

A BEGINNER'S GUIDE TO DEEP LEARNING

Irene Chen

@irenetrampoline

PyCon 2016

“A beginner’s guide to
deep learning”

“A **beginner’s** guide to
deep learning”

Convolutional nets

Backpropagation

Image recognition

Restricted Boltzmann machines

DeepMind's AlphaGo beating
professional Go player Lee Sedol
Nvidia and its latest GPU architecture
Toyota's \$1 billion AI investment
Facebook is building AI that builds AI

Geoff Hinton

Yann LeCun

Andrew Ng

Yoshua Bengio

deep learning

All

News

Videos

Books

Images

About 13,100,000 results (0.48 seconds)

The standard type of RBM has binary-valued ([Boolean/Bernoulli](#)) hidden and visible units, and consists of a [matrix](#) of weights $W = (w_{i,j})$ (size $m \times n$) associated with the connection between hidden unit h_j and visible unit v_i , as well as bias weights (offsets) a_i for the visible units and b_j for the hidden units. Given these, the *energy* of a configuration (pair of boolean vectors) (v, h) is defined as

$$E(v, h) = - \sum_i a_i v_i - \sum_j b_j h_j - \sum_i \sum_j v_i w_{i,j} h_j$$

or, in matrix notation,

$$E(v, h) = -a^T v - b^T h - v^T W h$$

This energy function is analogous to that of a [Hopfield network](#). As in general Boltzmann machines, probability distributions over hidden and/or visible vectors are defined in terms of the energy function:^[9]

$$P(v, h) = \frac{1}{Z} e^{-E(v, h)}$$

where Z is a [partition function](#) defined as the sum of $e^{-E(v, h)}$ over all possible configurations (in other words, just a [normalizing constant](#) to ensure the probability distribution sums to 1). Similarly, the ([marginal](#)) probability of a visible (input) vector of booleans is the sum over all possible hidden layer configurations:^[9]

$$P(v) = \frac{1}{Z} \sum_h e^{-E(v, h)}$$

```

1 import theano
2 from theano import tensor as T
3 from theano.sandbox.rng_mrg import MRG_RandomStreams as RandomStreams
4 import numpy as np
5 from load import mnist
6
7 srng = RandomStreams()
8
9 def floatX(X):
10     return np.asarray(X, dtype=theano.config.floatX)
11
12 def init_weights(shape):
13     return theano.shared(floatX(np.random.randn(*shape) * 0.01))
14
15 def rectify(X):
16     return T.maximum(X, 0.)
17
18 def softmax(X):
19     e_x = T.exp(X - X.max(axis=1).dimshuffle(0, 'x'))
20     return e_x / e_x.sum(axis=1).dimshuffle(0, 'x')
21
22 def RMSprop(cost, params, lr=0.001, rho=0.9, epsilon=1e-6):
23     grads = T.grad(cost=cost, wrt=params)
24     updates = []
25     for p, g in zip(params, grads):
26         acc = theano.shared(p.get_value() * 0.)
27         acc_new = rho * acc + (1 - rho) * g ** 2
28         gradient_scaling = T.sqrt(acc_new + epsilon)
29         g = g / gradient_scaling
30         updates.append((acc, acc_new))
31         updates.append((p, p - lr * g))
32     return updates
33
34 def dropout(X, p=0.):
35     if p > 0:
36         retain_prob = 1 - p
37         X *= srng.binomial(X.shape, p=retain_prob, dtype=theano.config.floatX)
38         X /= retain_prob
39     return X

```


Too much math


Too much code

Today

- Why now?
- Neural Networks in 7 minutes
- Deep nets in Caffe

WHY NOW?







Engine (neural network)



Engine (neural
network)

Fuel
(data)

Classifier

Input → **Classifier** → **Output**



```
graph LR; Input --> Classifier; Classifier --> Output;
```




Classifier



Ripe?



pe?



Ripe?

Trained Classifier

Logistic regression

Naïve Bayes

Support vector machine


K-nearest neighbors

Random forests



Ripe?

Trained Classifier



Lesson 1: Why now? Big data, big processing power, robust neural networks

NEURAL NETWORKS IN 7








Photo: Rebecca-Lee (Flickr)










Input Nodes




Output Nodes




Input Nodes




Output Nodes



Input Nodes




Input Nodes



Output Nodes

Input Nodes




A

B


C

Output Nodes




Sigmoid

$$f(x) = \frac{1}{1 + e^{-\beta x}}$$



Input Nodes




A

B


C

D

Output Nodes



Input Nodes




Output Nodes


Input Nodes




Hidden
layers



Output
Nodes





Input Nodes





Output Nodes


Input Nodes



Output Nodes




Input Nodes




Output Nodes

Input Nodes




Output Nodes

Input Nodes




Output Nodes

Input Nodes



Output Nodes



Input Nodes





Output Nodes

Forward propagation


Input Nodes




Output Nodes




Input Nodes



Output Nodes




Input Nodes



Output Nodes


Input Nodes



Output Nodes

No randomness!


Input Nodes





Output Nodes





Backpropagation
























1. $\frac{\partial E}{\partial w_{ij}} = \delta_j * x_i$ for all weights and biases
2. $o'_j = o_j * (1 - o_j)$ for output layer nodes using softmax
3. $\varphi'_j = (1 - \varphi_j) * (1 + \varphi_j)$ for hidden layer nodes using tanh
4. $\varphi'_j = \varphi_j * (1 - \varphi_j)$ for hidden layer nodes using logistic sigmoid
5. $e_j = (o_j - t_j)$ for hidden and output layer nodes
6. $\delta_j = e_j * o'_j$ if j is an output node
7. $\delta_j = (\sum \delta_j w_j) * \varphi'_j$ if j is a hidden node
8. $\Delta w_{ij} = \alpha * \frac{\partial E}{\partial w_{ij}}$ delta for all weights and biases
9. $w'_{ij} = w_{ij} + \Delta w_{ij}$ update for all weights and biases


Values of the nodes


Amount of error


Weights of edges

Learning rate
















Tuning parameters

Input Nodes



Output Nodes



Lesson 2: Neural networks
can be trained on labeled
data to classify avocados

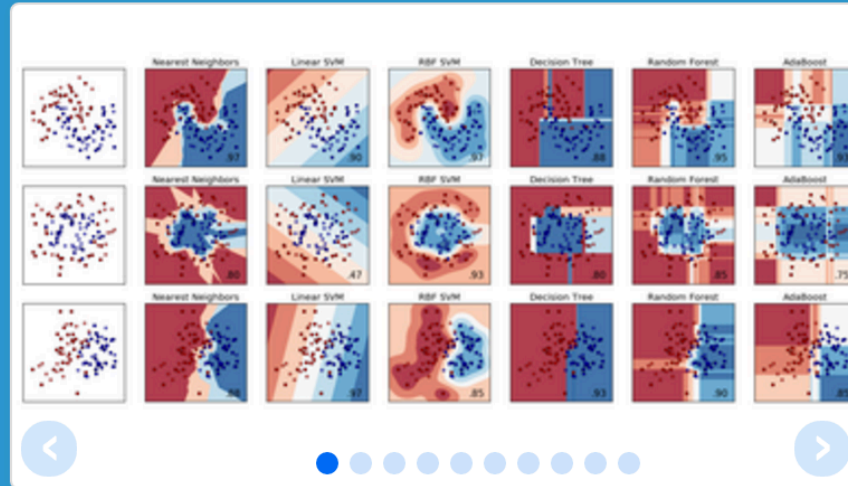
DEEP NETS ON CAFFE

Scikit-learn

Caffe

Theano

iPython Notebook



scikit-learn

Machine Learning in Python

- Simple and efficient tools for data mining and data analysis
- Accessible to everybody, and reusable in various contexts
- Built on NumPy, SciPy, and matplotlib
- Open source, commercially usable - BSD license

Classification

Identifying to which category an object belongs to.

Applications: Spam detection, Image recognition.

Algorithms: [SVM](#), [nearest neighbors](#), [random forest](#), ... — [Examples](#)

Regression

Predicting a continuous-valued attribute associated with an object.

Applications: Drug response, Stock prices.

Algorithms: [SVR](#), [ridge regression](#), [Lasso](#), ... — [Examples](#)

Clustering

Automatic grouping of data sets.

Applications: Customer segmentation, Grouping experiments

Algorithms: [k-Means](#), [mean-shift](#), ...

Dimensionality reduction

Reducing the number of random variables to

Model selection

Comparing, validating and choosing

Preprocessing

Feature extraction and

Caffe

Deep learning framework
by the [BVLC](#)

Created by
[Yangqing Jia](#)
Lead Developer
[Evan Shelhamer](#)

 [View On GitHub](#)

Caffe

Caffe is a deep learning framework made with expression, speed, and modularity in mind. It is developed by the Berkeley Vision and Learning Center ([BVLC](#)) and by community contributors. [Yangqing Jia](#) created the project during his PhD at UC Berkeley. Caffe is released under the [BSD 2-Clause license](#).

Check out our web image classification [demo](#)!

Why Caffe?

Expressive architecture encourages application and innovation. Models and optimization are defined by configuration without hard-coding. Switch between CPU and GPU by setting a single flag to train on a GPU machine then deploy to commodity clusters or mobile devices.

Extensible code fosters active development. In Caffe's first year, it has been forked by over 1,000 developers and had many significant changes contributed back. Thanks to these contributors the framework tracks the state-of-the-art in both code and models.

Speed makes Caffe perfect for research experiments and industry deployment. Caffe can process **over 60M images per day** with a single NVIDIA K40 GPU*. That's 1 ms/image for inference and 4 ms/image for learning. We believe that Caffe is the fastest convnet implementation available.

Community: Caffe already powers academic research projects, startup prototypes, and even large-scale industrial applications in vision, speech, and multimedia. Join our community of brewers on

Welcome

Theano is a Python library that allows you to define, optimize, and evaluate mathematical expressions involving multi-dimensional arrays efficiently. Theano features:

- **tight integration with NumPy** – Use `numpy.ndarray` in Theano-compiled functions.
- **transparent use of a GPU** – Perform data-intensive calculations up to 140x faster than with CPU.(float32 only)
- **efficient symbolic differentiation** – Theano does your derivatives for function with one or many inputs.
- **speed and stability optimizations** – Get the right answer for `log(1+x)` even when `x` is really tiny.
- **dynamic C code generation** – Evaluate expressions faster.
- **extensive unit-testing and self-verification** – Detect and diagnose many types of errors.

Theano has been powering large-scale computationally intensive scientific investigations since 2007. But it is also approachable enough to be used in the classroom (University of Montreal's [deep learning/machine learning](#) classes).

News

- 2016/05/09: New technical report on Theano: [Theano: A Python framework for fast computation of mathematical expressions](#). This is the new preferred reference.
- 2016/04/21: Release of Theano 0.8.2, adding support for [CuDNN v5](#).
- 2016/03/29: Release of Theano 0.8.1, fixing a compilation issue on MacOS X with XCode 7.3.
- 2016/03/21: Release of Theano 0.8. Everybody is encouraged to update.
- Multi-GPU.
- We added support for `CNMeM` to speed up the GPU memory allocation.
- Theano 0.7 was released 26th March 2015. Everybody is encouraged to update.
- We support [cuDNN](#) if it is installed by the user.
- Open Machine Learning Workshop 2014 [presentation](#).
- Colin Raffel [tutorial on Theano](#).
- Ian Goodfellow did a [12h class with exercises on Theano](#).
- New technical report on Theano: [Theano: new features and speed improvements](#).

theano

Table Of Contents

[Welcome](#)[News](#)[Download](#)[Status](#)[Citing Theano](#)[Documentation](#)[Community](#)[Help!](#)

- [How to Seek Help](#)
- [How to provide help](#)

Next topic

Release Notes

This Page

Show Source

Quick search

Go

Enter search terms or a module, class or function name.

IP[y]: IPython

Interactive Computing

[Install](#) · [Documentation](#) · [Project](#) · [Jupyter](#) · [News](#) · [Cite](#) · [Donate](#)

The Jupyter Notebook

(Formerly known as the IPython Notebook)

The IPython Notebook is now known as the Jupyter Notebook. It is an interactive computational environment, in which you can combine code execution, rich text, mathematics, plots and rich media. For more details on the Jupyter Notebook, please see the [Jupyter](#) website.

Google™ Custom Search

NOTEBOOK VIEWER

Share your notebooks



COMMUNITY

[Stack Overflow](#)

[Mailing list](#)

[File a bug](#)

[Reddit](#)

[Twitter](#)

FOR DEVELOPERS

[Mailing list](#)

Scikit-learn

Caffe

Theano

iPython Notebook

Loading a pre-trained
network into Caffe



Large Scale Visual
Recognition Challenge 2010
(ILSVRC 2010)

10 million images
10,000 object classes
310,000 iterations

Branch: master ▾

caffe / **examples** / 00-classification.ipynb

Find file

Copy path

 **longjon** [example] improve classification notebook

bd6b03f Feb 24, 2016

3 contributors



781 lines (780 sloc) | 794 KB

Raw

Blame

History



Classification: Instant Recognition with Caffe

In this example we'll classify an image with the bundled CaffeNet model (which is based on the network architecture of Krizhevsky et al. for ImageNet).

We'll compare CPU and GPU modes and then dig into the model to inspect features and the output.

1. Setup

- First, set up Python, numpy, and matplotlib.

```
In [1]: # set up Python environment: numpy for numerical routines, and matplotlib for plotting
import numpy as np
import matplotlib.pyplot as plt
```

2. Load net and set up input preprocessing

- Set Caffe to CPU mode and load the net from disk.

```
In [4]: caffe.set_mode_cpu()


model_def = caffe_root + 'models/bvlc_reference_caffenet/deploy.prototxt'
model_weights = caffe_root + 'models/bvlc_reference_caffenet/bvlc_reference_caffenet.caffemodel'

net = caffe.Net(model_def,          # defines the structure of the model
                 model_weights,     # contains the trained weights
                 caffe.TEST)        # use test mode (e.g., don't perform dropout)
```

- Set up input preprocessing. (We'll use Caffe's `caffe.io.Transformer` to do this, but this step is independent of other parts of Caffe, so any custom preprocessing code may be used).

Our default CaffeNet is configured to take images in BGR format. Values are expected to start in the range [0, 255] and then have the mean ImageNet pixel value subtracted from them. In addition, the channel dimension is expected as the first (*outermost*) dimension.

As matplotlib will load images with values in the range [0, 1] in RGB format with the channel as the *innermost* dimension, we are arranging for the needed transformations here.



```
In [9]: # load ImageNet labels
labels_file = caffe_root + 'data/ilsvrc12/synset_words.txt'
if not os.path.exists(labels_file):
    !../data/ilsvrc12/get_ilsvrc_aux.sh

labels = np.loadtxt(labels_file, str, delimiter='\t')

print 'output label:', labels[output_prob.argmax()]

output label: n0212304 tabby, tabby cat
```

Tabby cat

- "Tabby cat" is correct! But let's also look at other top (but less confident predictions).

```
In [10]: # sort top five predictions from softmax output
top_inds = output_prob.argsort()[::-1][:5] # reverse sort and take five largest items

print 'probabilities and labels:'
zip(output_prob[top_inds], labels[top_inds])
```

Tabby cat

Tiger cat

Egyptian cat

Red fox

Lynx

```
Out[10]: probabilities and labels:
[(0.31243637, 'n0212304 5 tabby, tabby cat'),
 (0.2379719, 'n0212315 1 tiger cat'),
 (0.12387239, 'n0212405 5 Egyptian cat'),
 (0.10075711, 'n0211902 2 red fox, Vulpes vulpes'),
 (0.070957087, 'n0212752 52 lynx, catamount')]
```


Lesson 3: Caffe provides pre-trained networks to jumpstart learning

Today

- **Lesson 1:** Why now? Big data, big processing power, robust neural networks
- **Lesson 2:** Neural networks can be trained on labeled data to classify avocados
- **Lesson 3:** Caffe provides pre-trained networks to jumpstart learning

What do **you** go from here?

Today

- Lesson 1: Why now? Big data, big processing power, robust neural networks
- Lesson 2: Neural networks can be trained on labeled data to classify avocados
- Lesson 3: Caffe provides pre-trained networks to jumpstart learning

Cuda implementations
Theano, Tensorflow, etc

Today

- Lesson 1: Why now? Big data, big processing power, robust neural networks
- Lesson 2: Neural networks can be trained on labeled data to classify avocados
- Lesson 3: Caffe provides pre-trained networks to jumpstart learning

Restricted Boltzmann Machines


Recurrent network

Convolutional network

Today

- Lesson 1: Why now? Big data, big processing power, robust neural networks
- Lesson 2: Neural networks can be trained on labeled data to classify avocados
- Lesson 3: Caffe provides pre-trained networks to jumpstart learning

Caffe iPython notebooks
Kaggle competitions



Thank you!

irenetrampoline@gmail.com