

Applied Computer Vision

David Vernon
Carnegie Mellon University Africa

vernon@cmu.edu
www.vernon.eu

Lecture 28

Computer Vision and Deep Learning III

Convolutional Neural Networks

Differences from traditional neural networks

- Last layer is the only fully connected FC layer
- New type of layer – convolution CONV – filters the images
 - Hundreds to thousands of filters
 - Filters are learned
 - Most filters in CNNs are square $n \times n$ matrices, where n is odd
- Non-linear activation function, e.g. RELU, is applied to CONV layers
- Multiple sequences of CONV => ReLU layers

Convolutional Neural Networks

Differences from traditional neural networks

- Layers in a CNN are arranged in a 3D volume in three dimensions
 - Width
 - Height
 - Depth
 - Number of channels in an image
 - Number of filters in a layer
- For example, CIFAR10:
 - Input layer $32 \times 32 \times 3$
 - Output layer $1 \times 1 \times 10$... Single vector with ten class scores (or probabilities)

Convolutional Neural Networks

Key benefits

- **Local invariance** / translation invariance: classification not sensitive to position of the object in the image
 - Achieve using pooling layers POOL
- **Compositionality**: later layers build **increasingly rich** features [representations] by building on features detected in earlier layers

Convolutional Neural Networks

Layer Types

Convolutional (CONV)

Activation (ACT or RELU)

Pooling (POOL)


Fully-connected (FC)

Batch normalization (BN)

Dropout (DO)

Stacking layers yields a CNN

INPUT => CONV => RELU => FC => SOFTMAX



Only layers that have parameters
that are learned during the training process

Convolutional Neural Networks

Layer Types

Convolutional (CONV)

Activation (ACT or RELU)

Pooling (POOL)

Fully-connected (FC)

Batch normalization (BN)

Dropout (DO)

Activation and Dropout layers
are not considered "true" layers
but are often included
to make the architecture explicit

Stacking layers yields a CNN

INPUT => CONV => RELU => FC => SOFTMAX



Convolutional Neural Networks

Layer Types

Convolutional (CONV)

Activation (ACT or RELU)

Pooling (POOL)

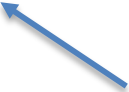
Fully-connected (FC)

Batch normalization (BN)

Dropout (DO)

Stacking layers yields a CNN

INPUT => CONV => RELU => FC => SOFTMAX



Often omitted and
simply assumed it follows FC

Convolutional Neural Networks

Layer Types

Convolutional (CONV)

Activation (ACT or RELU)

Pooling (POOL)

Fully-connected (FC)

Batch normalization (BN)

Dropout (DO)



The most important when defining the network architecture

Stacking layers yields a CNN

INPUT => CONV => RELU => FC => SOFTMAX

Convolutional Neural Networks

Layer Types

Convolutional (CONV)

Practically assumed to be part of the architecture

Activation (ACT or RELU)

Often omitted from architecture table/diagram to save space

Pooling (POOL)

Fully-connected (FC)

But are implicitly assumed to be part of the architecture

Batch normalization (BN)

Dropout (DO)

Stacking layers yields a CNN

INPUT => CONV => RELU => FC => SOFTMAX


Convolutional Neural Networks

Convolutional Layers CONV

- The CONV layer parameters consist of a set of K learnable filters (kernels)
- Width normally equals height (i.e. square kernel)
- Usually small size (e.g. 3×3 , 5×5 , ...)
- Extend throughout the full depth of the volume
 - Number of channels (input layer)
 - Number of filters applied in previous layer (deeper in the network)

Convolutional Neural Networks

Convolutional Layers CONV



Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Convolutional Neural Networks


Convolutional Layers CONV

- Depth
 - Number of filters K in the CONV layer that connect to a local region in the input volume
 - The set of filters that are “looking at” the same (x, y) location of the input is call the **depth column**
- Stride
 - The **step size S** (in pixels) between each application of a convolution kernel
 - Typically, $S = 1$ or $S = 2$
 - Larger strides imply less overlap in receptive field
 - Larger strides generate outputs with **reduced spatial dimensions**

Convolutional Neural Networks

Convolutional Layers CONV

- Padding
 - Enlarge the input to ensure that the **output spatial dimensions = input spatial dimensions**
 - Assuming stride of 1



- To ensure the can be tiled such that they fit across the input volume

$$(W - F + 2 P) / S + 1$$
must be an integer

Size of [square] input Receptive Field Padding Stride

Convolutional Neural Networks

Convolutional Layers CONV

To summarize, the CONV layer in the same, elegant manner as Karpathy [121]:

- Accepts an input volume of size $W_{input} \times H_{input} \times D_{input}$ (the input sizes are normally square, so it's common to see $W_{input} = H_{input}$).
- Requires four parameters:
 1. The number of filters K (which controls the *depth* of the output volume).
 2. The receptive field size F (the size of the K kernels used for convolution and is nearly always *square*, yielding an $F \times F$ kernel).
 3. The stride S .
 4. The amount of zero-padding P .
- The output of the CONV layer is then $W_{output} \times H_{output} \times D_{output}$, where:
 - $W_{output} = ((W_{input} - F + 2P)/S) + 1$
 - $H_{output} = ((H_{input} - F + 2P)/S) + 1$
 - $D_{output} = K$


121. Andrej Karpathy. *Convolutional Networks*. <http://cs231n.github.io/convolutional-networks/>

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Convolutional Neural Networks

Activation Layers **ACT, RELU, ELU, ...**

- Apply a nonlinear activation function after each CONV layer
- Technically not "layers" since no weights are learned during training
- Sometimes omitted (and assumed to be present)



INPUT => CONV => RELU => FC

INPUT => CONV => FC

- Dimensions of the output volume is the same as the input volume
 - Activation function is applied to each neuron individually

Convolutional Neural Networks

Pooling Layers **POOL**

- Used to reduce size (width, height) of an input volume
- Can also use CONV with a stride > 2
- Insert POOL in between consecutive CONV layers


INPUT \Rightarrow CONV \Rightarrow RELU \Rightarrow POOL \Rightarrow CONV \Rightarrow RELU \Rightarrow POOL \Rightarrow FC

- Also helps control overfitting
- Operate on each depth slice independently
 - Max function
 - Average function

Convolutional Neural Networks

Pooling Layers **POOL**

- Typically, pool size of 2×2 ; images > 200 pixels may use 3×3
- Stride of 1 or 2



Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Convolutional Neural Networks

Fully-connected Layers **FC**

- Neurons in FC layers are fully-connected to all activation outputs in the previous layer
- FC are always placed at the end of the network (and only at the end)
- It is common to use two FC layers prior to applying the softmax classifier

INPUT => CONV => RELU => POOL => CONV => RELU => POOL => FC => FC

Convolutional Neural Networks

Batch Normalization BN

- Batch normalization layers are used to normalize the activations of a given input volume before passing them to the next layer
 - Reduces the number of epochs needed to train the network
 - “Stabilizes” training: makes learning rate and regularization easier to tune
 - Lower final loss & more stable loss curve
 - Helps prevent overfitting
 - Use it in nearly every situation
- Apply BN after RELU: INPUT => CONV => RELU => BN ...

Convolutional Neural Networks

Batch Normalization BN

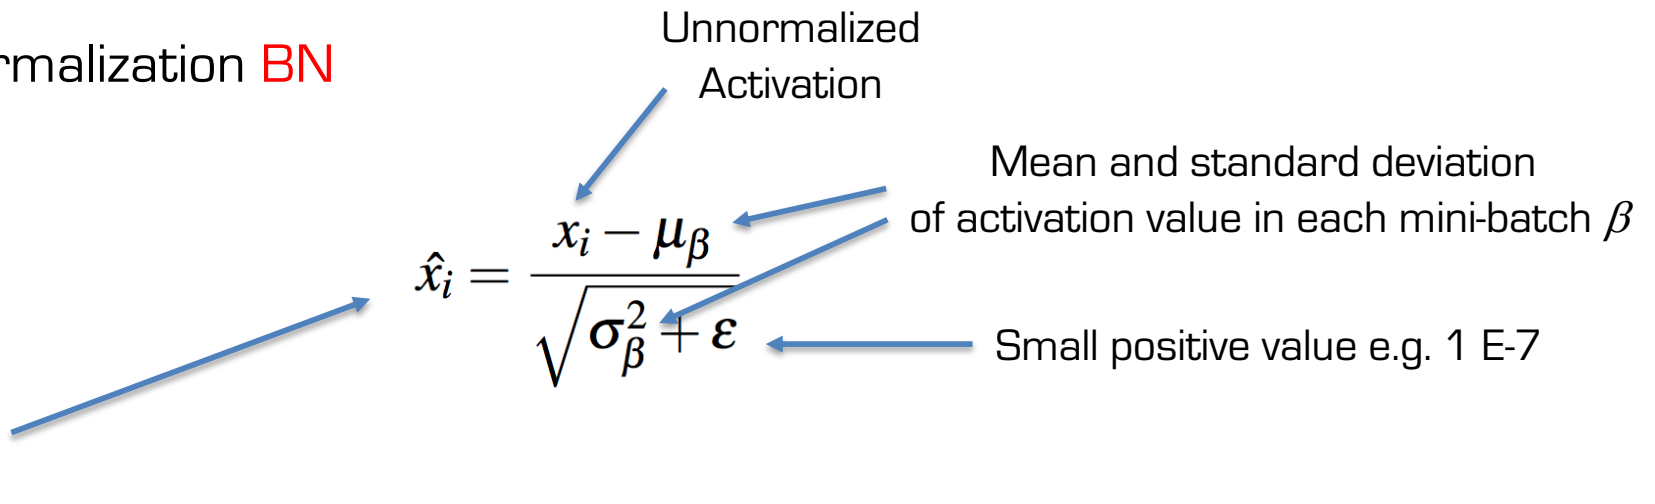
Unnormalized Activation

Mean and standard deviation of activation value in each mini-batch β

Small positive value e.g. 1 E-7

$$\hat{x}_i = \frac{x_i - \mu_\beta}{\sqrt{\sigma_\beta^2 + \epsilon}}$$

Normalized Activation



Produces values with approx. zero mean and unit variance, i.e. zero-centred values

$$\mu_\beta = \frac{1}{M} \sum_{i=1}^m x_i$$


At testing time, replace mini-batch μ_β and σ_β with running averages of μ_β and σ_β computed during the training process

$$\sigma_\beta^2 = \frac{1}{m} \sum_{i=1}^m (x_i - \mu_\beta)^2$$

Convolutional Neural Networks

Dropout DO

- Regularization technique to help prevent overfitting
 - Increasing testing accuracy
 - Possibly at the expense of training accuracy
- Randomly disconnect connections between two FC layers; probability p (e.g. 0.5)
- ... CONV => RELU => POOL => FC => DO => FC => DO => FC




Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Convolutional Neural Networks

Dropout DO

- Randomly dropping connections ensures no single node in the network is responsible for “activating” when presented with a given pattern
- Ensures there are multiple redundant nodes
- Help the network to **generalize**



Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Convolutional Neural Networks

Common Architectures and Training Patterns

INPUT \Rightarrow $[[\text{CONV} \Rightarrow \text{RELU}] * N \Rightarrow \text{POOL?}] * M \Rightarrow [\text{FC} \Rightarrow \text{RELU}] * K \Rightarrow \text{FC}$

Typically:

$$0 \leq N \leq 3$$

$$M \geq 0$$

$$0 \leq K \leq 2$$

AlexNet-like:

INPUT \Rightarrow $[\text{CONV} \Rightarrow \text{RELU} \Rightarrow \text{POOL}] * 2 \Rightarrow [\text{CONV} \Rightarrow \text{RELU}] * 3 \Rightarrow \text{POOL} \Rightarrow$
 $[\text{FC} \Rightarrow \text{RELU} \Rightarrow \text{DO}] * 2 \Rightarrow \text{SOFTMAX}$

VGGNet

INPUT \Rightarrow $[\text{CONV} \Rightarrow \text{RELU}] * 2 \Rightarrow \text{POOL} \Rightarrow [\text{CONV} \Rightarrow \text{RELU}] * 2 \Rightarrow \text{POOL} \Rightarrow$
 $[\text{CONV} \Rightarrow \text{RELU}] * 3 \Rightarrow \text{POOL} \Rightarrow [\text{CONV} \Rightarrow \text{RELU}] * 3 \Rightarrow \text{POOL} \Rightarrow$
 $[\text{FC} \Rightarrow \text{RELU} \Rightarrow \text{DO}] * 2 \Rightarrow \text{SOFTMAX}$

Convolutional Neural Networks

Common Architectures and Training Patterns

Apply deeper networks when we have

lots of training data

a sufficiently challenging classification problem

Deep Learning with Keras and Python

Installation instructions

<https://www.pyimagesearch.com/2017/09/25/configuring-ubuntu-for-deep-learning-with-python/>

Possible issues (these arose with Ubuntu 18 but may not with Ubuntu 16.04)

Ubuntu system dependencies

```
$ sudo apt-get install libjpeg-dev libtiff5-dev libjasper-dev libpng12-dev
```

Error: No installable candidates

Solution: before this command, execute:

```
$ sudo add-apt-repository "deb http://security.ubuntu.com/ubuntu xenial-security main"
```

```
$ sudo apt-get install libjpeg-dev libtiff5-dev libjasper-dev libpng12-dev
```

Compiling OpenCV

Error: Possible removal of the OpenCV folder

Do not execute the direct command (provided in the instructions)

```
$ rm -rf opencv-3.3.0 opencv.zip
```

```
$ rm -rf opencv_contrib-3.3.0 opencv_contrib.zip
```

Solution:

```
$ rm -rf opencv.zip
```

```
$ rm -rf opencv_contrib-3.3.0 opencv_contrib.zip
```


Deep Learning with Keras and Python

Software		✓	+	⋮
⋮	📄 Software Development Environment	✓	⋮	
⋮	📎 ACV.zip	✓	⋮	
⋮	📎 opencv-2.4.10.exe	✓	⋮	
⋮	📎 chapter12-first_cnn.zip	✓	⋮	
⋮	📎 chapter15-minivggnet.zip	✓	⋮	
⋮	📎 chapter20-out_of_the_box.zip	✓	⋮	

Deep Learning with Keras and Python

When you import the Keras library into a Python script, it generate a keras.json configuration file in `~/.keras/keras.json`

```
1  {  
2      "epsilon": 1e-07,  
3      "floatx": "float32",  
4      "image_data_format": "channels_last",  
5      "backend": "tensorflow"  
6  }
```



rows, columns, channels

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

ShallowNet with Keras and Python

ShallowNet: INPUT => CONV => RELU => FC

```
--- pyimagesearch
|   |--- __init__.py
|   |--- datasets
|   |--- nn
|   |   |--- __init__.py
...
|   |   |--- conv
|   |   |   |--- __init__.py
|   |   |   |--- shalownet.py
|   |--- preprocessing
```

New

We will put the CNN implementations here

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

ShallowNet with Keras and Python

ShallowNet: INPUT => CONV => RELU => FC

`shallownet.py`

```
1 # import the necessary packages
2 from keras.models import Sequential
3 from keras.layers.convolutional import Conv2D
4 from keras.layers.core import Activation
5 from keras.layers.core import Flatten
6 from keras.layers.core import Dense
7 from keras import backend as K
```

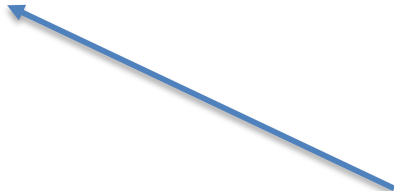

Keras implementation
of the convolutional layer



Activation layer



The Flatten classes takes our
multi-dimensional volume and
“flattens” it into a 1D array
prior to feeding the inputs into
the Dense (i.e, fully-connected)
layers.



Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

ShallowNet with Keras and Python

ShallowNet: INPUT => CONV => RELU => FC

shallownet.py

```
9  class ShallowNet:
10      @staticmethod
11      def build(width, height, depth, classes):
12          # initialize the model along with the input shape to be
13          # "channels last"
14          model = Sequential()
15          inputShape = (height, width, depth)
16
17          # if we are using "channels first", update the input shape
18          if K.image_data_format() == "channels_first":
19              inputShape = (depth, height, width)
```

Image dimensions

Number of classes

Defensive programming

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

ShallowNet with Keras and Python

ShallowNet: INPUT => CONV => RELU => FC

shallownet.py

```
21 # define the first (and only) CONV => RELU layer
22 model.add(Conv2D(32, (3, 3), padding="same",
23                 input_shape=inputShape))
24 model.add(Activation("relu"))
26 # softmax classifier
27 model.add(Flatten())
28 model.add(Dense(classes))
29 model.add(Activation("softmax"))
30
31 # return the constructed network architecture
32 return model
```

32 filters, all 3x3, padding ensure size of output = size of input

To apply the fully-connected layer:

1. flatten the multi-dimensional representation into a 1D list
2. Add a dense layer using the same number of nodes as our output class labels
3. Apply a softmax activation function to will generate the class label probabilities for each class


Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

ShallowNet with Keras and Python

CIFAR-10

- 60,000 images
- 32 x 32 x 3 (RGB) ...
3072 element feature vector
- 10 classes: airplanes, automobiles, birds, cats, deer, dogs, frogs, horses, ships, and trucks


airplane




automobile




bird




cat




deer




dog




frog




horse



ship



truck



Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

ShallowNet with Keras and Python

Apply to the CIFAR-10 dataset:

`shallownet_cifar10.py`

```
1 # import the necessary packages
2 from sklearn.preprocessing import LabelBinarizer
3 from sklearn.metrics import classification_report
4 from pyimagesearch.nn.conv import ShallowNet
5 from keras.optimizers import SGD
6 from keras.datasets import cifar10
7 import matplotlib.pyplot as plt
8 import numpy as np
9
10 # load the training and testing data, then scale it into the
11 # range [0, 1]
12 print("[INFO] loading CIFAR-10 data...")
13 ((trainX, trainY), (testX, testY)) = cifar10.load_data()
14 trainX = trainX.astype("float") / 255.0
15 testX = testX.astype("float") / 255.0
```

Driver script to load a dataset,
preprocess it, and then train
the network

Preprocessing and channel ordering handled automatically in this function. If this is the first time calling `cifar10.load_data()`, the function will load the dataset for you. The file is ~170Mb so be patient. Once downloaded, it is cached and doesn't need to be downloaded again.

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

ShallowNet with Keras and Python

Apply to the CIFAR-10 dataset: `shallownet_cifar10.py`

```
17 # convert the labels from integers to vectors
```

```
18 lb = LabelBinarizer()
```

```
19 trainY = lb.fit_transform(trainY)
```

```
20 testY = lb.transform(testY)
```

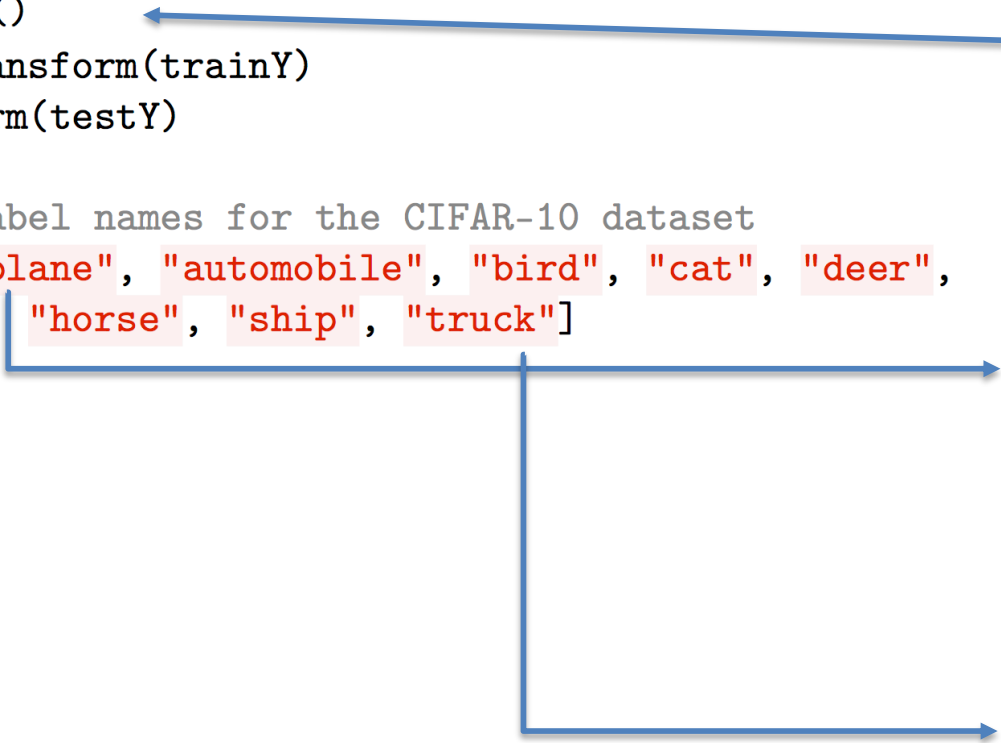
```
21
```

```
22 # initialize the label names for the CIFAR-10 dataset
```

```
23 labelNames = ["airplane", "automobile", "bird", "cat", "deer",
```

```
24               "dog", "frog", "horse", "ship", "truck"]
```

Labels are **one-hot encoded**, i.e. represented as a vector of 0 and 1, where 1 represents the correct class



```
[1 0 0 0 0 0 0 0 0 0]
[0 1 0 0 0 0 0 0 0 0]
[0 0 1 0 0 0 0 0 0 0]
[0 0 0 1 0 0 0 0 0 0]
[0 0 0 0 1 0 0 0 0 0]
[0 0 0 0 0 1 0 0 0 0]
[0 0 0 0 0 0 1 0 0 0]
[0 0 0 0 0 0 0 1 0 0]
[0 0 0 0 0 0 0 0 1 0]
[0 0 0 0 0 0 0 0 0 1]
```

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

ShallowNet with Keras and Python

Apply to the CIFAR-10 dataset: `shallownet_cifar10.py`

```
26 # initialize the optimizer and model
27 print("[INFO] compiling model...")
28 opt = SGD(lr=0.01)
29 model = ShallowNet.build(width=32, height=32, depth=3, classes=10)
30 model.compile(loss="categorical_crossentropy", optimizer=opt,
31               metrics=["accuracy"])
32
33 # train the network
34 print("[INFO] training network...")
35 H = model.fit(trainX, trainY, validation_data=(testX, testY),
36               batch_size=32, epochs=40, verbose=1)
```

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

ShallowNet with Keras and Python

Apply to the CIFAR-10 dataset: `shallownet_cifar10.py`

```
38 # evaluate the network
39 print("[INFO] evaluating network...")
40 predictions = model.predict(testX, batch_size=32)
41 print(classification_report(testY.argmax(axis=1),
42                             predictions.argmax(axis=1), target_names=labelNames))
```

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

ShallowNet with Keras and Python

Apply to the CIFAR-10 dataset: `shallownet_cifar10.py`

```
44 # plot the training loss and accuracy
45 plt.style.use("ggplot")
46 plt.figure()
47 plt.plot(np.arange(0, 40), H.history["loss"], label="train_loss")
48 plt.plot(np.arange(0, 40), H.history["val_loss"], label="val_loss")

49 plt.plot(np.arange(0, 40), H.history["acc"], label="train_acc")
50 plt.plot(np.arange(0, 40), H.history["val_acc"], label="val_acc")
51 plt.title("Training Loss and Accuracy")
52 plt.xlabel("Epoch #")
53 plt.ylabel("Loss/Accuracy")
54 plt.legend()
55 plt.show()
```

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

ShallowNet with Keras and Python

To train ShallowNet on CIFAR-10: `$ python shallownet_cifar10.py`

	precision	recall	f1-score	support
airplane	0.62	0.68	0.65	1000
automobile	0.79	0.64	0.71	1000
bird	0.43	0.46	0.44	1000
cat	0.42	0.38	0.40	1000
deer	0.52	0.51	0.52	1000
dog	0.44	0.57	0.50	1000
frog	0.74	0.61	0.67	1000
horse	0.71	0.61	0.66	1000
ship	0.65	0.77	0.70	1000
truck	0.67	0.66	0.66	1000
avg / total	0.60	0.59	0.59	10000




Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

ShallowNet with Keras and Python

To train ShallowNet on CIFAR-10: `$ python shallownet_cifar10.py`

	precision	recall	f1-score	support
airplane	0.70	0.54	0.61	1000
automobile	0.68	0.77	0.72	1000
bird	0.53	0.37	0.44	1000
cat	0.36	0.57	0.45	1000
deer	0.60	0.45	0.51	1000
dog	0.54	0.44	0.49	1000
frog	0.64	0.76	0.69	1000
horse	0.64	0.70	0.67	1000
ship	0.75	0.69	0.72	1000
truck	0.67	0.71	0.69	1000
avg / total	0.61	0.60	0.60	10000

[43 s / epoch on Ubuntu virtual machine]



Recall: 3072-1024-512-10 MLP

CIFAR-10

	precision	recall	f1-score	support
airplane	0.63	0.66	0.64	1000
automobile	0.69	0.65	0.67	1000
bird	0.48	0.43	0.45	1000
cat	0.40	0.38	0.39	1000
deer	0.52	0.51	0.51	1000
dog	0.48	0.47	0.48	1000
frog	0.64	0.63	0.64	1000
horse	0.63	0.62	0.63	1000
ship	0.64	0.74	0.69	1000
truck	0.59	0.65	0.62	1000
avg / total	0.57	0.57	0.57	10000



Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

MiniVGGNet

VGGNet (or simply VGG)

- Introduced by Simonyan and Zisserman in 2014: "Very Deep Learning Convolutional Neural Networks for Large-Scale Image Recognition"
- Convolution filters are 3 x 3 throughout the architecture
- Stacking multiple CONV => RELU layer sets
 - more repetitions deeper in the architecture
- 16 & 19 layers
- 2nd in ImageNet classification challenge

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

MiniVGGNet

MiniVGGNet

- INPUT => CONV => ACT => BN => CONV => ACT => BN => POOL => DO => CONV => ACT => BN => CONV => ACT => BN => POOL => DO => FC => ACT => BN => DO => FC => SOFTMAX
- First two CONV layers learn 32 filters
- Second two CONV layer learn 64 filters
- POOL layers perform max pooling over a 2 x 2 window with a 2 x 2 stride
- Input is 32 x 32 x 3 ... colour CIFAR-10 images

MiniVGGNet

MiniVGGNet


Layer Type	Output Size	Filter Size / Stride
INPUT IMAGE	$32 \times 32 \times 3$	
CONV	$32 \times 32 \times 32$	$3 \times 3, K = 32$
ACT	$32 \times 32 \times 32$	
BN	$32 \times 32 \times 32$	
CONV	$32 \times 32 \times 32$	$3 \times 3, K = 32$
ACT	$32 \times 32 \times 32$	
BN	$32 \times 32 \times 32$	
POOL	$16 \times 16 \times 32$	2×2
DROPOUT	$16 \times 16 \times 32$	
CONV	$16 \times 16 \times 64$	$3 \times 3, K = 64$
ACT	$16 \times 16 \times 64$	
BN	$16 \times 16 \times 64$	
CONV	$16 \times 16 \times 64$	$3 \times 3, K = 64$
ACT	$16 \times 16 \times 64$	
BN	$16 \times 16 \times 64$	
POOL	$8 \times 8 \times 64$	2×2
DROPOUT	$8 \times 8 \times 64$	
FC	512	
ACT	512	
BN	512	
DROPOUT	512	
FC	10	
SOFTMAX	10	

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

MiniVGGNet

MiniVGGNet

```
--- pyimagesearch
|   |--- __init__.py
|   |--- nn
|   |   |--- __init__.py
...
|   |   |--- conv
|   |   |   |--- __init__.py
|   |   |   |--- lenet.py
|   |   |   |--- minivggnet.py
|   |   |   |--- shallownet.py
```



MiniVGGNet

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

MiniVGGNet

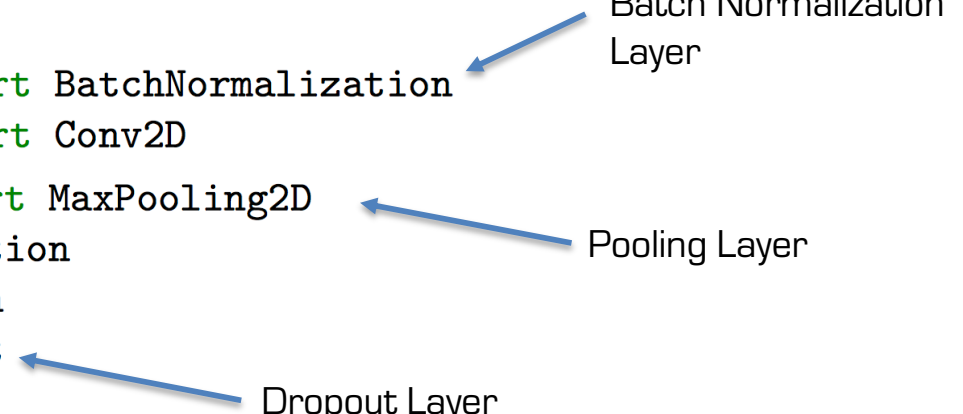
MiniVGGNet `minivggnet.py`

```
1  # import the necessary packages
2  from keras.models import Sequential
3  from keras.layers.normalization import BatchNormalization
4  from keras.layers.convolutional import Conv2D
5  from keras.layers.convolutional import MaxPooling2D
6  from keras.layers.core import Activation
7  from keras.layers.core import Flatten
8  from keras.layers.core import Dropout
9  from keras.layers.core import Dense
10 from keras import backend as K
```

Batch Normalization Layer

Pooling Layer

Dropout Layer



Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

MiniVGGNet

MiniVGGNet

minivggnet.py

Batch normalization operates on the features, i.e. channels, so in order to apply BN, we need to know which axis to normalize over (i.e. the channels axis)

```
12 class MiniVGGNet:
13     @staticmethod
14     def build(width, height, depth, classes):
15         # initialize the model along with the input shape to be
16         # "channels last" and the channels dimension itself
17         model = Sequential()
18         inputShape = (height, width, depth)
19         chanDim = -1
20
21         # if we are using "channels first", update the input shape
22         # and channels dimension
23         if K.image_data_format() == "channels_first":
24             inputShape = (depth, height, width)
25             chanDim = 1
```

Setting chanDim = -1 implies that the index of the channel dimension last in the input shape (i.e., channels last ordering).

Create the model

However, if we are using channels first ordering, we need need to update the inputShape and set chanDim = 1

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

MiniVGGNet

MiniVGGNet minivggnet.py

```
27       # first CONV => RELU => CONV => RELU => POOL layer set
28       model.add(Conv2D(32, (3, 3), padding="same",
29                        input_shape=inputShape))
30       model.add(Activation("relu"))
31       model.add(BatchNormalization(axis=chanDim))
32       model.add(Conv2D(32, (3, 3), padding="same"))
33       model.add(Activation("relu"))
34       model.add(BatchNormalization(axis=chanDim))
35       model.add(MaxPooling2D(pool_size=(2, 2)))
36       model.add(Dropout(0.25))
```

32 3x3 filters

Dropout propability of 0.25

[CONV => RELU => BN] * 2 => POOL => DO

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

MiniVGGNet

MiniVGGNet minivggnet.py

64 3x3 filters

```
38     # second CONV => RELU => CONV => RELU => POOL layer set
39     model.add(Conv2D(64, (3, 3), padding="same"))
40     model.add(Activation("relu"))
41     model.add(BatchNormalization(axis=chanDim))
42     model.add(Conv2D(64, (3, 3), padding="same"))
43     model.add(Activation("relu"))
44     model.add(BatchNormalization(axis=chanDim))
45     model.add(MaxPooling2D(pool_size=(2, 2)))
46     model.add(Dropout(0.25))
```

Dropout propability of 0.25

[CONV => RELU => BN] * 2 => POOL => DO

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

MiniVGGNet

MiniVGGNet minivggnet.py

```
48         # first (and only) set of FC => RELU layers
49         model.add(Flatten())
50         model.add(Dense(512))
51         model.add(Activation("relu"))
52         model.add(BatchNormalization())
53         model.add(Dropout(0.5))
54
55         # softmax classifier
56         model.add(Dense(classes))
57         model.add(Activation("softmax"))
58
59         # return the constructed network architecture
60         return model
```

FC => ACT => BN => DO => FC => SOFTMAX

Dropout probability of 0.5


Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

MiniVGGNet

CIFAR-10

- 60,000 images
- 32 x 32 x 3 (RGB) ...
3072 element feature vector
- 10 classes: airplanes, automobiles, birds, cats, deer, dogs, frogs, horses, ships, and trucks


airplane




automobile




bird




cat




deer




dog




frog




horse



ship



truck



Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

MiniVGGNet

MiniVGGNet

`minivggnet_cifar10.py`

Driver script to load a dataset, preprocess it, and then train the network

```
1 # set the matplotlib backend so figures can be saved in the background
2 import matplotlib
3 matplotlib.use("Agg")
4
5 # import the necessary packages
6 from sklearn.preprocessing import LabelBinarizer
7 from sklearn.metrics import classification_report
8 from pyimagesearch.nn.conv import MiniVGGNet
9 from keras.optimizers import SGD
10 from keras.datasets import cifar10
11 import matplotlib.pyplot as plt
12 import numpy as np
13 import argparse
```

Non-interactive: save plot to file


Parse command line arguments

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

MiniVGGNet

MiniVGGNet minivggnet_cifar10.py

```
15 # construct the argument parse and parse the arguments
16 ap = argparse.ArgumentParser()
17 ap.add_argument("-o", "--output", required=True,
18                 help="path to the output loss/accuracy plot")
19 args = vars(ap.parse_args())
```



Add a single argument: the path and filename of the file to which the output training and loss plot is saved

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

MiniVGGNet

MiniVGGNet minivggnet_cifar10.py

```
21 # load the training and testing data, then scale it into the
22 # range [0, 1]
23 print("[INFO] loading CIFAR-10 data...")
24 ((trainX, trainY), (testX, testY)) = cifar10.load_data()
25 trainX = trainX.astype("float") / 255.0
26 testX = testX.astype("float") / 255.0
27
28 # convert the labels from integers to vectors
29 lb = LabelBinarizer()
30 trainY = lb.fit_transform(trainY)
31 testY = lb.transform(testY)
32
33 # initialize the label names for the CIFAR-10 dataset
34 labelNames = ["airplane", "automobile", "bird", "cat", "deer",
35               "dog", "frog", "horse", "ship", "truck"]
```

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

MiniVGGNet

MiniVGGNet minivggnet_cifar10.py

```
37 # initialize the optimizer and model
38 print("[INFO] compiling model...")
39 opt = SGD(lr=0.01, decay=0.01 / 40, momentum=0.9, nesterov=True)
40 model = MiniVGGNet.build(width=32, height=32, depth=3, classes=10)
41 model.compile(loss="categorical_crossentropy", optimizer=opt,
42               metrics=["accuracy"])
43
44 # train the network
45 print("[INFO] training network...")
46 H = model.fit(trainX, trainY, validation_data=(testX, testY),
47               batch_size=64, epochs=40, verbose=1)
```

Set learning rate, momentum, and acceleration for stochastic gradient descent. The decay argument causes the learning rate to reduce with time: typically set to learning rate / # epochs

Set loss function, optimizer (defined above), and metric

Train the model over 40 epochs with batch size of 64

MiniVGGNet

MiniVGGNet `minivggnet.py`

```
49 # evaluate the network
50 print("[INFO] evaluating network...")
51 predictions = model.predict(testX, batch_size=64)
52 print(classification_report(testY.argmax(axis=1),
53                             predictions.argmax(axis=1), target_names=labelNames))
```

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

MiniVGGNet

MiniVGGNet `minivggnet.py`

```
55 # plot the training loss and accuracy
56 plt.style.use("ggplot")
57 plt.figure()
58 plt.plot(np.arange(0, 40), H.history["loss"], label="train_loss")
59 plt.plot(np.arange(0, 40), H.history["val_loss"], label="val_loss")
60 plt.plot(np.arange(0, 40), H.history["acc"], label="train_acc")
61 plt.plot(np.arange(0, 40), H.history["val_acc"], label="val_acc")
62 plt.title("Training Loss and Accuracy on CIFAR-10")
63 plt.xlabel("Epoch #")
64 plt.ylabel("Loss/Accuracy")
65 plt.legend()
66 plt.savefig(args["output"])
```

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017


MiniVGGNet

MiniVGGNet

```
$ python minivggnet_cifar10.py --output output/cifar10_minivggnet_with_bn.png
```

	precision	recall	f1-score	support
airplane	0.88	0.81	0.85	1000
automobile	0.93	0.89	0.91	1000
bird	0.83	0.68	0.75	1000
cat	0.69	0.65	0.67	1000
deer	0.74	0.85	0.79	1000
dog	0.72	0.77	0.74	1000
frog	0.85	0.89	0.87	1000
horse	0.85	0.87	0.86	1000
ship	0.89	0.91	0.90	1000
truck	0.88	0.91	0.90	1000
avg / total	0.83	0.82	0.82	10000

↙ Drops to 0.79 if we leave out
batch normalization



Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Pre-trained CNNs in Keras

The Keras library ships with five CNNs that have been pre-trained on the ImageNet dataset:


1. VGG16
2. VGG19
3. ResNet50
4. Inception V3
5. Xception

The goal of the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) is to train a model that can correctly classify an input image into **1000 separate object categories**

Pre-trained CNNs in Keras



ImageNet Object Recognition Challenge:
1.2 million training images, 1000 classes



[Deng et al. CVPR 2009]

Pre-trained CNNs in Keras

The Keras library ships with five CNNs that have been pre-trained on the ImageNet dataset:

1. VGG16
2. VGG19
3. ResNet50
4. Inception V3
5. Xception


The goal of the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) is to train a model that can correctly classify an input image into **1000 separate object categories**

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Pre-trained CNNs in Keras

VGG -16

K. Simonyan and A. Zisserman, “Very Deep Convolutional Networks for Large- Scale Image Recognition” , arXiv technical report, 2014.



Credit: <https://www.cs.toronto.edu/~frossard/post/vgg16/>

Pre-trained CNNs in Keras



VGG16 / VGG19

- Only 3 x 3 convolutional layers
- Volume size is reduced by max pooling
- Two fully-connected layers each with 4096 nodes
- Softmax classifier
- Very slow to train
- Network weights file size: 533 MB (VGG16) and 574 MB (VGG19)

Pre-trained CNNs in Keras

ResNet50

K. He et al. “Deep Residual Learning for Image Recognition”,
arXiv technical report, 2015.



Pre-trained CNNs in Keras

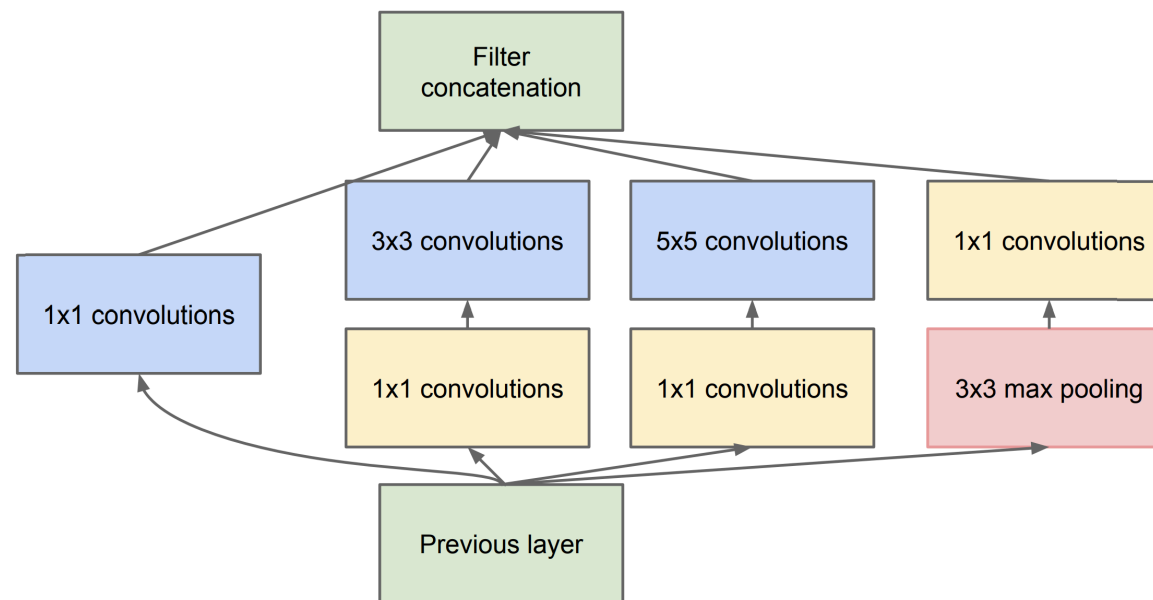
ResNet

- ResNet50 in Keras
- Can be much deeper: up to 200 for ImageNet and 1000 for CIFAR-10
- Network weights file size: 102 MB for ResNet50

Pre-trained CNNs in Keras

Inception V3

C. Szegedy et al. "Going Deeper with Convolutions". In: *Computer Vision and Pattern Recognition*, 2015.



Inception Module

Credit: <https://arxiv.org/abs/1409.4842>

Pre-trained CNNs in Keras

Inception V3

- The inception module is a multi-level feature extractor
- Output of each module is stacked along the channel dimension before being fed into the next layer
- Originally called GoogLeNet
- Subsequently called Inception vN, where N denotes the version number
- Keras version is V3 from C. Szegedy et al. “Rethinking the Inception Architecture for Computer Vision”, 2015
- Network weights file size: 96 MB

Pre-trained CNNs in Keras

Xception

- François Chollet, creator and chief maintainer of the Keras Library
- François Chollet. “Xception: Deep Learning with Depthwise Separable Convolutions”
- Network weights file size: 91 MB

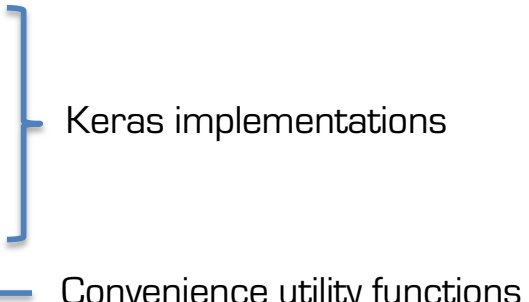
Pre-trained CNNs in Keras

imagenet_pretrained.py

```
1  # import the necessary packages
2  from keras.applications import ResNet50
3  from keras.applications import InceptionV3
4  from keras.applications import Xception # TensorFlow ONLY
5  from keras.applications import VGG16
6  from keras.applications import VGG19
7  from keras.applications import imagenet_utils
8  from keras.applications.inception_v3 import preprocess_input
9  from keras.preprocessing.image import img_to_array
10 from keras.preprocessing.image import load_img
11 import numpy as np
12 import argparse
13 import cv2
```

Keras implementations

Convenience utility functions


A blue bracket on the right side of the code block groups lines 2 through 6 under the label "Keras implementations". A blue arrow points from the label "Convenience utility functions" to line 7, which imports 'imagenet_utils'.

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Pre-trained CNNs in Keras

imagenet_pretrained.py

```
15 # construct the argument parse and parse the arguments
16 ap = argparse.ArgumentParser()
17 ap.add_argument("-i", "--image", required=True,
18                 help="path to the input image")
19 ap.add_argument("-model", "--model", type=str, default="vgg16",
20                 help="name of pre-trained network to use")
21 args = vars(ap.parse_args())
```



Add a two arguments:
the path and filename of the file to use as input
and the model to use

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Pre-trained CNNs in Keras

imagenet_pretrained.py

```
23 # define a dictionary that maps model names to their classes
24 # inside Keras
25 MODELS = {
26     "vgg16": VGG16,
27     "vgg19": VGG19,
28     "inception": InceptionV3,
29     "xception": Xception, # TensorFlow ONLY
30     "resnet": ResNet50
31 }
32
33 # ensure a valid model name was supplied via command line argument
34 if args["model"] not in MODELS.keys():
35     raise AssertionError("The --model command line argument should "
36                          "be a key in the 'MODELS' dictionary")
```

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Pre-trained CNNs in Keras

imagenet_pretrained.py

```
38 # initialize the input image shape (224x224 pixels) along with
39 # the pre-processing function (this might need to be changed
40 # based on which model we use to classify our image)
41 inputShape = (224, 224)
42 preprocess = imagenet_utils.preprocess_input
43
44 # if we are using the InceptionV3 or Xception networks, then we
45 # need to set the input shape to (299x299) [rather than (224x224)]
46 # and use a different image processing function
47 if args["model"] in ("inception", "xception"):
48     inputShape = (299, 299)
49     preprocess = preprocess_input
```

VGG16, VGG19, and ResNet
use 224x224 images

But InceptionV3 and Xception
use 229x229

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Pre-trained CNNs in Keras

imagenet_pretrained.py

```
51 # load our the network weights from disk (NOTE: if this is the
52 # first time you are running this script for a given network, the
53 # weights will need to be downloaded first -- depending on which
54 # network you are using, the weights can be 90-575MB, so be
55 # patient; the weights will be cached and subsequent runs of this
56 # script will be *much* faster)
57 print("[INFO] loading {}".format(args["model"]))
58 Network = MODELS[args["model"]]
59 model = Network(weights="imagenet")
```

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Pre-trained CNNs in Keras

imagenet_pretrained.py

```
61 # load the input image using the Keras helper utility while ensuring
62 # the image is resized to 'inputShape', the required input dimensions
63 # for the ImageNet pre-trained network
64 print("[INFO] loading and pre-processing image...")
65 image = load_img(args["image"], target_size=inputShape)
66 image = img_to_array(image)
67
68 # our input image is now represented as a NumPy array of shape
69 # (inputShape[0], inputShape[1], 3) however we need to expand the
70 # dimension by making the shape (1, inputShape[0], inputShape[1], 3)
71 # so we can pass it through thenetwork
72 image = np.expand_dims(image, axis=0)
73
74 # pre-process the image using the appropriate function based on the
75 # model that has been loaded (i.e., mean subtraction, scaling, etc.)
76 image = preprocess(image)
```

Images are trained/classified in batches with these CNNs so we need to add an extra dimension; failure to do so will cause an error

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Pre-trained CNNs in Keras

imagenet_pretrained.py

```
78 # classify the image
79 print("[INFO] classifying image with '{}...'".format(args["model"]))
80 preds = model.predict(image)
81 P = imagenet_utils.decode_predictions(preds)
82
83 # loop over the predictions and display the rank-5 predictions +
84 # probabilities to our terminal
85 for (i, (imagenetID, label, prob)) in enumerate(P[0]):
86     print("{} . {}: {:.2f}%".format(i + 1, label, prob * 100))
```

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Pre-trained CNNs in Keras

imagenet_pretrained.py

```
88 # load the image via OpenCV, draw the top prediction on the image,  
89 # and display the image to our screen  
90 orig = cv2.imread(args["image"])  
91 (imagenetID, label, prob) = P[0][0]  
92 cv2.putText(orig, "Label: {}".format(label), (10, 30),  
93           cv2.FONT_HERSHEY_SIMPLEX, 0.8, (0, 255, 0), 2)  
94 cv2.imshow("Classification", orig)  
95 cv2.waitKey(0)
```

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Pre-trained CNNs in Keras

`imagenet_pretrained.py`

```
$ python imagenet_pretrained.py \
    --image example_images/example_01.jpg --model vgg16
$ python imagenet_pretrained.py \
    --image example_images/example_02.jpg --model vgg19
$ python imagenet_pretrained.py \
    --image example_images/example_03.jpg --model inception
$ python imagenet_pretrained.py \
    --image example_images/example_04.jpg --model xception
$ python imagenet_pretrained.py \
    --image example_images/example_05.jpg --model resnet
```

Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

Pre-trained CNNs in Keras

`imagenet_pretrained.py`



Credit: Adrian Rosebrock, *Deep Learning for Computer Vision*, PyImageSearch, 2017

