

A PRACTICAL INTRODUCTION TO DEEP LEARNING FOR THE EARTH SYSTEM SCIENCES, USING PYTORCH

David Hall, Senior Data Scientist, NVIDIA NOAA-AI Workshop Tutorial, Nov 2020



WE WILL

Identify tropical storms and hurricanes with deep learning using PyTorch



WE WILL Examine Artic Sea-Ice Extent as a 'warm-up' exercise







WE CAN

Apply the same techniques to detect features on the sun

(AIA 193Å) BCE loss = 0.00137



LEARNING GOALS



In 2 hours, we can only travel so far

Main goal: Become familiar with ideas and process

A starting point for solving your own problems

SUGGESTED PRE-REQUISTITES Familiarity with Python and NumPy are helpful







WHY PYTORCH?

PyTorch is the Preferred Framework for Research



Repository Creation Date

https://paperswithcode.com/trends

🕺 NVIDIA.



AGENDA

- Intro
- Activity: Intro to Colab
- What is Deep Learning?
- Activity : Curve Fitting in PyTorch
- Fully Connected Networks
- Break
- Classification and CNNs
- Activity : Cyclone Classification
- Wrap up

5 min 15 min 15 min 15 min

20 min 5 min

20 min 20 min 5 min

DOWNLOAD THIS PRESENTATION https://github.com/halldm2000/NOAA-AI-2020-TUTORIAL https://github.com/halldm2000/NOAA-AI-2020-TU ດ Why GitHub? ~ Team Enterprise Explore V Marketplace Pricing halldm2000 / NOAA-AI-2020-TUTORIAL NOAA-AI-2020-TUTORIAL

A Practical Introduction to Deep Learning for the Earth System Sciences, using PyTorch

1

In this tutorial, we will learn how to build deep learning applications from the ground up using PyTorch. We will begin simply and build toward full-fledged solutions for detecting tropical cyclones and other strong storms in model data and satellite observations. The primary goal of this tutorial is to familiarize you with each of the concepts and tools needed to begin building your own deep learning applications. Previous experience with Python and Numpy are desired, but not required.

Click here to download the tutorial PDF

TORIAL	Ċ								
Search	7 Sign in Sign up								
	⊙ Watch 1 ☆ Star 0 ♀ Fork 0								

HAVING TROUBLE WITH GOOGLE COLAB? https://www.google.com/chrome/



Google Colab works best with Google Chrome Browser



ACTIVITY: INTRO TO COLAB Click here to launch the notebook





11

🕺 NVIDIA.





DEEP LEARNING ANALOGIES What is this deep learning thing, anyway?

A GENERALIZATION OF CURVE FITTING (SIMPLE!) A NEW TYPE OF SOFTWARE (POWERFUL!)



DEEP LEARNING CAN DO IMPRESSIVE THINGS











OPERATE VEHICLES AUTONOMOUSLY

GENERATE ORIGINAL CONTENT

BUILD FUNCTIONS FROM EXAMPLES

Find *f*, given *x* and *y*



SUPERVISED DEEP LEARNING INPUTS x_1 x_3 x_4 $\left(x_{5}\right)$ x_6

OUTPUTS



A NEW WAY TO BUILD SOFTWARE



LEARNED FUNCTION

|--|

- A = relu(w1 * [T,P,Q] + b1)
- B = relu(w2 * A + b2)
- C = relu(w3 * B + b3)
- D = relu(w4 * C + b4)
- E = relu(w5 * D + b5)
- y = sigmoid(w6 * E + b6)

return y

Reverse-engineer a function from inputs / outputs

LEARNED FUNCTIONS ARE GPU ACCELERATED Next level software. No porting required.



GPU ACCELERATED FUNCTIONS







USE IT TO ENHANCE EXISTING APPLICATIONS Improve all stages of numerical weather prediction





USE IT TO BUILD NEW CAPABILITIES



WEATHER DETECTION



ENVIRONMENTAL MONITORING



DISASTER PLANNING



DATA FUSION



AUTONOMOUS SENSORS



DATA ENHANCEMENT



NEAR-EARTH OBJECT DETECTION



DATA DRIVEN MODELS



PART 1: CURVE FITTING WITH PYTORCH



SEA ICE EXTENT VS TIME An interesting data set to practice curve-fitting



22

🕺 NVIDIA.

ANNUAL SEA ICE EXTENT 2D Find function: extent = f(day of year, year)



23

🕺 NVIDIA.

ANNUAL SEA ICE EXTENT 1D Find function: extent = f(day of the year)

Annual Artic Sea-ice Extent.



NVIDIA

TAYLOR SERIES

Function approximation using polynomial basis vectors

 $y = b_0 + w_0 x^1 + w_1 x^2 + w_2 x^3$



25 🕺 NVIDIA.

FOURIER SERIES

Function approximation using sinusoidal basis vectors

 $y = b_0 + w_0 \sin(x) + w_1 \sin(2x) + w_2 \sin(3x)$



26 INVIDIA

CURVE FIT



CURVE FIT

1	import torch			Annual Arti
2	<pre>torch.manual_seed(0)</pre>			
3		. [0000
4	# DATA	16 1		1000
5	x,y = load data()		1	1
6			1	
7	# MODEL	14 -		
8	N = 3			<i>R</i>
9	<pre>x0 = torch.zeros(1, requires_grad=True)</pre>		and the second	
10	<pre>w = torch.zeros(N+1, requires_grad=True)</pre>	12 -		
11	<pre>def model(x): return sum(w[i]*(x-x0)**i for i in range(N+1))</pre>			
12		Ę		
13	# CONFIGURE	sq.		
14	<pre>optimizer = torch.optim.AdamW(params = [x0,w], lr=2e-3)</pre>	5 ¹⁰		
15	<pre>loss_fcn = torch.nn.MSELoss()</pre>	illi		
16		ع		
17	for epoch in range(2000+1):	8 -		
18				
19	# TRAIN			
20	prediction = model(x)	6 -		
21	optimizer.zero_grad()	Ĩ		
22	loss = loss_fcn(prediction, y)			
23	loss.backward()			
24	optimizer.step()	4 -		
25		L	0	50
26	<pre>with torch.no_grad(): plot(interval=100)</pre>		~	50



Sea-ice Extent Epoch=1990 Training loss=0.096 $y = -1.20X^{0} - 0.98X^{1} + 1.51X^{2} + 0.72X^{3}$ X = (x - 0.488)

ACTIVITY: IMPROVE THE CURVE FIT Click here to open notebook

import torch, numpy as np torch.manual see Train the Model Below, we have approximated the seasonal sea-ice extent data using a 3rd order Taylor series. Can we do better? DATA x,y = load data(**TODO:** MODEL • Try changing the order of the model to N=2, N=5, etc 3 = Add more epochs if neccessary to get a better fit x0 = torch.zeros = torch.zeros • Adjust the learning rate (Ir=...) def model(x): re Try replacing MSELoss with L1Loss Increase the plotting interval to reduce training time # CONFIGURE optimizer = torc • See how small you can get the training loss loss_fcn = torc • Try replacing the Taylor series with a Fourier series for epoch in ran import torch # TRAIN torch.manual_seed(0) prediction = π 3 optimizer.zero Δ # DATA loss = loss fo x,y = load_data() 5 loss.backward(6 optimizer.step()

26 with torch.no grad(): plot(interval=100)

2

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

#



50





FOURIER SERIES SOLUTION



LOSS FUNCTION

Error surface, measuring how well prediction matches targets



31

🤒 NVIDIA

OPTIMIZER

Strategy for adjusting model weights, to minimize loss



32 <

🥯 NVIDIA.

BACKPROPAGATION Compute gradient of loss function with respect to each parameter





Compute Loss

33 🕺 NVIDIA.

AUTOGRAD

Let a framework keep track of your gradient, so you don't have to

PyTorch Autograd

```
from torch.autograd import Variable
```

```
x = Variable(torch.randn(1, 10))
prev h = Variable(torch.randn(1, 20))
W h = Variable(torch.randn(20, 20))
W x = Variable(torch.randn(20, 10))
```

```
i2h = torch.mm(W x, x.t())
h2h = torch.mm(W_h, prev_h.t())
next h = i2h + h2h
next h = next h.tanh()
```

next_h.backward(torch.ones(1, 20))





UNDERFITTING AND OVERFITTING A good model is one that generalizes to new data



https://scikit-learn.org/stable/auto_examples/model_selection/plot_underfitting_overfitting.html

🤒 NVIDIA
EPISTEMIC ERROR

Error can be large, far from training data



EPISTEMIC ERROR

Error grows, in regions far from our training data



https://www.researchgate.net/figure/Illustration-of-Gaussian-process-regression-inone-dimension-for-the-target-test_fig1_327613136

DATA SPLITTING

Split your examples into train, validation, and test sets



For model training

For final evaluation For hyperparameter tuning



SPLITTING TIME SERIES DATA

Avoid correlations and extrapolation



40

🐵 NVIDIA.

```
import torch, numpy as np
torch.manual seed(0)
```

DATA

```
x,y = load data()
```

```
shuffle blocks(blocksize=200)
ntrain, nval, ntest = (len(x)*torch.tensor([0.50, 0.25, 0.25])).int()
x_train, x_val, xtest = x[:ntrain], x[ntrain:-ntest], x[-ntest:]
y_train, y_val, ytest = y[:ntrain], y[ntrain:-ntest], y[-ntest:]
```

MODEL

```
Ν
      = 5
      = torch.randn(N, requires grad=True)
\mathbf{x0}
     = torch.randn(N, requires grad=True)
freq
      = torch.zeros(N, requires_grad=True)
ω
def model(x):
 return sum(w[i]*torch.sin(freq[i]*np.pi*(x-x0[i])) for i in range(N) )
```

CONFIGURE

```
optimizer = torch.optim.AdamW(params = [freq,x0,w], lr=5e-3)
loss_fcn = torch.nn.MSELoss()
```

```
train loss, val loss=[],[]
for epoch in range(200+1):
```

TRAIN

```
prediction = model(x train)
optimizer.zero grad()
loss = loss fcn(prediction, y train)
loss.backward()
optimizer.step()
train_loss.append(loss.item())
```

VALIDATE

```
with torch.no_grad():
```

```
p val = model(x val)
loss_val = loss_fcn(p_val, y_val)
val_loss.append(loss_val.item())
```





VALIDATION **Block shuffled**

```
import torch, numpy as np
     torch.manual seed(0)
 3
     # DATA
    x,y = load data()
     shuffle blocks(blocksize=1)
     ntrain, nval, ntest = (len(x) * torch.tensor([0.50, 0.25, 0.25])).int()
     x_train, x_val, xtest = x[:ntrain], x[ntrain:-ntest], x[-ntest:]
 8
     y train, y val, ytest = y[:ntrain], y[ntrain:-ntest], y[-ntest:]
 9
                                                                                             16
10
11
     # MODEL
                                                                                             14
12
     N
           = 5
                                                                                         million sq km
10
           = torch.randn(N, requires grad=True)
13
     \mathbf{x}\mathbf{0}
          = torch.randn(N, requires grad=True)
14
     freq
           = torch.zeros(N, requires grad=True)
15
     w
     def model(x):
16
       return sum(w[i]*torch.sin(freq[i]*np.pi*(x-x0[i])) for i in range(N) )
17
18
                                                                                              6
19
     # CONFIGURE
                                                                                                    training data
     optimizer = torch.optim.AdamW(params = [freq,x0,w], lr=5e-3)
20
                                                                                                    validation data
     loss_fcn = torch.nn.MSELoss()
21
                                                                                                           50
                                                                                                  0
22
23
     train_loss, val_loss=[],[]
                                                                                            100
     for epoch in range(200+1):
24
25
       # TRAIN
26
       prediction = model(x_train)
27
                                                                                         average error
       optimizer.zero grad()
28
       loss = loss fcn(prediction, y train)
29
                                                                                           10-1
       loss.backward()
30
31
       optimizer.step()
                                                                                                    Too good to be true:
       train loss.append(loss.item())
32
33
       # VALIDATE
34
       with torch.no grad():
35
                                                                                           10^{-2}
36
                                                                                                  0
                                                                                                           25
         p val = model(x val)
37
         loss val = loss fcn(p val, y val)
38
         val loss.append(loss val.item())
39
40
         plot(interval=10)
```

VALIDATION Fully shuffled: self-correlated

Annual Artic Sea-ice Extent Epoch=200 Training loss=0.057 $y = -0.17X^{0} + 0.34X^{1} + 0.44X^{2} - 0.41X^{3} - 0.49X^{4}$ X = (x - -0.305)



```
import torch, numpy as np
     torch.manual_seed(0)
     # DATA
     x,y = load data()
     ntrain, nval, ntest = (len(x) * torch.tensor([0.50, 0.25, 0.25])).int()
     x_train, x_val, xtest = x[:ntrain], x[ntrain:-ntest], x[-ntest:]
 8
     y_train, y_val, ytest = y[:ntrain], y[ntrain:-ntest], y[-ntest:]
 9
                                                                                              16
10
11
     #
       MODEL
                                                                                              14
     Ν
            = 5
12
           = torch.randn(N, requires grad=True)
                                                                                             12
13
                                                                                           million sq km
     \mathbf{x}\mathbf{0}
     freq = torch.randn(N, requires grad=True)
14
                                                                                              10
            = torch.zeros(N, requires grad=True)
15
     w
     def model(x):
16
                                                                                              8
       return sum(w[i]*torch.sin(freq[i]*np.pi*(x-x0[i])) for i in range(N) )
17
18
                                                                                              6
19
     # CONFIGURE
                                                                                                     training data
     optimizer = torch.optim.AdamW(params = [freq,x0,w], lr=5e-3)
                                                                                                     validation data
20
21
     loss fcn = torch.nn.MSELoss()
                                                                                                            50
                                                                                                   0
22
     train_loss, val_loss=[],[]
23
                                                                                             10
24
     for epoch in range(200+1):
25
26
       # TRAIN
27
       prediction = model(x train)
                                                                                          average error
       optimizer.zero_grad()
28
       loss = loss_fcn(prediction, y_train)
29
                                                                                           10<sup>-1</sup>
       loss.backward()
30
       optimizer.step()
31
32
       train_loss.append(loss.item())
33
       # VALIDATE
34
       with torch.no_grad():
35
                                                                                            10^{-2}
36
                                                                                                            25
                                                                                                   0
         p val = model(x val)
37
         loss val = loss fcn(p val, y val)
38
         val_loss.append(loss_val.item())
39
```

plot(interval=10)

40

VALIDATION Sequential: future = extrapolation

Annual Artic Sea-ice Extent Epoch=200 Training loss=0.016 $y = +0.22X^{0} - 0.40X^{1} - 0.40X^{2} - 0.45X^{3} + 0.39X^{4}$ X = (x - 1.881)





FULLY CONNECTED NETWORKS

```
import torch, numpy as np
     torch.manual seed(0)
 3
     # DATA
     x,y = load data()
     shuffle blocks(blocksize=100)
     ntrain, nval, ntest = (len(x)*torch.tensor([0.50, 0.25, 0.25])).int()
    x_train, x_val, xtest = x[:ntrain], x[ntrain:-ntest], x[-ntest:]
    y_train, y_val, ytest = y[:ntrain], y[ntrain:-ntest], y[-ntest:]
 9
                                                                                             16
10
     # MODEL
11
                                                                                             14
     def relu(x): return x*(x>0)+ 0.1*x*(x<0)</pre>
12
     N = 20
13
                                                                                             12
                                                                                           million sq km
     b = torch.randn(N+1, requires_grad=True)
14
                                                                                              10
     w1= torch.randn(N,
                           requires grad=True)
15
     w2= torch.randn(N,
                           requires grad=True)
16
     def model(x): return sum(w2[i]*relu(w1[i]*x+b[i]) for i in range(N))+b[N]
17
18
                                                                                              6
     # CONFIGURE
19
     optimizer = torch.optim.AdamW(params = [b, w1, w2], lr=1e-2)
20
     loss_fcn = torch.nn.MSELoss()
21
                                                                                                   0
22
     train_loss, val_loss=[],[]
23
     for epoch in range(1000+1):
24
25
       # TRAIN
26
                                                                                            10<sup>0</sup>
       prediction = model(x train)
27
                                                                                          average error
10<sup>-1</sup>
       optimizer.zero_grad()
28
       loss = loss_fcn(prediction, y_train)
29
       loss.backward()
30
       optimizer.step()
31
       train_loss.append(loss.item())
32
33
       # VALIDATE
34
       with torch.no_grad():
35
                                                                                            10-2
                                                                                                   0
36
         p val = model(x val)
37
         loss_val = loss_fcn(p_val, y_val)
38
         val_loss.append(loss_val.item())
39
         plot(interval=20)
40
```

SHALLOW NEURAL NET 1D ReLU basis function expansion



```
import torch, torch.nn as nn, numpy as np
     torch.manual_seed(0)
 2
 3
     # DATA
     x,y = load_data()
     shuffle blocks(blocksize=100)
 6
     ntrain, nval, ntest = (len(x) * torch.tensor([0.50, 0.25, 0.25])).int()
     x_train, x_val, xtest = x[:ntrain], x[ntrain:-ntest], x[-ntest:]
 8
                                                                                         16
     y_train, y_val, ytest = y[:ntrain], y[ntrain:-ntest], y[-ntest:]
 9
                                                                                         14
10
     # MODEL
                                                                                         12
                                                                                       million sq km
     N = 20
12
     model = nn.Sequential(nn.Linear(1,N), nn.ReLU(), nn.Linear(N,1))
13
                                                                                          10
14
     # CONFIGURE
15
     optimizer = torch.optim.AdamW(params = model.parameters(), lr=1e-2)
16
                                                                                          6
     loss_fcn = torch.nn.MSELoss()
17
18
     train_loss, val_loss=[],[]
19
                                                                                              0
     for epoch in range(1000+1):
20
21
       # TRAIN
22
       prediction = model(x_train)
23
                                                                                        10<sup>0</sup>
       optimizer.zero_grad()
24
       loss = loss_fcn(prediction, y_train)
25
                                                                                      average error
10<sup>-1</sup>
       loss.backward()
26
27
       optimizer.step()
       train_loss.append(loss.item())
28
29
       # VALIDATE
30
       with torch.no_grad():
31
32
                                                                                        10^{-2}
         p val = model(x val)
                                                                                              0
33
         loss_val = loss_fcn(p_val, y_val)
34
         val_loss.append(loss_val.item())
35
         plot(interval=20)
36
```

SHALLOW NEURAL NET 1D Using nn.sequential

Annual Artic Sea-ice Extent Epoch=1000 Training loss=0.050



RELU = PIECEWISE-LINEAR ReLU basis functions can represent any function using straight line-segments



RELU (X-3) + RELU (X-4)



ACTIVATION FUNCTIONS

Many to choose from. But most use ReLU or LeakyReLU





🐵 NVIDIA

48

RELU IN 2D model = Sequential(Linear(2,N), ReLU(), Linear(N,1))



49

🐵 NVIDIA.

LINEAR, RELU, LINEAR Linear transformations rotate, shift, and scale the activation function





RELU IN 2D Find extent = f(year, day of year). <u>Click here for the code</u>.



🕺 NVIDIA.

INFERENCE

Once training is complete, we can load and use our model

Reconstruct the same model, then load its parameters from file





ARTIFICIAL NEURONS Simple equations with adjustable parameters

Terminal branches of axon (form junctions with other cells)

Biological neuron

Dendrites (receive messages

from other cells)

Axon (passes messages away from the cell body to other neurons, muscles, or glands)

Neural impulse

down the axon)

Cell body (the cell's lifesupport center)

> Myelin sheath (covers the axon of some neurons and helps speed neural impulses) (electrical signal traveling

X₁

https://towardsdatascience.com/the-differences-between-artificial-and-biological-neural-networks-a8b46db828b7

Artificial neuron



 $y = f(w_1 x_1 + w_2 x_2 + w_3 x_3)$

54 🕺 NVIDIA.

FULLY CONNECTED NETWORKS A given neuron is connected to every neuron in the previous layer



LAYER 4 OUTPUT

SINGLE LAYER NEURAL NETWORKS

A series expansion over basis functions ϕ .



🤒 NVIDIA.

TWO LAYER NEURAL NETWORKS Learn the function and the basis functions at the same time



L2: LEARNED BASIS FCNS

L1: RELU BASIS FCNS

DEEPER NEURAL NETWORKS More layers allows for more levels of abstraction

Input

Low-level features

Mid-level features









https://web.eecs.umich.edu/~honglak/icml09-ConvolutionalDeepBeliefNetworks.pdf

High-level features



Output

Probability Image is A Face

🕺 NVIDIA.



PART 2: CLASSIFICATION

GOAL: AUTOMATED STORM DETECTION Automatically detect and classify dangerous tropical cyclones

Epsilon 11 AM ET Oct 19, 2020 **Tropical storm**

https://www.nesdis.noaa.gov/content/goes-east-spies-newly-formed-tropical-storm-epsilon

NOAA: GOES-16



COMPUTER VISION TASKS We will start with classification, as it is simplest.



Instance Segmentation



DOG, DOG, CAT



This image is CC8 public domain

NVIDIA.

CYCLONE CLASSIFICATION Given images of various storms, how can we determine the storm type?



TRADITIONAL APPROACH (like TECA)

- Search for low-pressure center \bullet
- Measure maximum sustained winds \bullet
- Classify storm by location and wind-speed \bullet
- Strengths: simple, interpretable ٠
- Requires wind s unavailable. Lo



DEEP LEARNING

Gather many labelled examples of each storm Let optimizer search for solution

• Strengths: Can use water vapor only • No expert knowledge needed to build the software. Faster development. Less fragile. Works for most any phenomena. GPU accelerated.

many labelled examples.

NVIDIA

IBTRACS Provides expert labels for historical tropical cyclones

ibtracks: tropical storm tracks 2016-2017



https://www.ncdc.noaa.gov/ibtracs/index.php?name=introduction

GFS ANALYSIS

Model data makes a good starting point for this task.



https://www.ncdc.noaa.gov/data-access/model-data/model-datasets/global-forcast-system-gfs

NOAA: GFS ANALYSIS



MAKING A LABELLED DATASET Download the data and convert it to desired format

Download GFS Analysis Data and IBTRACs data from NOAA



Convert GFS data from GRIB format to NetCDF

```
from pathlib import Path
import pygrib
import numpy as np
import os
          = Path(os.environ['HOME'])
home
           = home/"data/original/gfs_analysis_allfields/nomads.ncdc.noaa.gov"
#source
          = home/"data/original/gfs_analysis_allfields/nomads.ncdc.noaa.gov/data/gfsanl/201010"
source
def read_grib(gribfile):
    grbs = pygrib.open(str(gribfile))
   arche cook (A)
```

MAKING A LABELLED DATASET Use IBTRACs to extract labelled examples

For each GFS file, find IBTRACs storms and extract cropped regions

21	<pre># get list of gfs netcdf files</pre>		
22			
23	gfs_files = sorted(list(gfs_dir.glob('**/*.nc')))		
24			
25	<pre># read expert-labelled storm-track database</pre>		
26			
27	<pre>ib = xr.open_dataset(Path(os.environ['HOME'])/"data/original/</pre>		
28	# masses such CEC file		
29	# process each GFS file		
20	for afe file in afe files[1].		
32	IVI gis_litte III gis_littes[.].		
33	<pre>nrint("reading".gfs_file)</pre>		
34	afs = xr.open dataset(afs file)		
35	gio miopon_databot(gio_iiito)		
36	<pre># plot GFS water vapor field as an image (optional)</pre>		
37			
38	<pre>if plot_gfs:</pre>		
39			
40	gfs.pwat.plot(size=10, yincrease= False , cmap='bone')		
41	<pre>fig=plt.gcf(); fig.set_size_inches(20,10)</pre>		
42	plt.show()		
43			
44	# read time t from GFS file		
45	afa time - detetime etantime (afa time 10) (- 0, m, 0, d, 0, 1), 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0,		

btracs/IBTrACS.ALL.v04r00.nc")

NVIDIA

CYCLONE CLASSIFICATION Find function: category = f(water vapor image)







67

AVAILABLE FIELDS

Each NetCDF field contains the following fields

file = crop_64/CAT2/crop64__2010-09-03 00:00:00_EARL_row60_col546.nc



file = crop_64/CAT2/crop64__2010-10-29 03:00:00_CHABA_row96_col222.nc

precipitable water

wind speed









file = crop_64/CAT2/crop64__2010-10-15 21:00:00_MEGI_row106_col241.nc







wind speed





file = crop_64/CAT2/crop64__2010-06-24 06:00:00_CELIA_row141_col487.nc



INVIDIA.

```
from torch.utils.data import DataLoader
    from torch.utils.tensorboard import SummaryWriter
    import torch, torch.nn as nn
 3
    from time import time
    t start = time()
    # DATA
 7
    def get data(): --
 8 >
    train_loader, val_loader = get_data()
17
18
    # MODEL
19
    def get model(): ---
20 >
    model = get model().to(device)
42
    count parameters(model)
43
44
    # CONFIGURE
45
    optimizer = torch.optim.AdamW(model.parameters(), lr=1e-3)
46
    loss fcn = nn.CrossEntropyLoss()
47
    writer
               = SummaryWriter()
48
49
    for epoch in range(30) :
50
51
       # VALIDATE
52
      with torch.no grad():
53
         model.eval()
54
        val loss = running mean()
55
        val_acc = accuracy_metric()
56
57 >
        for i, data in enumerate(val_loader, 0): --
64
       # TRAIN
65
      model.train()
66
      train_acc = accuracy_metric()
67
      train_loss = running_mean()
68
                  = running mean()
69
       gpu
                  = time()
       t0
70
       for i, data in enumerate(train_loader, 0): "
71 >
83
       # display training / validation progress
84
```

Input



CYCLONE CLASSIFICATION Training

http://alexlenail.me/NN-SVG/AlexNet.html

GPU ACCELERATION

GPUs make deep learning practical.



device = torch.device("cuda" if torch.cuda.is_available() else "cpu") model = get_model().to(device) inputs, labels = data[0].to(device), data[1].to(device)





NVIDIA

SUMMARY WRITER

Uses Tensorboard to keep track of training losses

TensorBoard SCALARS		
 Show data download links Ignore outliers in chart scaling 	Q Filter tags (regular expressions so	upported)
method: default -	train tag: Accuracy/train	val tag: Accuracy/val
Smoothing	0.92	0.9
•	0.88	0.85
	0.84	0.8
Horizontal Axis	0.8	0.75
STEP RELATIVE WALL	0.76	0.7
	0.72	0.65
	0.68	
Runs Write a regex to filter runs	0 5 10 15	20 25 0 []
Cct30_17-00-25_964fd419cda	Loss	
TOGGLE ALL RUNS		
runs		

%tensorboard --logdir runs

from torch.utils.tensorboard import SummaryWriter
writer = SummaryWriter()
writer.add_scalar('Accuracy/val', val_acc.value(), epoch)
writer.add_scalar('Accuracy/train', train_acc.value(), epoch)



71

唑 NVIDIA

DATASET AND DATALOADER Classes used to organize and return examples during training

DATASET: Holds a set of examples. (Might read them from files.)



DATALOADER: Returns a batch of examples each time it's called

from torch.utils.data import DataLoader = Dataset(path, inputs, outputs) train data train_loader = DataLoader(train_data, batch_size=200, shuffle=True) for i, data in enumerate(train_loader):


DATA CACHE

Store as many examples as possible in RAM to minimize disc access



cache = Cache(loading_function, device="cuda", memory_cap=80)
if path in cache: return cache[path]
else: load(file, device)

MODEL: CONVOLUTIONAL NEURAL NETWORKS

A SIMPLE CLASSIFIER A basic encoder, using convolution layers

Previous Model: extent = f(year, day of year)

model = nn.Sequential(Linear(2,N), ReLU(), Linear(N,N), ReLU(), Linear(N,1))

Current model: prob of storm = f(water-vapor image)

model = nn.Sequential(Conv2d(Cin,C, 3), ReLU(), MaxPool2d(2) Conv2d(C , C, 3), ReLU(), MaxPool2d(2) Flatten() Linear(C*nrow*ncol//(4*4*4), N), ReLU(), Linear(N, Nout))







SIMPLE CLASSIFIER

Represented as a diagram



http://alexlenail.me/NN-SVG/LeNet.html



CONVOLUTION-2D

A small matrix transformation, applied at each point of the image



77

🕺 NVIDIA.

CONVOLUTION-2D When we have multiple channels, learned kernels are 3d matrices

Filter 1 Input Δ -1 -1 -1 4 x 4 3 x 3 x 3 * Filter 2 6 x 6 x 3 -1 -1 -1 4 x 4 3 x 3 x 3



78 📀 NVIDIA.

CONVOLUTION EXAMPLE: SOBEL FILTER



Image source: https://en.wikipedia.org/wiki/Sobel_operator



 $G = \sqrt{G_x^2 + G_y^2}$



- 79

CONVOLUTION EXAMPLE: SOBEL FILTER



Image source: https://en.wikipedia.org/wiki/Sobel_operator



 $G = \sqrt{G_x^2 + G_y^2}$

0 <

🕺 NVIDIA

MAXPOOL-2D

Pooling reduces the size of the image, so we can detect large-scale features

12	20	30	0	
8	12	2	0	2×2 Max
34	70	37	4	
112	100	25	12	





81

TRANSLATIONAL INVARIANCE Objects in nature look the same from place to place











WHAT ARE CNNS USED FOR?

Problems with translational invariance in variance in time and/or space

CONV_2D



CONV_3D



Computer Vision Invariance in 2d space Computational Physics Invariance in 3d space

FOR? e in time and/or space

CONV_1D



Audio and Time Series Invariance in time

IMBALANCED DATA Some features occur more frequently than others



REGULARIZATION **BatchNorm and Dropout**





https://www.kdnuggets.com/2018/09/dropout-convolutional-networks.html



(b) After applying dropout.

ACTIVITY: CLASSYING TROPICAL CYCLONES Click here to launch the notebook

Hurricane Iselle

System

Train a simple CNN to classify storms

1.1.1

Hawaii

Tropical Depression Genevieve

Genevieve Depression

```
TODO: Some activities to try
     1. Measure training time: CPU
                                            use cache=False. runtime->type = None
       Measure training time: CPU + cache use cache=True. runtime->type = None
       Measure training time: GPU
                                            use_cache=False. runtime->type = GPU
       Measure training time: GPU + cache use_cache=True. runtime->type = GPU
        (Note: when you change runtime type, it deletes your files. So you have to re-download.
 9
10
        (Use nepoch = 2)
11
12
       training time for epoch 0 = ??, ??, ??, ?? seconds
13
       training time for epoch 1 = ??, ??, ??, ?? seconds
14
       training time for 100 epochs = ??, ??, ??, ?? minutes (use: epoch0 + 99 * epoch1)
15
       %GPU utilization
                                   = ??, ??, ??, ??
16
17
       (from now on, always use GPU + cache)
18
    2. How important is the quantity of training data? (use nepochs = 30.)
19
20
       ftrain
                   = 0.1 \quad 0.2 \quad 0.5 \quad 0.9
21
       val accuracy = ?? ?? ?? ??
22
23
    3. How much can we improve the result by including wind and pressure data?
       fields = ["pwat","u","v","mslp"]
24
25
       val_accuracy = ??
26
27
     4. How well can we classify 10 categories, using water vapor alone? fields=['pwat']
28
       Recall, a random guess should get you around 10%
29
       categories = ['TD','TS','CAT1','CAT2','CAT3','CAT4','CAT5','DS','SS','NONE']
30
       val accuracy = ??
31 ''';
```



Tropical Storm Julio



CLASSIFYING TROPICAL STORMS

My results on Google Colab using a T4 GPU

<u>Training time for 100 epochs</u>						
CPU GPU 5000s 3400 (1.5x ~1%GPU -> Problem is I/0 bound)						
+cache +cache 614 (8x) 79s (63x ~55% GPU)						
Using the GPU effectively reduces training time from 1.5 hou						
<u>Accuracy vs data quantity</u>						
f = 90% 50% 20% 10% 0% Acc=.91 .87 .81 .78 .66						
Data quantity and quality are key for high accuracy						
<u>Accuracy vs inputs</u>						
U,v,p pwat pwat,u,v pwat,u,v,p all fields 0.85 0.91 0.93 0.94 ??						
Can detect storms with water vapor alone. But if you have additional information … use it!						

On your own: Augment and balance your samples to increase classification accuracy

urs -> 1 minute !

ADDITIONAL TOPICS TO EXPLORE We've only scratched the surface. There is *much* more to learn.

Learn by doing!

Data augmentation and balancing Regularization, dropout, batch norm, weight decay Skip connections and Resnets Segmentation, detection, instance segmentation Transfer learning Multi-gpu training, data parallel, ddp GPU profiling and tuning using nsight Tensor-RT for runtime optimization Encoder-decoder, autoencoder Anomaly detection Generative models, gans, and vaes Masked convolutions and inpainting Uncertainty quantification

Transformers Self-supervision Pytorch lightning and bolts HPC + AI coupling Deep learning with julia Reinforcement learning Active learning Fast.ai Rapids for ML Etc.

- Explainable / trustworthy AI
- Automatic hyperparameter tuning
- Physics informed neural nets
- Docker, containers, and NGC Jetson nano, edge computing



SOME RECORDED TALKS I'VE GIVEN

<u>Frontiers of Deep Learning (2020)</u>	Recent breakthro
• Deep Learning Architectures (2020)	Overview of majo
• Machine Learning and the Future of the Earth (2020	Can we save the
• NEDTalk: AI and NOAA's Mission (2019)	How AI can be use
• <u>Al for Science (2018/19)</u>	Prototype deep le
• Nonhydrostatic Global Climate Modeling (2015)	Making climate m

- nroughs in AI + geoscience
- ajor deep-learning architectures
- ne planet, using AI?
- used for geosciences
- o learning applications
- e models better





Summary

- Deep learning is powerful
- It's a generalization of curve fitting
- PyTorch is like GPU accelerated numpy
- Most popular DL framework for research
- GPUs are necessary for real-world apps
- We can use DL to build next-level apps
- There's a ton left to learn

Please click here to help us improve the tutorial ! (Two-minute survey)

<u>dhall@nvidia.com</u> <u>david-matthew-hall</u> @halldm2000



ACTIVE REGIONS

HELIOPHYSICS APPLICATIONS

NASA GODDARD ALTAMIRA & NVIDIA

Feature detection can be applied to detect features on the Sun and other astrophysical bodies. In particular, we can apply AI to solar flares and coronal mass ejections in order to predict the influx of highly charged particles on Earth's atmosphere.





SOLAR DYNAMICS OBSERVATORY

- 1.5 TB Data / Day ullet
- **Operational Since 2010**
- AIA: 10 Wavelength Channels ullet
- 150M Images To Be Labelled
- 30k Images Labelled so far
- Coronal Holes
- Active Regions
- Sunspots \bullet
- Solar Flares \bullet
- **Coronal Mass Ejections** \bullet
- Filaments







AIA 4500 Å 6000 Kelvin Photosphere

AIA 1600 Å 10.000 Kelvin Upper photosphere/ Transition region



AIA 211 Å 2 million Kelvin Active regions







AIA 304 Å 50,000 Kelvin Transition region/ Chromosphere



AIA 171 Å 600.000 Kelvin Upper transition Region/quiet corona



2.5 million Kelvin Active regions



AIA 094 Å 6 million Kelvin Flaring regions



AIA 131 Å 10 million Kelvin Flaring regions

RESULTS: CORONAL HOLES

NASA Goddard

Michale Kirk, Barbara Thompson, Jack Ireland, Raphael Attie

NVIDIA David Hall

Altamira Matt Penn, James Stockton,

SOURCE Solar Dynamics Observatory AIA Imager

TARGET Hand-crafted detection algorithm

> Ground Truth **Prob of Detection**



SUNSPOT PREDICTIONS Highly imbalanced dataset. Needs special care.

Predicts all 0s unless special care is taken

- Super-sample minority class
- Under-sample majority class
- Use focal loss

Select small crops from high-res imagery Pos : crops w/large fraction sunspot pixels Neg : randomly selected crops

Train conv net on small crops only Predict on full-resolution images









NVIDIA

RESULTS: SUNSPOTS

NASA Goddard Michale Kirk, Barbara Thompson, Jack Ireland, Raphael Attie

NVIDIA David Hall

Altamira Matt Penn, James Stockton,

SOURCE Solar Dynamics Observatory AIA Imager

TARGET Hand-crafted detection algorithm

Ground Truth Prob of Detection



RESULTS: ACTIVE REGIONS

NASA Goddard

Michale Kirk, Barbara Thompson, Jack Ireland, Raphael Attie

NVIDIA David Hall

Altamira Matt Penn, James Stockton,

SOURCE Solar Dynamics Observatory AIA Imager

TARGET Hand-crafted detection algorithm

Ground Truth Prob of Detection

