detailed above. However, depending on what you have on your machine and how much experience you have on macOS, it could be easy or difficult. Even I myself encountered some difficulties when I was trying to port YOLOv3 from my newer MacBook Pro to an older MacBook Pro. Here are some practices you might want to keep in mind:

- The *cmake* version might matter, so please use the latest *cmake*. I encountered many issues in building openCV on my older MacBook, which were resolved simply by upgrading *cmake* from 3.6.1 to 3.11.4. On my newer MacBook Pro, I have *cmake* 3.10.2, which also works.
- Install the latest llvm by executing the command of “*brew install llvm*.” The openMP is included in newer versions of llvm and does not require a separate flag like `-fopenmp` to enable omp.
- The python version also matters. I have python3.6 on my newer MacBook and 3.7.0 on my older MacBook. I had significant difficulties while building opencv on my older MacBook, as described next.

After successfully building openCV on my older MacBook, I got the following error when verifying the installation of openCV:

```
henrys-MBP-2:build henryliu$ python3
Python 3.6.5 (v3.6.5:f59c0932b4, Mar 28 2018, 05:52:31)
[GCC 4.2.1 Compatible Apple LLVM 6.0 (clang-600.0.57)] on darwin
Type "help", "copyright", "credits" or "license" for more information.
```

```
>>> import cv2
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
ModuleNotFoundError: No module named 'cv2'
```

Apparently, python 3.7 on my machine did not seem to be able to find where openCV library was. The blog at <https://www.codingforentrepreneurs.com/blog/install-opencv-3-for-python-on-mac/> helped me resolve the issue.

First I figured out the PATH for my python 3.7 as follows:

```
Python 3.7.0 (default, Jun 29 2018, 20:13:13)
[Clang 9.1.0 (clang-902.0.39.2)] on darwin
>>> import sys
>>> print(sys.path)
['', '/usr/local/Cellar/python/3.7.0/Frameworks/Python.framework/Versions/3.7/lib/python37.zip',
'/usr/local/Cellar/python/3.7.0/Frameworks/Python.framework/Versions/3.7/lib/python3.7',
'/usr/local/Cellar/python/3.7.0/Frameworks/Python.framework/Versions/3.7/lib/python3.7/lib-dynload',
'/usr/local/lib/python3.7/site-packages']
>>>
```

Then I checked the path of `/usr/local/lib/python3.7/site-packages` and found no openCV library there. The following command fixed the issue (all in one line) by creating a soft link to link the built `cv2.cpython-36m-darwin.so` to `/usr/local/lib/python3.7/site-packages` with an alias of `cv2.so`:

```
In -s /usr/local/Cellar/opencv/3.4.1_2/lib/python3.6/site-packages/cv2.cpython-36m-darwin.so
/usr/local/lib/python3.7/site-packages/cv2.so
```

I observed that the YOLOv3 optimized on my MacBook Pro takes anywhere between 80 – 200 seconds per batch, depending on if there are any other jobs running on my machine. I compared the original and optimized YOLOv3 on my MacBook Pro again and got the following results:

With ACCEL=1

```
henryliu:darknet henryliu$ ./darknet_detector train cfg/coco.data cfg/yolov3.cfg darknet53.conv.74
1: 862.024475, 862.024475 avg, 0.000000 rate, 109.445253 seconds, 64 images, 08-27-2018 15:26:12.000
```

With ACCEL=0

```
henryliu:darknet henryliu$ ./darknet_detector train cfg/coco.data cfg/yolov3.cfg darknet53.conv.74
1: 1447.881836, 1447.881836 avg, 0.000000 rate, 2373.783949 seconds, 64 images, 08-27-2018 16:08:49.000
```

This represents a $2373/105=22.6x$ speed-up. You can try it out on your machine and it will be very impressive. For your reference, refer to Figures C.18 - 20 for how COCO training performed on my MacBook Pro with Apple's Accelerate framework enabled as described above.

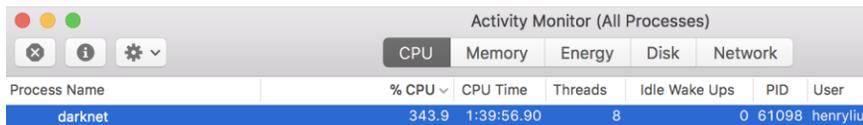


Figure C.17 YOLOv3 running on another MacBook Pro with all 8 CPU cores in use.

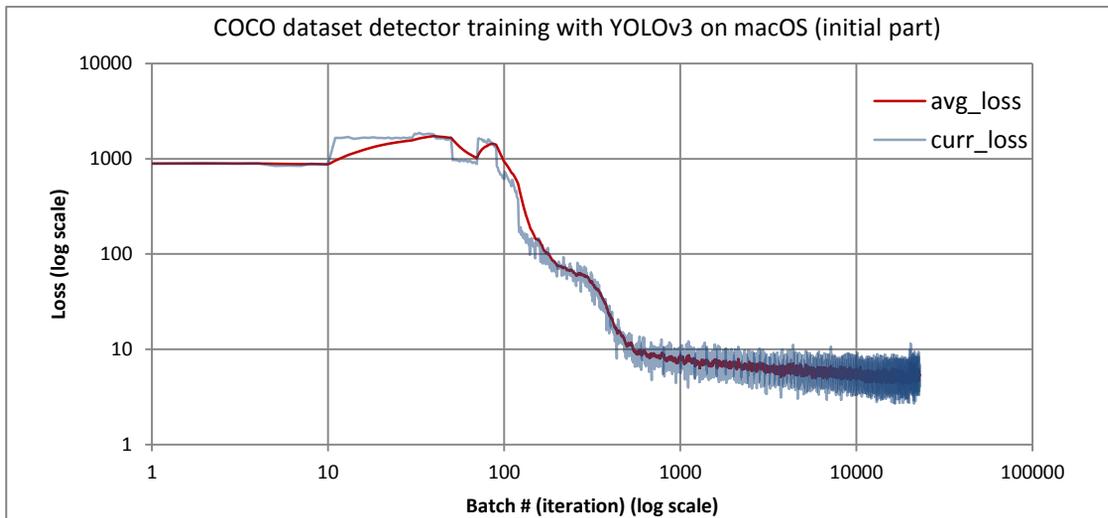


Figure C.18 COCO detector training with YOLOv3 optimized on macOS: current loss and average loss evolved with iterations from 1 up to 22756 with a batch size of 64 and subdivision of 16. The average loss started from over 1000 during the first 100 iterations and settled quickly down to around 8.0 at iteration 1000, which slowly progressed to 5.0 – 6.0 at the end of 22785 iterations.

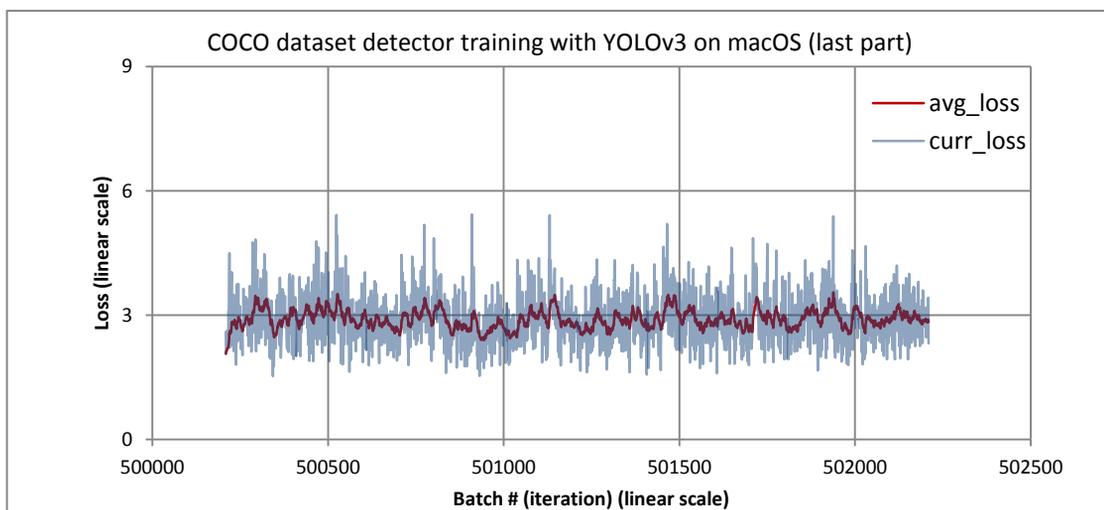


Figure C.19 COCO detector training with YOLOv3 optimized on macOS: current loss and average loss evolved with iterations from 500209 up to 502209 with a batch size of 64 and subdivision of 16. The losses settled down around 3.0.

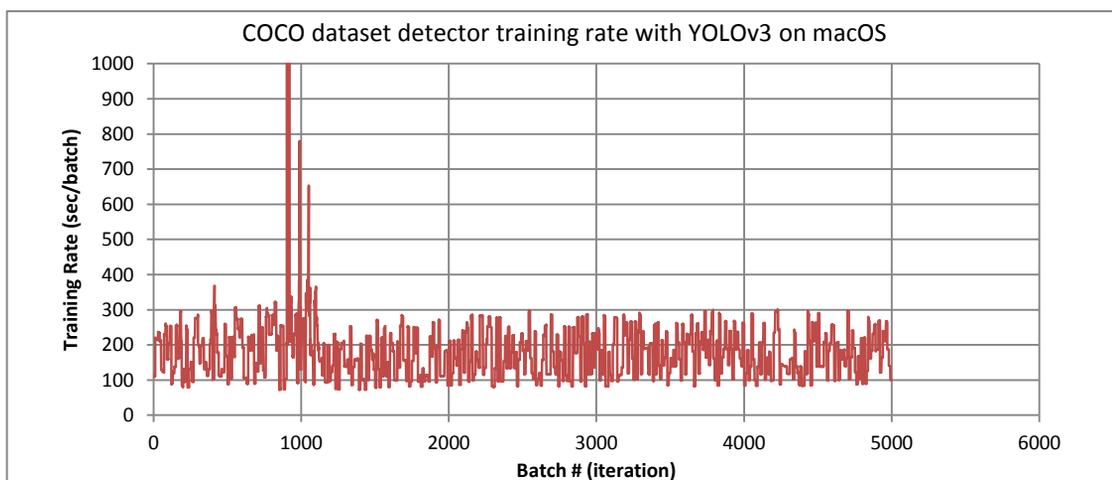


Figure C.20 COCO detector training with YOLOv3 optimized on macOS: training performance in terms of seconds per batch, with a batch size of 64 and subdivision of 16.

C.2.10 HOW YOLOV3 TRAINING IS KICKED OFF

It is interesting to explore how YOLOv3 training is kicked off exactly with the COCO dataset as an example. Most of the logistics is entailed in the `train_detector` function in `detector.c`. This function is shown in Listing C.28. It is a bit lengthy, with the main logic coded in the `while`-loop from lines 71-156.

The loop condition is that *the current batch is smaller than the max_batches* parameter specified in the *yolov3.cfg* file introduced earlier. If you have basic knowledge about C and CNN models as introduced in the main text, it's not hard to understand this piece of code, which is left as an exercise for those who are interested in such details.

Listing C.28 `train_detector` function (in `examples/detetcor.c`)

```

15 void train_detector(char *datacfg, char *cfgfile, char *weightfile, int
    *gpus, int ngpus, int clear)
16 {
17     list *options = read_data_cfg(datacfg);
18     char *train_images = option_find_str(options, "train",
        "data/train.list");
19     char *backup_directory = option_find_str(options, "backup",
        "/backup/");
20
21     srand(time(0));
22     char *base = basecfg(cfgfile);
23     printf("%s\n", base);
24     float avg_loss = -1;
25     network **nets = calloc(ngpus, sizeof(network));
26
27     srand(time(0));
28     int seed = rand();
29     int i;
30     for(i = 0; i < ngpus; ++i){
31         srand(seed);
32 #ifdef GPU
33         cuda_set_device(gpus[i]);
34 #endif
35         nets[i] = load_network(cfgfile, weightfile, clear);
36         nets[i]->learning_rate *= ngpus;
37     }
38     srand(time(0));
39     network *net = nets[0];
40
41     int imgs = net->batch * net->subdivisions * ngpus;
42     printf("Learning Rate: %g, Momentum: %g, Decay: %g\n",
        net->learning_rate, net->momentum, net->decay);
43     data train, buffer;
44
45     layer l = net->layers[net->n - 1];
46
47     int classes = l.classes;
48     float jitter = l.jitter;
49
50     list *plist = get_paths(train_images);
51     //int N = plist->size;
52     char **paths = (char **)list_to_array(plist);
53
54     load_args args = get_base_args(net);

```

```

55     args.coords = l.coords;
56     args.paths = paths;
57     args.n = imgs;
58     args.m = plist->size;
59     args.classes = classes;
60     args.jitter = jitter;
61     args.num_boxes = l.max_boxes;
62     args.d = &buffer;
63     args.type = DETECTION_DATA;
64     //args.type = INSTANCE_DATA;
65     args.threads = 64;
66
67     pthread_t load_thread = load_data(args);
68     double time;
69     int count = 0;
70     //while(i*imgs < N*120){
71     while(get_current_batch(net) < net->max_batches){
72         if(l.random && count++%10 == 0){
73             printf("Resizing\n");
74             int dim = (rand() % 10 + 10) * 32;
75             if (get_current_batch(net)+200 > net->max_batches) dim = 608;
76             //int dim = (rand() % 4 + 16) * 32;
77             printf("%d\n", dim);
78             args.w = dim;
79             args.h = dim;
80
81             pthread_join(load_thread, 0);
82             train = buffer;
83             free_data(train);
84             load_thread = load_data(args);
85
86             #pragma omp parallel for
87             for(i = 0; i < ngpus; ++i){
88                 resize_network(nets[i], dim, dim);
89             }
90             net = nets[0];
91         }
92         time=what_time_is_it_now();
93         pthread_join(load_thread, 0);
94         train = buffer;
95         load_thread = load_data(args);
96
97         /*
98         int k;
99         for(k = 0; k < l.max_boxes; ++k){
100             box b = float_to_box(train.y.vals[10] + 1 + k*5);
101             if(!b.x) break;
102             printf("loaded: %f %f %f %f\n", b.x, b.y, b.w, b.h);
103         }
104         */
105         /*
106         int zz;

```

```

107         for(zz = 0; zz < train.X.cols; ++zz){
108             image im = float_to_image(net->w, net->h, 3, train.X.vals[zz]);
109             int k;
110             for(k = 0; k < l.max_boxes; ++k){
111                 box b = float_to_box(train.y.vals[zz] + k*5, 1);
112                 printf("%f %f %f %f\n", b.x, b.y, b.w, b.h);
113                 draw_bbox(im, b, 1, 1,0,0);
114             }
115             show_image(im, "truth11");
116             cvWaitKey(0);
117             save_image(im, "truth11");
118         }
119     */
120
121     printf("Loaded: %lf seconds\n", what_time_is_it_now()-time);
122
123     time=what_time_is_it_now();
124     float loss = 0;
125 #ifdef GPU
126     if(ngpus == 1){
127         loss = train_network(net, train);
128     } else {
129         loss = train_networks(nets, ngpus, train, 4);
130     }
131 #else
132     loss = train_network(net, train);
133 #endif
134     if (avg_loss < 0) avg_loss = loss;
135     avg_loss = avg_loss*.9 + loss*.1;
136
137     i = get_current_batch(net);
138     printf("%ld: %f, %f avg, %f rate, %lf seconds, %d images\n",
139           get_current_batch(net), loss, avg_loss, get_current_rate(net),
140           what_time_is_it_now()-time, i*imgs);
139     if(i%100==0){
140 #ifdef GPU
141         if(ngpus != 1) sync_nets(nets, ngpus, 0);
142 #endif
143         char buff[256];
144         sprintf(buff, "%s/%s.backup", backup_directory, base);
145         save_weights(net, buff);
146     }
147     if(i%10000==0 || (i < 1000 && i%100 == 0)){
148 #ifdef GPU
149         if(ngpus != 1) sync_nets(nets, ngpus, 0);
150 #endif
151         char buff[256];
152         sprintf(buff, "%s/%s_%d.weights", backup_directory, base, i);
153         save_weights(net, buff);
154     }
155     free_data(train);
156 }
157 #ifdef GPU

```

```

158     if(ngpus != 1) sync_nets(nets, ngpus, 0);
159 #endif
160     char buff[256];
161     sprintf(buff, "%s/%s_final.weights", backup_directory, base);
162     save_weights(net, buff);
163 }

```

C.2.11 OBJECT DETECTION WITH UNDER-TRAINED YOLOV3

If you use the pre-trained weights named `yolov3.weights` demonstrated in §C2.3, you would find that the loss arrived at after 500200 batches should be in the range of 2.5 – 3. Assuming that we can achieve a 100 seconds per batch or 0.64 images/sec training speed on macOS with optimized YOLOv3, completing the entire training of 500200 iterations would take $500200 * 100 / 3600 / 24 = 579$ days or 1 year and 7 months! Obviously, we cannot run it for so long, so let's look at what kind of results we could get with a significantly under-trained YOLOv3. Some of the results are shown next.

C.2.11.1 BOUNDING BOX DETECTION WITH YOLOV3 AFTER TRAINED ~500K ITERATIONS ON MACOS

First of all, let's repeat the results with a well-trained YOLOv3, which reached a loss of ~2.5 after 500300 iterations, as shown by the last output line below:

```
500300: 2.530798, 2.527618 avg, 0.000010 rate, 260.651288 seconds, 32019200 images, 08-20-2018 16:16:45.000
```

Next, I tried the `dog.jpg` file and `horses.jpg` file, which are shown below in Figure C.21, respectively. Note that the detection now takes under one second with the optimized YOLOv3 on macOS vs ~8s with the original YOLOv3 compiled on macOS. Also note that the detections exceeded above 90% accuracy in both cases. Keep in mind that above 90% confidence is a very high confidence score.

```
henryliu:darknet henryliu$ ./darknet_mac_accel detect cfg/yolov3.cfg backup/yolov3_500300.weights
data/dog.jpg
```

```
.....
```

```
data/dog.jpg: Predicted in 0.909023 seconds.
```

```
bicycle: 99%
```

```
truck: 94%
```

```
dog: 99%
```

```
henryliu:darknet henryliu$ ./darknet_mac_accel detect cfg/yolov3.cfg backup/yolov3_500300.weights
data/horses.jpg
```

```
data/horses.jpg: Predicted in 0.902591 seconds.
```

```
horse: 98%
```

```
horse: 97%
```

```
horse: 92%
```

```
horse: 90%
```

C.2.11.2 BOUNDING BOX DETECTION WITH YOLOV3 TRAINED FOR 10000 ITERATIONS ON MACOS

Now, let's get back to a significantly under-trained model, which was trained with 10000 iterations only or 50x less trained than the previous case of being trained to the completion of 500k iterations. As shown in Figure C.22, with 10000 batches of training, the dog was detected as a cat, and one horse was detected as a sheep, although the bicycle, the automobile and two horses have been detected correctly.

```
henryliu:darknet henryliu$ ./darknet_mac_accel detect cfg/yolov3.cfg backup/yolov3_10000.weights data/dog.jpg
data/dog.jpg: Predicted in 0.879459 seconds.
```

cat: 71%

car: 56%

truck: 72%

bicycle: 65%

```
henryliu:darknet henryliu$ ./darknet_mac_accel detect cfg/yolov3.cfg backup/yolov3_10000.weights
data/horses.jpg
data/horses.jpg: Predicted in 0.873518 seconds.
```

sheep: 55%

horse: 72%

horse: 71%

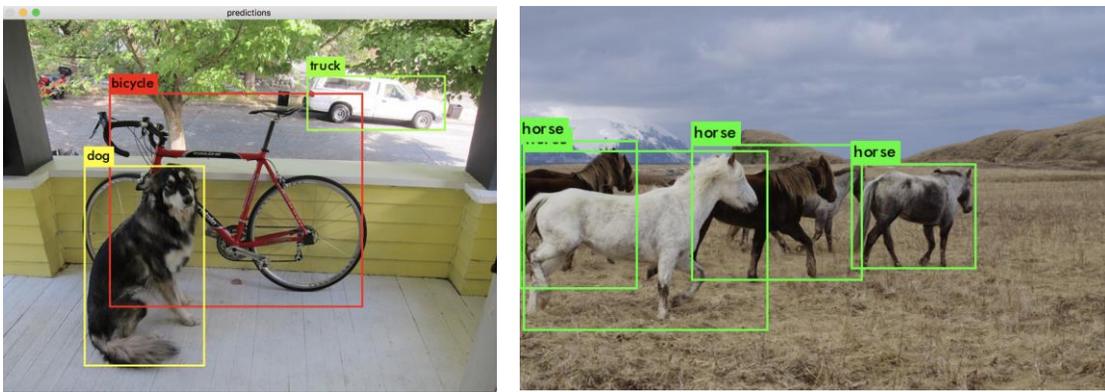


Figure C.21 Object detection test with YOLOv3 after fully trained with over 500k iterations, using the *dog.jpg* and *horses.jpg* image. All objects are detected correctly.

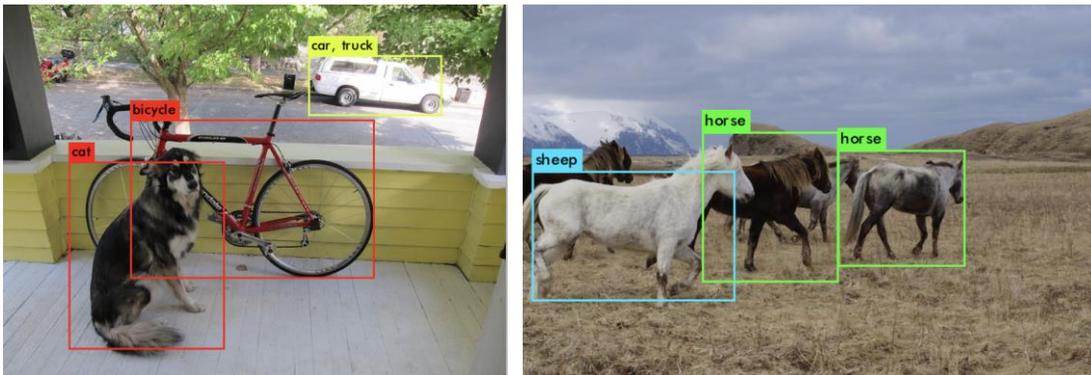


Figure C.22 Object detection test with YOLOv3 after trained for 10k iterations only, using: (a) the *dog.jpg* image, with the dog being mis-detected as a cat, and (b) using the *horses.jpg* image, with a horse mis-detected as a sheep.

C.2.11.3 BOUNDING BOX DETECTION WITH YOLOV3 TRAINED FOR 20000 ITERATIONS ON MACOS

What if we use a model trained with 20k iterations, or twice more iterations than the previous case? As shown in Figure C.23, with 20000 batches of training, the dog was not detected, and two horses were detected as cows.

```
henryliu:darknet henryliu$ ./darknet_mac_accel detect cfg/yolov3.cfg backup/yolov3_20000.weights data/dog.jpg
data/dog.jpg: Predicted in 0.848138 seconds.
bicycle: 76%
car: 58%
```

```
henryliu:darknet henryliu$ ./darknet_mac_accel detect cfg/yolov3.cfg backup/yolov3_20000.weights
data/horses.jpg
Loading weights from backup/yolov3_20000.weights...Done!
data/horses.jpg: Predicted in 0.922860 seconds.
cow: 69%
cow: 62%
horse: 55%
```

C.2.11.4 BOUNDING BOX DETECTION WITH YOLOV3 TRAINED FOR 1000 ITERATIONS ON MACOS

What if the COCO model were trained with only 1000 batches? Figure C.24 shows that no bounding boxes were detected with the *dog.jpg* image and a “person” was actually detected out of a horse! These examples show why a model needs to be thoroughly trained.

```
henryliu:darknet henryliu$ ./darknet_mac_accel detect cfg/yolov3.cfg backup/yolov3_1000.weights data/dog.jpg
Loading weights from backup/yolov3_1000.weights...Done!
data/dog.jpg: Predicted in 1.346573 seconds.
```

```
henryliu:darknet henryliu$ ./darknet_mac_accel detect cfg/yolov3.cfg backup/yolov3_900.weights data/horses.jpg
Loading weights from backup/yolov3_900.weights...Done!
data/horses.jpg: Predicted in 1.237775 seconds.
person: 65%
```

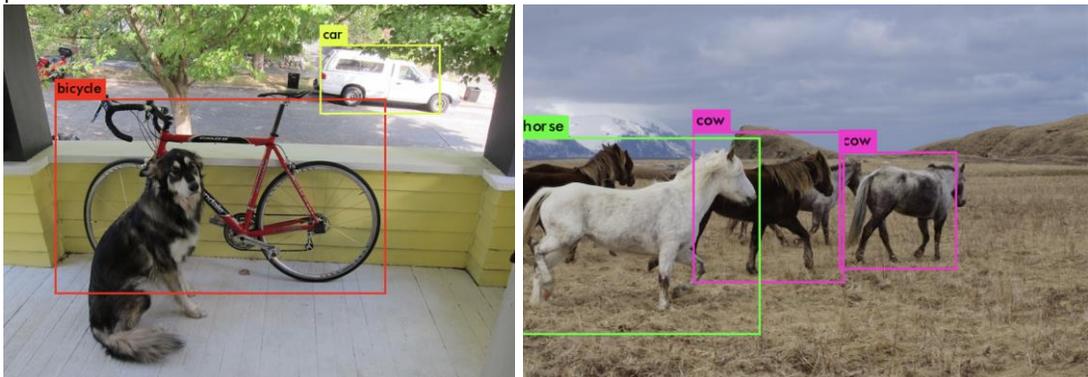


Figure C.23 Object detection test with YOLOv3 after trained for 20k iterations, using: (a) the *dog.jpg* image, with the dog not detected, and (b) the *horses.jpg* image, with two horses mis-detected as cows.



Figure C.24 Object detection test with YOLOv3 after trained for 1k iterations only, using: (a) the *dog.jpg* image, with no bounding boxes detected at all, and (b) the *horses.jpg* image, strangely enough that it actually mis-detected a “person” out of a horse.

C.3 IMAGE PROCESSING BASICS

In order to better understand how YOLOv3 works in particular and how deep learning CNN models work in general, it’s beneficial to have a basic understanding of how software processes images in general. In this section, I’ll show you an example in C of how OpenCV processes images, and then a more advanced image processing example in Java using the Marvin image processing framework. Such knowledge would enhance your understanding of YOLOv3 in particular and CNN models in general.

Let’s start with OpenCV first.

C.3.1 HOW OPENCV PROCESSES IMAGES IN C

First of all, this is how the website <https://opencv.org/> introduces OpenCV:

*OpenCV (Open Source Computer Vision Library) is released under a BSD license and hence it’s free for both academic and commercial use. It has **C++**, **Python** and **Java** interfaces and supports **Windows**, **Linux**, **Mac OS**, **iOS** and **Android**. OpenCV was designed for computational efficiency and with a strong focus on real-time applications. **Written in optimized C/C++**, the library can **take advantage of multi-core processing**. Enabled with OpenCL, it can take advantage of the **hardware acceleration** of the underlying heterogeneous compute platform.*

*Adopted all around the world, OpenCV has more than 47 thousand people of user community and estimated number of downloads exceeding **14 million**. Usage ranges from interactive art, to mines inspection, stitching maps on the web or through advanced robotics.*

I have emphasized some highlights in bold type in the above citation, which gives you a quick glance of OpenCV’s strengths and popularity.

Listing C.29 shows a C program that reads and displays the eagle image as shown in Figure C.25 that comes with the YOLOv3 distribution. If you have a basic understanding of how a C program works, it

should be easy for you to understand how this simple C program works. I suggest that you build and run this program by executing the commands given as follows, respectively.

```
gcc -o z.out `pkg-config --libs opencv` `pkg-config --cflags opencv` opencv_test.c -w
./z.out /Users/henryliu/mspc/devs/ws_cpp/ml01/images/eagle.jpg
```

Then, verify that you get the same results as shown in Figure C.25.

Listing C.29 opencv_test.c program

```
1  include <stdlib.h>
2  #include <stdio.h>
3  #include <math.h>
4  //These two lines cannot be recognized by Eclipse CDT
5  // #include <cv.h>
6  // #include <highgui.h>
7  //use absolute path
8  // #include "/usr/local/include/opencv/cv.h"
9  // #include "/usr/local/include/opencv/highgui.h"
10 //Use this include in place of the above two
11 #include "/usr/local/include/opencv2/core/core_c.h"
12 // resolve segmentation 11 runtime error
13 #include "/usr/local/include/opencv2/imgcodecs/imgcodecs_c.h"
14
15 // show an IplImage with a name for the window at x, y offsets relative
   to the upper left corner
16 void show_image (IplImage* img, char *window_name, int offset_x, int
   offset_y) {
17     cvNamedWindow(window_name, 1);
18     cvMoveWindow(window_name, offset_x, offset_y);
19     cvShowImage(window_name, img );
20 }
21
22 int main(int argc, char *argv[])
23 {
24     IplImage* img = 0;
25
26     char *image_file =
27         "/Users/henryliu/mspc/devs/ws_cpp/ml01/images/eagle.jpg";
28
29     // load the image with "-1" - as-is
30     Img = cvLoadImage(image_file, -1);
31     if(!img){
32         printf("Failed to load image file: %s\n",image_file);
33         exit(0);
34     }
35     int height, width, widthStep, channels;
36     uchar *data;
37
38     // get the image data
39     height = img->height;
```

```

40 width      = img->width;
41 widthStep  = img->widthStep;
42 channels    = img->nChannels;
43 data       = (uchar *)img->imageData;
44 printf("Image properties: (widthStep x width x height) = %dx%dx%d with
      %d channels\n", widthStep, width, height, channels);
45 printf("On macOS: locate images from a Terminal icon on the Dock
      ...\n");
46 printf("To exit: click ^c or any key after clicking the Terminal icon on
      the Dock \n");
47
48 show_image(img, "Original", 0, 0);
49 int i,j,k;
50
51 // invert the image
52 for(i = 0; i < height; i++)
53     for(j = 0; j < width; j++)
54         for(k = 0; k < channels; k++)
55             data[i * widthStep + j * channels + k] = 255 - data[i *
      widthStep + j * channels + k];
56
57 show_image(img, "Inverted", width / 2, height / 2);
58
59 // wait for a key or ^c and then release the image
60 cvWaitKey(0);
61 cvReleaseImage(&img );
62 return 0;
63 }

```

Now, focus on the two blocks shown from lines 35 – 43 and lines 52 – 55 in Listing C.29, respectively. The first block of code shows that an image has properties such as height, width, widthStep, nChannels and a data array. The height and width properties define the number of pixels, the widthStep property defines the number of data points for each row, the nChannels property defines the number of channels, while the data property define a 1D array that represents all pixels, including color information for each pixel. For example, running the above program would give the following output:

```
henryliu:src henryliu$ ./z.out /Users/henryliu/mspc/devs/ws_cpp/ml01/images/eagle.jpg
```

Image properties: (widthStep x width x height) = 2320x773x512 with 3 channels

On macOS: locate images from a Terminal icon on the Dock ...

To exit: click ^c or any key after clicking the Terminal icon on the Dock

which shows that the image of eagle.jpg has 773×512 pixels and 3 channels of (red, green, blue). The widthStep parameter has a value of 2320, since each row has 773 pixels, each of which has three values for the colors of red, green and blue, i.e., each pixel is defined with 3 values.

The second block from lines 52 - 55 shows how the data array is indexed: The variables i, j and k represent row, column and channel, respectively. This is exactly how an image is represented in YOLOv3, as YOLOv3 uses OpenCV as one of its image libraries.

However, image matrices cannot be used directly for convolutional computations. You may recall a function named `im2col_cpu`, as shown in Figures C. 15 and C. 17(b). This function converts image

matrices to column matrices to facilitate convolutional computations. Next, I'll show you an example to help you understand how this function works.



Figure C.25 An image read and displayed with OpenCV in C in its original and inverted forms.

C3.2. CONVERTING IMAGE MATRICES TO COLUMN MATRICES WITH THE `IM2COL` FUNCTION

First, visit the article available at <https://www.mathworks.com/help/images/ref/im2col.html> to become familiar with what the `im2col` function is about. That article has a very good example as shown in Figure C.26, illustrating how the function `im2col` works. Keeping sliding the matrix A with 2×2 sub-matrices and arranging them in column would result in the matrix B as shown there. For example, take the first upper left 2×2 sub-matrix of the matrix A and put it as the first column of matrix B, which changes the column-wise elements of $\{0, 0.2667, 0.0667, 0.3333\}$ in A into the first column of B. Then, slide down one row for the second sub matrix of $\{0.2667, 0.5333, 0.3333, 0.6000\}$ and place it in B as the second column, and so on. Since A is a 4×4 matrix, we would have a $4 \times (4-1) \times (4-1) = 4 \times 9$ column matrix B with a stride of 1.

A = 4x4								
0	0.0667	0.1333	0.2000					
0.2667	0.3333	0.4000	0.4667					
0.5333	0.6000	0.6667	0.7333					
0.8000	0.8667	0.9333	1.0000					
Rearrange the values into a column-wise arrangement.								
B = <code>im2col(A, [2 2])</code>								
B = 4x9								
0	0.2667	0.5333	0.0667	0.3333	0.6000	0.1333	0.4000	0.6667
0.2667	0.5333	0.8000	0.3333	0.6000	0.8667	0.4000	0.6667	0.9333
0.0667	0.3333	0.6000	0.1333	0.4000	0.6667	0.2000	0.4667	0.7333
0.3333	0.6000	0.8667	0.4000	0.6667	0.9333	0.4667	0.7333	1.0000

Figure C.26 A 4×4 matrix is converted to a 4×9 matrix with *mathworks'* `im2col`.

Now let's explore how this is done in YOLOv3. In fact, YOLOv3 has a function named `im2col_cpu` in `im2col.c` that does that kind of matrix conversion when an image is loaded and before being fed to a convolutional layer. Listing C.30 shows how this function is coded in C, together with a driver program I wrote for testing it, as explained next.

First, let's start with the driver function named `mathworks()`. Note that line 34 defines the parameters of `channels`, `height`, `width`, `ksize` and `stride`, where `ksize` is the convolution kernel size. Next, the `pad` parameter is computed with the formula of `pad = (ksize - 1) / 2`, which is derived from the following formula that if `stride = 1` and the `pad` parameter is chosen that way, the condition $n_{out} = n_{in}$ will always hold:

$$n_{out} = \frac{(n_{in} + 2 * pad - ksize)}{stride} + 1 \quad (C.3)$$

Apparently, if `stride > 1`, then the number of outputs will be reduced roughly proportionally.

Next, line 36 defines a 4x4 input matrix using the `mathworks` example mentioned previously. Line 37 computes the number of the elements for the column matrix to be obtained, and line 38 allocates the memory for the column matrix named `b_col`. Line 39 calls the `im2col_cpu` function, using those parameters given. The remaining lines are just for printing the column matrix and free its memory at the end. Now if you compile and run this program with the instructions given in the `im2col_test.c` program available from this book's download website, you should get the following output:

```
mathworks calling im2col_cpu ...
height_col = 3, width_col = 3
```

```
Input 4x4 matrix (manually added here for convenience)
```

```
0.0000, 0.0667, 0.1333, 0.2000,
0.2667, 0.3333, 0.4000, 0.4667,
0.5333, 0.6000, 0.6667, 0.7333,
0.8000, 0.8667, 0.9333, 1.0000
```

```
Column matrix (3x3x4=36)
```

```
0.000000, 0.333300, 0.666700, 0.333300, 0.666700, 0.400000, 0.800000, 0.400000, 0.733300,
0.066700, 0.400000, 0.066700, 0.400000, 0.733300, 0.533300, 0.866700, 0.466700, 0.866700,
0.133300, 0.533300, 0.133300, 0.466700, 0.266700, 0.600000, 0.933300, 0.600000, 0.933300,
0.266700, 0.600000, 0.200000, 0.600000, 0.333300, 0.666700, 0.333300, 0.666700, 1.000000
```

Note the differences between this column matrix and the one from `mathworks` discussed above. First, each sub-matrix is traversed in row-major fashion rather than in column-major fashion; secondly, in machine learning, when we specify a 2x2 kernel, the associated sub-matrix is actually a 3x3 matrix, as a 4 grid cell would need 3 connection points on each side. Other than those differences, this column matrix and the previous column matrix are essentially the same.

Now let's check the `im2col_cpu` function and see exactly how it works. First, lines 16 through 18 compute the dimensions for the height, width and total number of channels of the column matrix. Note that the total number of channels is equal to the number of image colors multiplied by the kernel size squared. Next, note that the 3 embedded `for`-loops follow the sequence of `channels` → `heights` → `width` top-to-bottom. Lines 20-22 calculate the offsets for the `channel`, `height` and `width` loop variables, based on the expanded channels and kernel size. Finally, line 39 maps the original image array to the

column matrix array, with the help of another function named `im2col_get_pixel` shown from lines 3-9. Keep in mind that with zero-padding, `pad` rows/columns are applied on each side of the four sides of an image, so `pad` is subtracted from the `row` and `col` variables as shown by lines 5 and 6. Line 7 shows exactly how zero-padding is applied.

This is how the `im2col` function is implemented in YOLOv3 and I hope it helps you gain insight into how images are turned into column matrices so that they will be more treatable numerically. There is also a function named `col2im`, which does the opposite, but we will not elaborate it here.

Next, I'll show you an example in Java to demonstrate how images can be manipulated interestingly.

Listing C.30 YOLOv3's `im2col` function

```

1  #include "im2col.h"
2  #include <stdio.h>
3  float im2col_get_pixel(float *im, int height, int width, int channels,
4      int row, int col, int channel, int pad)
5  {
6      row -= pad;
7      col -= pad;
8
9      if (row < 0 || col < 0 || row >= height || col >= width) return 0;
10     return im[col + width*(row + height*channel)];
11 }
12
13 //From Berkeley Vision's Caffe!
14 //https://github.com/BVLC/caffe/blob/master/LICENSE
15 void im2col_cpu(float* data_im, channels, int height, int width,
16     int ksize, int stride, int pad, float* data_col)
17 {
18     int c,h,w;
19     int height_col = (height + 2*pad - ksize) / stride + 1;
20     int width_col = (width + 2*pad - ksize) / stride + 1;
21     int channels_col = channels * ksize * ksize;
22     for (c = 0; c < channels_col; ++c) {
23         int w_offset = c % ksize;
24         int h_offset = (c / ksize) % ksize;
25         int c_im = c / ksize / ksize;
26         for (h = 0; h < height_col; ++h) {
27             for (w = 0; w < width_col; ++w) {
28                 int im_row = h_offset + h * stride;
29                 int im_col = w_offset + w * stride;
30                 int col_index = (c * height_col + h) * width_col + w;
31                 data_col[col_index] = im2col_get_pixel(data_im, height,
32                     width, channels, im_row, im_col, c_im, pad);
33             }
34         }
35     }
36 }
37
38 void mathworks () {
39     int channels = 1, height = 4, width = 4, ksize = 2, stride = 1;

```

```

35  int pad = (ksize - 1) / 2;
36  float a[16] = {0, 0.0667, 0.1333, 0.2000,
                 0.2667, 0.3333, 0.4000, 0.4667,
                 0.5333, 0.6000, 0.6667, 0.7333,
                 0.8000, 0.8667, 0.9333, 1.0000};
37  int num_elements = channels * ksize * ksize * height * width;
38  float *b_col = calloc(num_elements, sizeof(float));

39  im2col_cpu(a, channels, width, height, ksize, stride, pad, b_col);

40  int i, j, k;
41  printf("\nmathworks calling im2col_cpu ...\n");
42  int height_col = (height + 2*pad - ksize) / stride + 1;
43  int width_col = (width + 2*pad - ksize) / stride + 1;
44  printf("height_col = %d, width_col = %d\n", height_col, width_col);
45  int channels_col = channels * ksize * ksize;
46  int c, w, h;
47  for (c = 0; c < channels_col; ++c) {
48      int w_offset = c % ksize;
49      int h_offset = (c / ksize) % ksize;
50      int c_im = c / ksize / ksize;
51      for (h = 0; h < height_col; ++h) {
52          for (w = 0; w < width_col; ++w) {
53              int im_row = h_offset + h * stride;
54              int im_col = w_offset + w * stride;
55              int col_index = (c * height_col + h) * width_col + w;
56              printf("%f, ", b_col[col_index]);
57          }
58      }
59      free(b_col);
60 }

```

C.3.3 A MORE ADVANCED IMAGE PROCESSING EXAMPLE IN JAVA

Listing C.31 shows how the same eagle.jpg image can be manipulated to yield different effects. Figure C.27 shows the original image and three images processed with 3 Marvin plug-ins of `prewitt`, `errorDiffusion` and `emboss`, respectively. Since it uses the Marvin Java image processing library, the program needs to be compiled with `marvin_1.5.5.jar`. Once you built the program, you could run and get processed images as shown in Figure C.27.

Listing C.31 `JavaImageProcessing` program

```

1  import java.awt.GridLayout;
2  import javax.swing.JFrame;
3
4  import marvin.color.MarvinColorModelConverter;
5  import marvin.gui.MarvinImagePanel;
6  import marvin.image.MarvinImage;
7  import marvin.io.MarvinImageIO;
8  import marvin.plugin.MarvinImagePlugin;
9  import marvin.util.MarvinPluginLoader;

```

```

10
11 public class JavaImageProcessing extends JFrame{
12
13     // Marvin plug-ins for image processing
14     MarvinImagePlugin prewitt =
        MarvinPluginLoader.loadImagePlugin("org.marvinproject.image.edge.prewitt")
        ;
15     MarvinImagePlugin errorDiffusion =
        MarvinPluginLoader.loadImagePlugin("org.marvinproject.image.halftone.error
        Diffusion");
16     MarvinImagePlugin emboss =
        MarvinPluginLoader.loadImagePlugin("org.marvinproject.image.color.emboss")
        ;
17
18     public JavaImageProcessing(){
19         super("Java Image Processing Examples");
20
21         // Layout
22         setLayout(new GridLayout(2,2));
23
24         // Load image
25         MarvinImage img_input =
            MarvinImageIO.loadImage("./images/eagle.jpg");
26
27         int w = img_input.getWidth();
28         int h = img_input.getHeight();
29
30         MarvinImage img_prewitt = new MarvinImage(w, h);
31         MarvinImage img_errorDiffusion = new MarvinImage(w, h);
32         MarvinImage img_emboss = new MarvinImage(w, h);
33
34         //Processing plug-ins
35         errorDiffusion.process(img_input, img_prewitt);
36         prewitt.process(img_input, img_errorDiffusion);
37         emboss.process(img_input, img_emboss);
38
39         MarvinImageIO.saveImage(img_errorDiffusion,"./images/eagle4.jpeg");
40         // Set panels (top left, top right, bottom left, bottom right)
41         addPanel(img_input);
42         addPanel(img_prewitt);
43         addPanel(img_errorDiffusion);
44         addPanel(img_emboss);
45
46         traceShape("./images/eagle.jpg", "eagle");
47
48         setSize(1200,800);
49         setVisible(true);
50     }
51
52     public void addPanel(MarvinImage image){
53         MarvinImagePanel imagePanel = new MarvinImagePanel();
54         imagePanel.setImage(image);

```

```

54     add(imagePanel);
55 }
56
57 public static void main(String[] args) {
58     new JavaImageProcessing().setDefaultCloseOperation
59         (JFrame.EXIT_ON_CLOSE);
60 }
61 // http://jaypthakkar.blogspot.com/2014/
62 public static void traceShape(String imageFile, String imageName) {
63     // Load Plug-in
64     MarvinImagePlugin boundary =
65     MarvinPluginLoader.loadImagePlugin("org.marvinproject.image.morphological.
66     boundary");
67     // Load image
68     MarvinImage image = MarvinImageIO.loadImage(imageFile);
69     // Binarize 145 better than 145, 100 better than 145
70     MarvinImage binImage = MarvinColorModelConverter.rgbToBinary(image,
71     100);
72     MarvinImageIO.saveImage(binImage, "./images/" + imageName +
73     "_bin.png");
74     // morphological boundary
75     boundary.process(binImage.clone(), binImage);
76     MarvinImageIO.saveImage(binImage, "./images/" + imageName +
77     "_boudary.png");
78 }
79 }

```

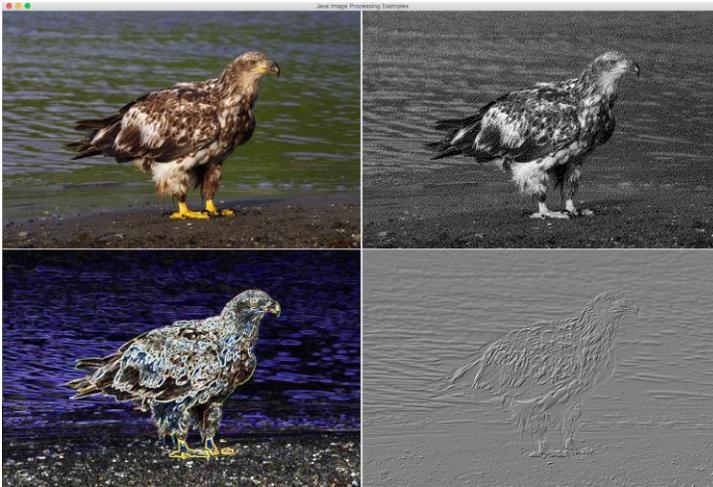


Figure C.27 An image read and processed with Marvin in Java in three different forms.

Note that the above Java program also calls a `traceShape` function, which yields binarized and morphological boundary images as shown in Figure C.28. This extra example is provided just to show you that images can be pre-processed in many different forms for machine learning training purposes.

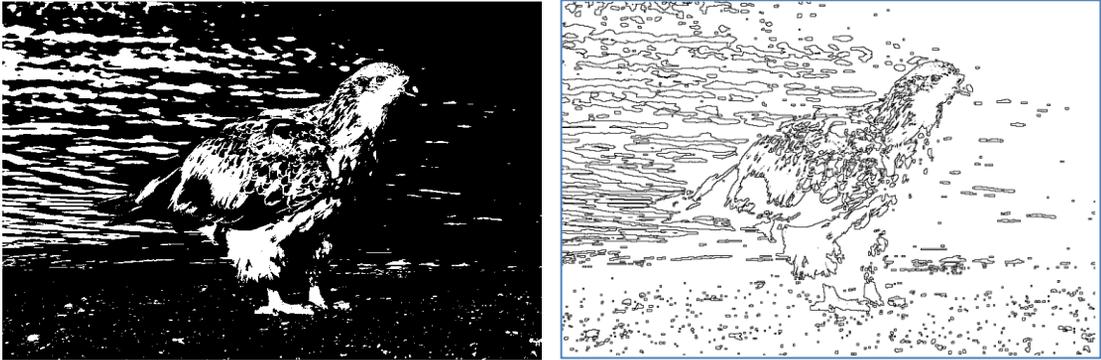


Figure C.28 The same eagle image in binary and morphological boundary forms.

C.3.4 RGB VERSUS HSV IMAGE FORMATS

There are two commonly used image formats, RGB and HSV, that can be interchanged with each other. The RGB format is simple, which is simply a mixing of the three different colors of red, green and blue. The HSV format is defined with three different attributes of *hue* or tint, *saturation* or amount of gray from completely gray to full color, and *value* or brightness from completely white to completely black, as shown in Figure C.29 known as the HSV color model. As is seen, hue is like the spectrum of color from red → yellow → green → cyan → blue → magenta, respectively. In fact, Adobe Photoshop calls HSV as HSB with *value* replaced with *brightness*. On macOS, you can open an image with Preview → Tools → Adjust Color..., as shown in Figure C.30, to check out the effects of *hue*, *saturation* and *value*, although we have *tint* for *hue*, *saturation*, but no *brightness*, which is more of a combo of *Exposure* (making all pixels lighter or darker), *Contrast* (making whites whiter or blacks blacker), *Highlights* (making white pixels whiter in lighter areas), *Shadows* (making darker pixels darker in darker areas), which are all located in the first section of the Preview Tools drop-down menu. YOLOv3 allows us to define hue, exposure, and saturation, with hue being an additive operation and other two multiplicative operations, defined in the `image.c` file.

C.4 THE ART AND SCIENCE OF DEEP LEARNING PERFORMANCE

Apparently, large configuration deep learning training requires GPUs, which means using NVidia GPUs most of the time. However, GPUs may not be easily available or affordable for individual AI/ML researchers. Still, knowing what performance one can expect with training large DL models on NVidia GPUs helps us understand how far we are - if we only have access to lower-end GPUs or CPUs on an optimized platform like macOS. For this purpose, I'd like to share some of the benchmarks conducted by NVidia from <https://www.nvidia.com/content/g/pdfs/DGX-Station-WP.pdf> so that we could get an assessment of the art and science of DL performance provided by an industrial leader like NVidia in the GPU industry.

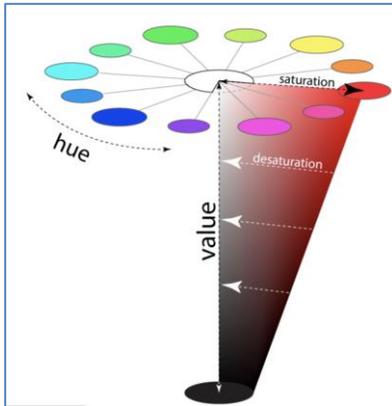


Figure C.29 The HSV color model.

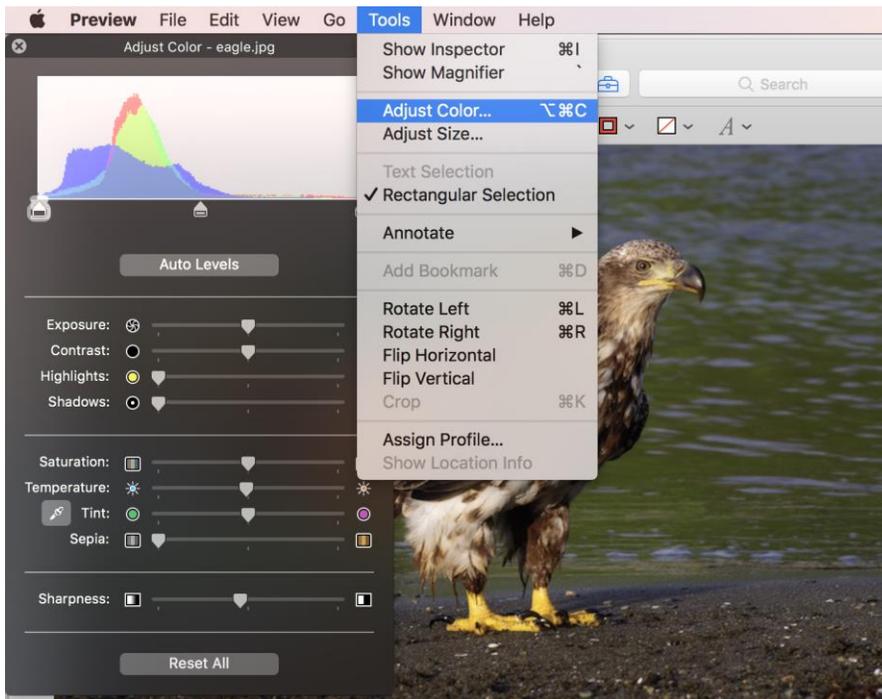


Figure C.30 Adjust color from Preview – Tools menu on macOS.

The above NVIDIA white paper is interesting to read in general. While I leave it up to you, I'd like to share two charts out of that white paper, one is 4x Tesla P100 vs 4x Tesla V100, as shown in Figure C.31, and the other is with 1-4 V100 GPUs for scalability, as shown in Figure C.32. As you see from Figure C.31, the ResNet-50 benchmark achieved ~900 images/second with the 4xP100 configuration and

~3000 images per second with the 4xV100 configuration. If the batch size is 64 images, it means that each batch or iteration took ~70 milliseconds only with the 4xP100 configuration or 7ms with the 4xV100 configuration, although such performance metric varies from model to model. On the other hand, Figure C.32 shows that a 4xP100 configuration slightly outperforms a 1xV100 configuration with ~900 images/second versus ~800 images/second. It is clear that newer GPUs may help speed up DL training significantly.

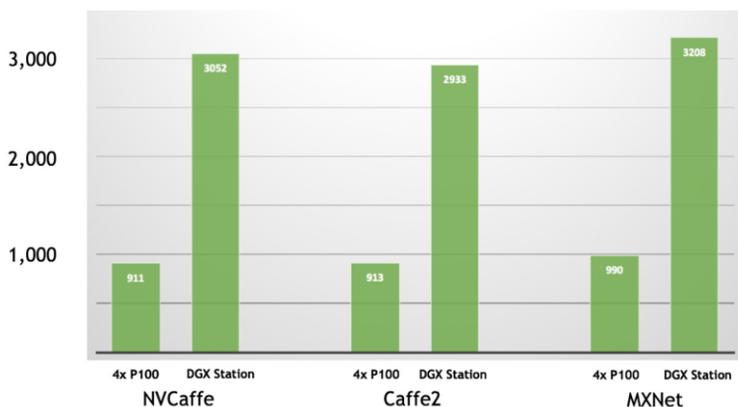


Figure C.31 NVidia 4x Tesla P100 (Pascal) versus DGX Station with 4x Tesla V100 (Volta). The y-axis represents the score in terms of images per second trained with the ResNet-50 training configuration with mixed precision FP16 and FP32.

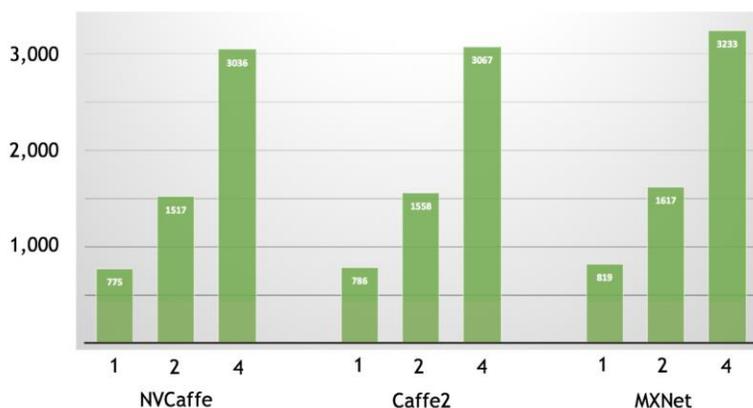


Figure C.32 NVidia 4x Tesla V100 (Pascal) scalability with 1, 2 and 4 GPUs. The y-axis represents the score in terms of images per second trained with the ResNet-50 training configuration with mixed precision FP16 and FP32.

C.5 TRAINING YOLOv3 ON GPUS

In this section, I share with you how I trained YOLOv3 with COCO on a 4xP100 NVidia GPU machine instance with Ubuntu 16.04 and NVidia cuda-9.1, after many trials to find out exactly what works and what doesn't. This may help save you a lot of time if you want to try the same. Then, I'll show you how to build YOLOv3 on Ubuntu 18.04 with CUDA-9.1 on a VM that has no GPU card.

C5.1 YOLOv3 COCO TRAINING ON GPUS

My situation was unique that I did not have a standalone GPU-empowered machine and the target machine was not a dev environment in the sense that I could only make a *darknet* binary and move it with COCO data to run it on the target machine with Ubuntu-16.04-cuda-9.1. It actually worked out well eventually and here is the first part of the output from a single GPU instance with the same COCO-YOLOv3 training runs I described in the previous sections:

```
1, 590.249084, 590.249084 avg, 0.000000 rate, 6.607559 seconds, 128 images
.....
200, 79.368683, 80.681160 avg, 0.000002 rate, 4.380903 seconds, 12800 images
.....
309, 51.029556, 53.454048 avg, 0.000009 rate, 7.997087 seconds, 19776 images
```

As you see, the same COCO training with YOLOv3 took 4-8 seconds per iteration on the single GPU machine instance I used, versus 80-200 seconds per iteration on my macOS with Apple's *Accelerate* framework enabled, which took up to 3200 seconds per iteration with YOLOv3's own `gemm` function. You can calculate how long it would take to go through the entire 520200 batches in each case.

I also ran the COCO training with 4 GPUs and 8 GPUs, respectively. This is how I kicked off a multi-GPU training, e.g., on a 4-GPU machine instance:

```
nohup ./darknet_gpu_omp detector train cfg/coco.data cfg/yolov3.cfg backup/yolov3.backup -gpus 0,1,2,3
1>coco_exp1_2d1_0927_0.txt 2>&1 &
```

Note from the above command on how it is prefixed with the `nohup` directive and how the 4 GPUs are specified. Besides, note the following:

- In the function `train_detector` in `detector.c`, there is a `for`-loop as shown in Listing C.32, which initializes the `nets` array for distributing training on multiple GPUs. Note at line 37 that the input learning rate is multiplied by the number of GPUs. This means that for an input learning rate of 0.001, each GPU would train with a learning rate of 0.004. To account for this, I specified the input learning rate as $0.001/ngpus$, where `ngpus` is the number of GPUs to train on the GPU instance. However, eventually, I figured out how to choose the input learning rate for the much more time-consuming, entire end-to-end training: *Try a few runs with increasing learning rates until you see it starts to diverge, as there is no formula for computing optimal learning rate.*
- I kept the same `max_batches` and `steps` settings as originally specified as follows, without dividing them by the number of GPUs:

```
learning_rate=0.00025
burn_in=1000
max_batches = 520200
policy=steps
```

```
steps=400000,450000
```

Listing C.32 nets array initialized with multi-GPUs in YOLOv3

```
31     for(i = 0; i < ngpus; ++i){
32         srand(seed);
33 #ifdef GPU
34         cuda_set_device(gpus[i]);
35 #endif
36         nets[i] = load_network(cfgfile, weightfile, clear);
37         nets[i]->learning_rate *= ngpus;
38     }
```

Finally, the output with multi-GPU training jumps by a number of $(ngpus - 1) \times 4$, as shown below that from batch 520199 to 520212, 12 batches or lines were jumped when my 4 GPU run was finishing up :

```
.....
520196: 3.128988, 2.981058 avg, 0.000010 rate, 8.054236 seconds, 133170176 images, 09-28-2018 02:00:34.000 , 6424 count, 1135.644165 epochs
520197: 2.959709, 2.978923 avg, 0.000010 rate, 7.526659 seconds, 133170432 images, 09-28-2018 02:00:41.000 , 6425 count, 1135.646240 epochs
520198: 3.189289, 2.999960 avg, 0.000010 rate, 7.594982 seconds, 133170688 images, 09-28-2018 02:00:49.000 , 6426 count, 1135.648438 epochs
520199: 3.212610, 3.021225 avg, 0.000010 rate, 7.330811 seconds, 133170944 images, 09-28-2018 02:00:56.000 , 6427 count, 1135.650635 epochs
Syncing ... Saving weights to backup/yolov3_final.weights
Done!
520212: 3.123556, 3.031458 avg, 0.000010 rate, 7.966230 seconds, 133174272 images, 09-28-2018 02:01:04.000 , 6428 count, 1135.679077 epochs
```

In the above output, the time stamp, count, and epochs terms were added by me in the source code for me to track the progress of training more easily. I added a `save_interval` net parameter so that a weights-save is performed whenever the condition `count%save_interval` equals 0 is met, where the count variable is always initialized to 0 every time when a new run is started. This 4 GPU run took about 7- 8 days total for 520200 batches or ~1136 epochs. Figure C.33 shows what were achieved with this run at the end, using the dog and horses images. It seems that the dog image had a “tvmonitor” object identified erroneously, most likely due to an over-trained model, resulting in the familiar issue of overfitting.

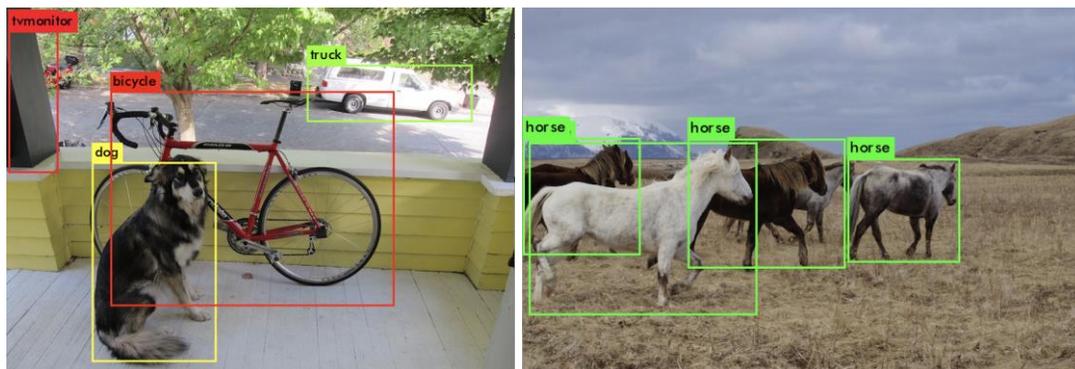


Figure C.33 Verification of the 4-GPU COCO training at the end of 520200 batches. The dog image had object:confidence scores of tvmonitor:67%, dog:98%, truck:83%, bicycle:99%, while the horses image had 99%, 97%, 95% and 85% for those four horses, respectively.

Next I share with you how I compiled YOLOv3 *darknet* binary to run on a Ubuntu 16.04-cuda-9.1 machine instance under the circumstance that I had no NVidia GPU empowered machine to compile YOLOv3 and the target machine does not support compiling YOLOV3 with cuda-9.1.

The procedure that worked out includes:

- Creating a Ubuntu-18.04 VM
- Installing cuda-9.1
- Making YOLOv3 on the Ubuntu-18.04-cuda-9.1 VM

Let's start with how to create a Ubuntu-18.04 VM first.

C.5.2 CREATING A UBUNTU-18.04 VM

I used VMware Fusion, which costs \$79.99 as of this writing, on my MacBook to create a Ubuntu-18.04 VM. I started with creating a Ubuntu 16.04 instead of 18.04, but had huge difficulties with the guest session login loop problem, and the issue was avoided only after I chose Ubuntu 18.04.

First download *ubuntu-18.04.1-desktop-amd64.iso* online. Use the desktop version instead of server version as we need the UI. Then on the VMware Fusion menu bar, select *File > New > Create a custom virtual machine > Continue*. Then select *Linux > Ubuntu 64-bit > Continue*. Select *Create a new virtual disk > Continue*. Choose *Customize Settings*. Enter a name like *Ubuntu-18.04.vmwarevm > Save*.

Now on the *Settings* panel, click *Processors & Memory*. Choose 2 or 4 processor cores and set memory to 8096 MB. Click *Show All* to go back to *Settings*. Click *Hard Disk* and change to 80 GB and *Save*. Now click on *CD/DVD Drive* and make sure you select the image you downloaded online as shown in Figure C.34 below. This tells Fusion where to find the Ubuntu image to create a VM from it. Now close the dialog and click the big triangle to start up your Ubuntu VM.

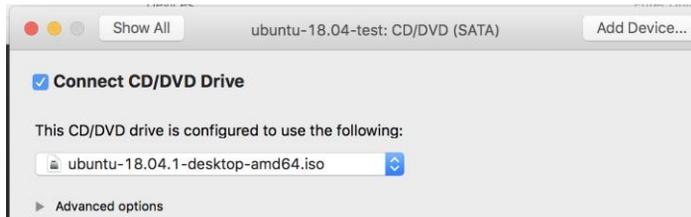


Figure C.34 Creating a Ubuntu-18.04 desktop VM.

Click *Install Ubuntu*. Select *English > Continue > Continue > Install Now > Continue*. Select your time zone *> Continue*. Enter your name, your computer's name, username and password and then *Continue*. The installation may take up to 10 minutes. Then, click *Restart Now*.

Now if it says "Please remove the installation medium, then reboot.", just Select *Virtual Machine > Restart* from the Fusion menu bar. You may need to press "*Ctrl + Cmd*" to get your mouse pointer back. Now click through a few initial dialogs and select *Upgrade Later*, and you should have your Ubuntu VM up and running now as shown in Figure C.35 below. Now open up a Terminal and execute the following commands:

```
$sudo apt-get install open-vm-tools-desktop
```

```
$sudo apt install net-tools
$sudo reboot
```

Now you can directly drag-drop files between your macOS host and Ubuntu VM.

Note: Tips with creating a Ubuntu VM on macOS using VMware Fusion: (a) If you lose your mouse pointer on the screen, press the “Do you want to enable Siri” button near the power button if you have a touch bar, or press “Ctrl + Cmd” key combo; and (b) After enabling Sharing between your Ubuntu VM and macOS host, you can change to `/mnt/hgfs/` to find your shared directories, but drag-and-drop is the easiest way to move files between the two systems.



Figure C.35 A Ubuntu-18.04 desktop VM created with VMware Fusion on macOS.

C.5.3 INSTALLING CUDA-9.1

First, download NVidia cuda-9.1 base and patch 3 from https://developer.nvidia.com/cuda-91-download-archive?target_os=Linux&target_arch=x86_64&target_distro=Ubuntu&target_version=1604&target_type=runfilelocal. Then install the base as follows:

```
$chmod +x cuda_9.1.85_387.26_linux.run
$./cuda_9.1.85_387.26_linux.run --override
```

```
Do you accept the previously read EULA?
accept/decline/quit: accept
```

```
You are attempting to install on an unsupported configuration. Do you wish to continue?
```

(y)es/(n)o [default is no]: y

Install NVIDIA Accelerated Graphics Driver for Linux-x86_64 387.26?

(y)es/(n)o/(q)uit: n

Install the CUDA 9.1 Toolkit?

(y)es/(n)o/(q)uit: y

Enter Toolkit Location

[default is /usr/local/cuda-9.1]:

/usr/local/cuda-9.1 is not writable.

Do you wish to run the installation with 'sudo'?

(y)es/(n)o: y

Please enter your password:

Do you want to install a symbolic link at /usr/local/cuda?

(y)es/(n)o/(q)uit: y

Install the CUDA 9.1 Samples?

(y)es/(n)o/(q)uit: n

Installing the CUDA Toolkit in /usr/local/cuda-9.1 ...

```
=====
= Summary =
=====
```

Driver: Not Selected

Toolkit: Installed in /usr/local/cuda-9.1

Samples: Not Selected

Then install Patch 3 as follows:

```
$chmod +x cuda_9.1.85.3_linux.run
```

```
$sudo ./cuda_9.1.85.3_linux.run
```

Do you accept the previously read EULA?

accept/decline/quit: accept

Enter CUDA Toolkit installation directory

[default is /usr/local/cuda-9.1]:

Installation complete!

Installation directory: /usr/local/cuda-9.1

Now add the following to the `~/.bashrc` file and then source it:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/lib:/usr/local/cuda-9.1/lib64
```

```
export PATH=$PATH:/usr/local/cuda-9.1/bin
```

Finally install gcc-5 as making YOLOv3 requires this particular gcc version, as follows:

```
$sudo apt-get install gcc-5 g++-5 g++-5-multilib gfortran-5
$sudo rm /usr/bin/gcc
$sudo ln -s gcc-5 /usr/bin/gcc
```

Now check the *nvcc* and *gcc* version and you should get the following outputs:

```
$ nvcc --version
nvcc: NVIDIA (R) Cuda compiler driver
Copyright (c) 2005-2017 NVIDIA Corporation
Built on Fri_Nov__3_21:07:56_CDT_2017
Cuda compilation tools, release 9.1, V9.1.85
$ gcc --version
gcc (Ubuntu 5.5.0-12ubuntu1) 5.5.0 20171010
Copyright (C) 2015 Free Software Foundation, Inc.
```

You are ready to make YOLOv3 on your Ubuntu-18.04 VM with NVidia cuda9.1 now.

C.5.4 MAKING YOLOV3 ON THE UBUNTU-18.04-CUDA-9.1 VM

In order to make YOLOv3 on your Ubuntu VM, first install/update the *make* utility as follows:

```
$sudo apt install make
$sudo apt install make-guile
```

Now copy the *darknet* download to your Ubuntu VM. You only need the following directories and files:

```
examples
include
Makefile
src
```

Make sure that you have the following settings in your *Makefile*:

```
GPU=1
CUDNN=0
OPENCV=0
ACCEL=0
OPENMP=0
DEBUG=0

ifeq ($(GPU), 1)
COMMON+= -DGPU -I/usr/local/cuda-9.1/include/
CFLAGS+= -DGPU
LDFLAGS+= -L/usr/local/cuda-9.1/lib64 -lcuda -lculdrt -lcublas -lcurand
```

```
endif
```

Now after executing “make clean” and “make” commands and if you get the following error:

```
/usr/bin/ld: cannot find -lcuda
collect2: error: ld returned 1 exit status
Makefile:90: recipe for target 'libdarknet.so' failed
make: *** [libdarknet.so] Error 1
```

Then, remove “-lcuda” flag for linking the *cuda* library from the LDFLAGS setting in your *Makefile* and re-make. The generated *darknet* executable file is your YOLOv3 binary built on Ubuntu 18.04 with *cuda-9.1*. You can test it with the CIFAR small configuration on a remote GPU-empowered machine instance and it should work. Here is the output I tried:

```
/cifar# ./darknet_ubtn_gpu classifier train cfg/cifar.data cfg/cifar_small.cfg
cifar_small
1
layer  filters  size      input       output
 0 conv   32 3x3/1  28x28x3 -> 28x28x32 0.001 BFLOPs
 1 max     2x2/2  28x28x32 -> 14x14x32
 2 conv   64 3x3/1  14x14x32 -> 14x14x64 0.007 BFLOPs
 3 max     2x2/2  14x14x64 -> 7x7x64
 4 conv  128 3x3/1   7x7x64 -> 7x7x128 0.007 BFLOPs
 5 conv   10 1x1/1   7x7x128 -> 7x7x10 0.000 BFLOPs
 6 avg              7x7x10 -> 10
 7 softmax                10
 8 cost                   10
Learning Rate: 0.1, Momentum: 0.9, Decay: 0.0005
50000
32 32
1, 0.003: 1.642243, 1.642243 avg, 0.099920 rate, 1.051387 seconds, 128 images, 08-31-2018 18:45:03.000
2, 0.005: 1.586244, 1.636643 avg, 0.099840 rate, 0.036392 seconds, 256 images, 08-31-2018 18:45:03.000
.....
9, 0.023: 1.522952, 1.553780 avg, 0.099282 rate, 0.033507 seconds, 1152 images, 08-31-2018 18:45:03.000
10, 0.026: 1.503915, 1.548793 avg, 0.099202 rate, 0.033441 seconds, 1280 images, 08-31-2018 18:45:03.000
```

As you see, it took 33 ms per batch or iteration on a 4xP100/4x12GB GPU machine instance, versus 710ms or 0.71s per iteration on my MacBook Pro built-as-is or versus 300ms or 0.3s per iteration on my MacBook Pro built with Apple’s *Accelerate* framework. The speed-up is about 10x if we compare the GPU version of YOLOv3 with the version of YOLOv3 optimized on macOS. For the COCO training model, the speed up is between $80/8 = 10x$ (at least) to $200/4 = 50x$ (at best), though.

C5.5 INSTALLING OPENCV ON UBUNTU-18.04

If you have a Ubuntu machine and you want to install OpenCV on it, you can follow the instructions from <https://www.learnopencv.com/install-opencv3-on-ubuntu/>. However, if you end up with no *cv2*.so* file created, here are some tips based on my real experience:

- Make sure your *cmake* is up-to-date.
- Make sure you use *gcc-5* and *g++-5*. For example, with *g++-5*, this is how you would install it:

```
$sudo apt-get install g++-5
$sudo ln -s g++-5 /usr/bin/g++
```

- Make sure that you have *python3* and *numpy* properly installed, as shown below:

```
$sudo apt-get install python-dev python-pip python3-dev python3-pip
$sudo -H pip3 install -U pip numpy
```

- Use a proper *cmake* command, e.g., this is what I used (all on one line):

```
cmake -D CMAKE_BUILD_TYPE=RELEASE -D CMAKE_INSTALL_PREFIX=/usr/local -D
OPENCV_EXTRA_MODULES_PATH=/home/henry/opencv/opencv_contrib-3.4.0/modules -D
BUILD_opencv_python2=OFF -D BUILD_opencv_python3=ON -D INSTALL_PYTHON_EXAMPLES=OFF -D
INSTALL_C_EXAMPLES=OFF -D BUILD_EXAMPLES=OFF -D WITH_CUDA=OFF ..
```

If it works out properly, you should see something similar to the following at the end of executing the above *cmake* command:

```
.....
[100%] Built target opencv_perf_stitching
[100%] Linking CXX shared module ../../lib/python3/cv2.cpython-36m-x86_64-linux-gnu.so
[100%] Built target opencv_python3
henry@ubuntu-18-04:~/opencv/opencv-3.4.0/build$
```

Finally, create a soft link as follows:

```
$ln -s /usr/local/lib/python3.6/dist-packages/cv2.cpython-36m-x86_64-linux-gnu.so cv2.so
```

However, it could be challenging to use OpenCV with YOLOv3 when your GPU-enabled Ubuntu machine is a remote container that does not have X11 built-in.

C.6 PYTORCH

PyTorch is one of the most popular deep learning frameworks for building dynamic neural networks in Python. In this section, I'll help you get a taste of how it can be easily installed on your machine and how easy it is to get it up and running with the simplest yet representative MNIST example you have been very familiar with at this point, given so much you have gone through with this text. You can learn more about it by visiting its website at <https://pytorch.org/>.

To get started, visit the PyTorch's website and decide how you can get it installed on your machine. For example, I made the following selections as shown in Figure C.36. As you see, you can get the installation command by choosing proper OS, Package Manager, Python version and CUDA, based on what you have on your machine.

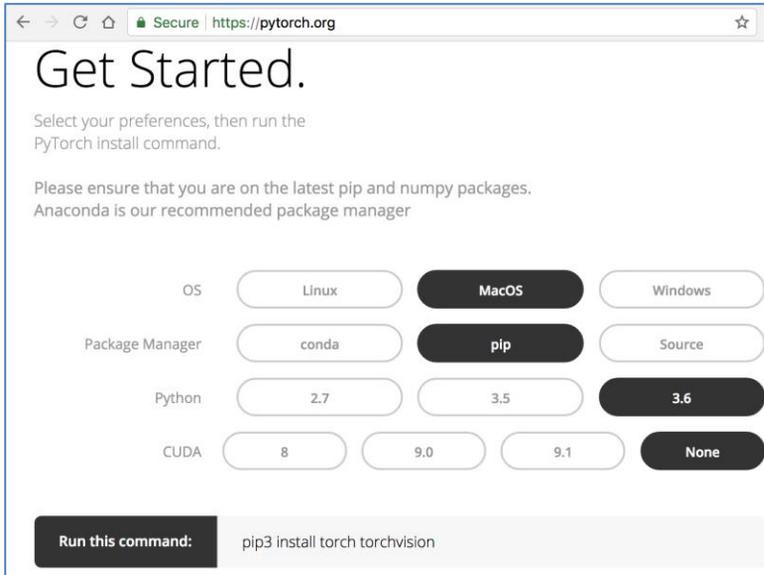


Figure C.36 Decide how to get PyTorch installed on your machine.

In my case, I selected *MacOS*, *pip*, *Python 3.6* and *None* for *CUDA*, as I do not have GPUs on my macOS or Linux VM. Then, I simply executed the following command on my machine to have PyTorch installed on my macOS machine:

```
$ pip3 install torch torchvision
```

Then, I downloaded the PyTorch examples from <https://github.com/pytorch/examples> and expanded it to a directory on my machine. To try out the MNIST example, I changed to the *mnist* directory and executed the following command, with some of the outputs shown following it:

```
henryliu:mnist henryliu$ time python3 main.py
Train Epoch: 1 [0/60000 (0%)] Loss: 2.376790
Train Epoch: 1 [640/60000 (1%)] Loss: 2.332813
...
Train Epoch: 1 [59520/60000 (99%)] Loss: 0.539000
Test set: Average loss: 0.2079, Accuracy: 9416/10000 (94%)
Train Epoch: 2 [0/60000 (0%)] Loss: 0.362165
...
Train Epoch: 2 [59520/60000 (99%)] Loss: 0.337212
Test set: Average loss: 0.1263, Accuracy: 9628/10000 (96%)
Train Epoch: 3 [0/60000 (0%)] Loss: 0.379173
...
Train Epoch: 3 [59520/60000 (99%)] Loss: 0.351524
Test set: Average loss: 0.0981, Accuracy: 9703/10000 (97%)
Train Epoch: 4 [0/60000 (0%)] Loss: 0.066369
...
```

```

Train Epoch: 4 [59520/60000 (99%)] Loss: 0.180767
Test set: Average loss: 0.0834, Accuracy: 9742/10000 (97%)
Train Epoch: 5 [0/60000 (0%)] Loss: 0.344416
...
Train Epoch: 5 [59520/60000 (99%)] Loss: 0.206268
Test set: Average loss: 0.0772, Accuracy: 9765/10000 (98%)
Train Epoch: 6 [0/60000 (0%)] Loss: 0.136501
...
Train Epoch: 6 [59520/60000 (99%)] Loss: 0.109242
Test set: Average loss: 0.0661, Accuracy: 9797/10000 (98%)
Train Epoch: 7 [0/60000 (0%)] Loss: 0.132029
...
Train Epoch: 7 [59520/60000 (99%)] Loss: 0.436856
Test set: Average loss: 0.0606, Accuracy: 9813/10000 (98%)
Train Epoch: 8 [0/60000 (0%)] Loss: 0.277644
...
Train Epoch: 8 [59520/60000 (99%)] Loss: 0.092792
Test set: Average loss: 0.0605, Accuracy: 9812/10000 (98%)
Train Epoch: 9 [0/60000 (0%)] Loss: 0.169455
...
Train Epoch: 9 [59520/60000 (99%)] Loss: 0.102595
Test set: Average loss: 0.0539, Accuracy: 9835/10000 (98%)
Train Epoch: 10 [0/60000 (0%)] Loss: 0.242028
...
Train Epoch: 10 [59520/60000 (99%)] Loss: 0.133413
Test set: Average loss: 0.0483, Accuracy: 9856/10000 (99%)
real3m54.478s
user 3m17.471s
sys 1m4.820s

```

As you see, it took about 4 minutes for 10 epochs to reach a test accuracy of 99%.

If you are curious about how the PyTorch code looks like, Listing C.33 shows the `main.py` script for the above MNIST example. As you see, behind the scene a package named `torch` does all the work. You can spend about ten minutes and find out exactly how PyTorch works at <https://pytorch.org/about/>, which explains all magic behind this great deep learning framework.

Listing C.33 `main.py` code the PyTorch MNIST example

```

1 from __future__ import print_function
2 import argparse
3 import torch
4 import torch.nn as nn
5 import torch.nn.functional as F
6 import torch.optim as optim
7 from torchvision import datasets, transforms
8
9 class Net(nn.Module):
10     def __init__(self):
11         super(Net, self).__init__()

```

```

12     self.conv1 = nn.Conv2d(1, 10, kernel_size=5)
13     self.conv2 = nn.Conv2d(10, 20, kernel_size=5)
14     self.conv2_drop = nn.Dropout2d()
15     self.fc1 = nn.Linear(320, 50)
16     self.fc2 = nn.Linear(50, 10)
17
18     def forward(self, x):
19         x = F.relu(F.max_pool2d(self.conv1(x), 2))
20         x = F.relu(F.max_pool2d(self.conv2_drop(self.conv2(x)), 2))
21         x = x.view(-1, 320)
22         x = F.relu(self.fc1(x))
23         x = F.dropout(x, training=self.training)
24         x = self.fc2(x)
25         return F.log_softmax(x, dim=1)
26
27 def train(args, model, device, train_loader, optimizer, epoch):
28     model.train()
29     for batch_idx, (data, target) in enumerate(train_loader):
30         data, target = data.to(device), target.to(device)
31         optimizer.zero_grad()
32         output = model(data)
33         loss = F.nll_loss(output, target)
34         loss.backward()
35         optimizer.step()
36         if batch_idx % args.log_interval == 0:
37             print('Train Epoch: {} [{}/{}] {:.0f}%)\tLoss:
38                 {:.6f}'.format(
39                 epoch, batch_idx * len(data), len(train_loader.dataset),
40                 100. * batch_idx / len(train_loader), loss.item()))
41
42 def test(args, model, device, test_loader):
43     model.eval()
44     test_loss = 0
45     correct = 0
46     with torch.no_grad():
47         for data, target in test_loader:
48             data, target = data.to(device), target.to(device)
49             output = model(data)
50             test_loss += F.nll_loss(output, target,
51                 size_average=False).item() # sum up batch loss
52             pred = output.max(1, keepdim=True)[1] # get the index of the
53             max log-probability
54             correct += pred.eq(target.view_as(pred)).sum().item()
55
56     test_loss /= len(test_loader.dataset)
57     print('\nTest set: Average loss: {:.4f}, Accuracy: {}/{}
58         {:.0f}%)\n'.format(
59         test_loss, correct, len(test_loader.dataset),
60         100. * correct / len(test_loader.dataset)))
61
62 def main():
63     # Training settings
64     parser = argparse.ArgumentParser(description='PyTorch MNIST Example')

```

```

61     parser.add_argument('--batch-size', type=int, default=64, metavar='N',
62                          help='input batch size for training (default:
63                               64)')
64     parser.add_argument('--test-batch-size', type=int, default=1000,
65                          metavar='N',
66                          help='input batch size for testing (default:
67                               1000)')
68     parser.add_argument('--epochs', type=int, default=10, metavar='N',
69                          help='number of epochs to train (default: 10)')
70     parser.add_argument('--lr', type=float, default=0.01, metavar='LR',
71                          help='learning rate (default: 0.01)')
72     parser.add_argument('--momentum', type=float, default=0.5,
73                          metavar='M',
74                          help='SGD momentum (default: 0.5)')
75     parser.add_argument('--no-cuda', action='store_true', default=False,
76                          help='disables CUDA training')
77     parser.add_argument('--seed', type=int, default=1, metavar='S',
78                          help='random seed (default: 1)')
79     parser.add_argument('--log-interval', type=int, default=10,
80                          metavar='N', help='how many batches to wait before logging training
81                          status')
82     args = parser.parse_args()
83     use_cuda = not args.no_cuda and torch.cuda.is_available()
84     torch.manual_seed(args.seed)
85     device = torch.device("cuda" if use_cuda else "cpu")
86
87     kwargs = {'num_workers': 1, 'pin_memory': True} if use_cuda else {}
88     train_loader = torch.utils.data.DataLoader(
89         datasets.MNIST('../data', train=True, download=True,
90                       transform=transforms.Compose([transforms.ToTensor(),
91                                                     transforms.Normalize((0.1307,), (0.3081,))])),
92         batch_size=args.batch_size, shuffle=True, **kwargs)
93     test_loader = torch.utils.data.DataLoader(
94         datasets.MNIST('../data', train=False,
95                       transform=transforms.Compose([
96                 transforms.ToTensor(),
97                 transforms.Normalize((0.1307,), (0.3081,))
98             ])),
99         batch_size=args.test_batch_size, shuffle=True, **kwargs)
100
101     model = Net().to(device)
102     optimizer = optim.SGD(model.parameters(), lr=args.lr,
103                           momentum=args.momentum)
104
105     for epoch in range(1, args.epochs + 1):
106         train(args, model, device, train_loader, optimizer, epoch)
107         test(args, model, device, test_loader)
108
109 if __name__ == '__main__':
110     main()

```