


Short developer's tutorial on Deep Learning and Neural Networks with PyTorch


Christian Wolf

December 12th, 2019



Decision making




{dog, cat, bike, avocado, blender, ...}



{0, 1, ... 26, 27, 28..., 98, 99, ...}





{Left, right, ahead, back, pick up, ...}




Motor commands

“A blue parrot with a yellow belly
sitting on a branch in a forest”



Deep Learning @ Tesla


<https://www.youtube.com/watch?v=oBklltKXtDE>



- (3, 960, 1280) input images
- “ResNet-50 like” dilated backbones
- FPN / DeepLabV3 / UNet -like heads
- ~15 tasks => “prototypes framework”


Deep Learning @ Tesla

<https://www.youtube.com/watch?v=oBklltKXtDE>



We would like to learn to predict a value y from observed input x

$$y = h(x, \theta)$$



Handcrafted from domain
knowledge

Learned from data or
interactions



The 3 fundamental problems of ML

1. Expressivity

- What is the complexity of functions my model can represent?

2. Trainability

- How easy is it to fit my model to my training data?

3. Generalization

- Does my model generalize to unseen data?
- In presence of shifts in distributions?

(After Eric Jang & Jascha Sohl-Dickstein)

1 Neural Networks & learning


2 Optimization / gradient descent

3 A simple linear example in PyTorch


4 A multi-layer model in PyTorch

5 Convolutional networks


Deep neural networks




Complex networks



Common activation functions




Hyperbolic tangent
 $\tanh$



Rectified Linear Unit
 ReLU

Universal approximation


We can approximate any $\psi \in \mathcal{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.



Slide: François Fleuret, IDIAP/EPFL
<https://www.idiap.ch/~fleuret/>

We can approximate any $\psi \in \mathcal{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.


$$f(x) = \sigma(w_1 x + b_1)$$



Slide: François Fleuret, IDIAP/EPFL
<https://www.idiap.ch/~fleuret/>

We can approximate any $\psi \in \mathcal{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.


$$f(x) = \sigma(w_1x + b_1) + \sigma(w_2x + b_2)$$



Slide: François Fleuret, IDIAP/EPFL
<https://www.idiap.ch/~fleuret/>

We can approximate any $\psi \in \mathcal{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.


$$f(x) = \sigma(w_1x + b_1) + \sigma(w_2x + b_2) + \sigma(w_3x + b_3)$$



Slide: François Fleuret, IDIAP/EPFL
<https://www.idiap.ch/~fleuret/>

We can approximate any $\psi \in \mathcal{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.


$$f(x) = \sigma(w_1x + b_1) + \sigma(w_2x + b_2) + \sigma(w_3x + b_3) + \dots$$



Slide: François Fleuret, IDIAP/EPFL
<https://www.idiap.ch/~fleuret/>

We can approximate any $\psi \in \mathcal{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.


$$f(x) = \sigma(w_1x + b_1) + \sigma(w_2x + b_2) + \sigma(w_3x + b_3) + \dots$$



Slide: François Fleuret, IDIAP/EPFL
<https://www.idiap.ch/~fleuret/>

We can approximate any $\psi \in \mathcal{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.


$$f(x) = \sigma(w_1x + b_1) + \sigma(w_2x + b_2) + \sigma(w_3x + b_3) + \dots$$



Slide: François Fleuret, IDIAP/EPFL
<https://www.idiap.ch/~fleuret/>

We can approximate any $\psi \in \mathcal{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.


$$f(x) = \sigma(w_1x + b_1) + \sigma(w_2x + b_2) + \sigma(w_3x + b_3) + \dots$$



Slide: François Fleuret, IDIAP/EPFL
<https://www.idiap.ch/~fleuret/>

We can approximate any $\psi \in \mathcal{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.


$$f(x) = \sigma(w_1x + b_1) + \sigma(w_2x + b_2) + \sigma(w_3x + b_3) + \dots$$



Slide: François Fleuret, IDIAP/EPFL
<https://www.idiap.ch/~fleuret/>

We can approximate any $\psi \in \mathcal{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.


$$f(x) = \sigma(w_1x + b_1) + \sigma(w_2x + b_2) + \sigma(w_3x + b_3) + \dots$$



Slide: François Fleuret, IDIAP/EPFL
<https://www.idiap.ch/~fleuret/>

We can approximate any $\psi \in \mathcal{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.


$$f(x) = \sigma(w_1x + b_1) + \sigma(w_2x + b_2) + \sigma(w_3x + b_3) + \dots$$



Slide: François Fleuret, IDIAP/EPFL
<https://www.idiap.ch/~fleuret/>

We can approximate any $\psi \in \mathcal{C}([a, b], \mathbb{R})$ with a linear combination of translated/scaled ReLU functions.


$$f(x) = \sigma(w_1x + b_1) + \sigma(w_2x + b_2) + \sigma(w_3x + b_3) + \dots$$



This is true for other activation functions u

Slide: François Fleuret, IDIAP/EPFL
<https://www.idiap.ch/~fleuret/>

Demo: Tensorflow Playground



<https://playground.tensorflow.org>

1 Neural Networks & learning

2 Optimization / gradient descent

3 A simple linear example in PyTorch

4 A multi-layer model in PyTorch


5 Convolutional networks

Learning by gradient descent

Iterative minimisation through **gradient descent**:

$$\theta^{[t+1]} = \theta^{[t]} + \boxed{\nu} \nabla \mathcal{L}(h(x, \theta), y^*)$$

└─ Learning rate



Can be blocked in a local minimum (not that it matters much ...)

[Figure: C. Bishop, 2006]

Gradient descent

One optimizer step:

$$\theta^{[t+1]} = \theta^{[t]} + \nu \nabla \mathcal{L}(h(x, \theta), y^*)$$

The gradient is a vector of partial derivatives:


$$\nabla \mathcal{L} = \begin{bmatrix} \frac{\partial \mathcal{L}}{\partial \theta_0} \\ \frac{\partial \mathcal{L}}{\partial \theta_1} \\ \vdots \\ \frac{\partial \mathcal{L}}{\partial \theta_N} \end{bmatrix}$$

Differentiate a multi-layer network

Consider a multi-layer network, in particular an arbitrary unit indexed by j and receiving inputs from units indexed by i , providing outputs z_i .

Its gradient is:

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial w_{ji}} &= \frac{\partial \mathcal{L}}{\partial a_j} \frac{\partial a_j}{\partial w_{ji}} \\ &= \delta_j z_i\end{aligned}$$



δ_j can be obtained with a recurrence:

$$\delta_j = h'(a_j) \sum_k w_{kj} \delta_k$$

The full backpropagation algorithm

1. **Forward pass:** stimulate the model with input x , calculate all a_j and z_j up to the output y .
2. Calculate the δ_j for the output units using the derivative of the loss function.
3. **Backward pass:** calculate all δ_j with

$$\delta_j = h'(a_j) \sum_k w_{kj} \delta_k$$

4. Calculate the gradients with

$$\frac{\partial \mathcal{L}}{\partial w_{ji}} = \delta_j z_i$$



Backpropagation in PyTorch

Autograd

In PyTorch (and some other frameworks), Autograd performs automatic differentiation through a sequence of tensor instructions of an imperative language.

Let's consider a simple linear operation:

$$w = [5 \ 3], \quad x = [7 \ 2], \quad y = wx^T$$

The gradient of y w.r.t to x is given as

$$\nabla = \left[\frac{\partial y}{\partial x_i} \right] = \begin{bmatrix} 5 \\ 3 \end{bmatrix}$$

The gradient of y w.r.t to w is given as

$$\nabla = \left[\frac{\partial y}{\partial w_i} \right] = \begin{bmatrix} 7 \\ 1 \end{bmatrix}$$

Autograd

In PyTorch, we will first create the tensors:

```
1 w = torch.tensor([5, 3], dtype=float, requires_grad=True)
2 x = torch.tensor([7, 1], dtype=float, requires_grad=True)
```

The `requires_grad` flag ensures that all calculations are tracked. We perform the linear operation:

```
1 y = torch.dot(w, x)
```

Since the tensor y has been calculated as result of operations on tracked tensors, it has a gradient function:

```
1 print (y)
```

```
1 tensor(38., dtype=torch.float64, grad_fn=<DotBackward>)
```

Autograd

We now run a backward pass on the variable y , which calculates gradients w.r.t. to all involved tensors:

```
1 y.backward()
```

The gradients are attached to each variable:

```
1 print (x.grad)
2 print (w.grad)
```

```
1 tensor([5., 3.], dtype=torch.float64)
2 tensor([7., 1.], dtype=torch.float64)
```

- 1 Neural Networks & learning
- 2 Optimization / gradient descent
- 3 A simple linear example in PyTorch
- 4 A multi-layer model in PyTorch
- 5 Convolutional networks

Example: Wisconsin breast-cancer

Id	Cl.thickness	Cell.size	Cell.shape	Marg.adhesion	Epith.c.size	Bare.nuclei	Bl.cromatin	Normal.nucleoli	Mitoses	Class
1000025	5	1	1	1	2	1	3	1	1	benign
1002945	5	4	4	5	7	10	3	2	1	benign
1015425	3	1	1	1	2	2	3	1	1	benign
1016277	6	8	8	1	3	4	3	7	1	benign
1017023	4	1	1	3	2	1	3	1	1	benign
1017122	8	10	10	8	7	10	9	7	1	malignant
1018099	1	1	1	1	2	10	3	1	1	benign
1018561	2	1	2	1	2	1	3	1	1	benign
1033078	2	1	1	1	2	1	1	1	5	benign
1033078	4	2	1	1	2	1	2	1	1	benign
1035283	1	1	1	1	1	1	3	1	1	benign
1036172	2	1	1	1	2	1	2	1	1	benign
1041801	5	3	3	3	2	3	4	4	1	malignant

```
chris@pytorch $ head breast-cancer-wisconsin.csv
1000025,5,1,1,1,2,1,3,1,1,2
1002945,5,4,4,5,7,10,3,2,1,2
1015425,3,1,1,1,2,2,3,1,1,2
1016277,6,8,8,1,3,4,3,7,1,2
1017023,4,1,1,3,2,1,3,1,1,2
1017122,8,10,10,8,7,10,9,7,1,4
1018099,1,1,1,1,2,10,3,1,1,2
1018561,2,1,2,1,2,1,3,1,1,2
1033078,2,1,1,1,2,1,1,1,5,2
1033078,4,2,1,1,2,1,2,1,1,2
```


Data loading and conversion

```
1 import numpy as np
2 from numpy import genfromtxt
3 import torch
4 from torch.nn import functional as F
5
6 # Import the text file into a numpy array
7 n = genfromtxt('breast-cancer-wisconsin-cleaned.csv',
8               delimiter=',')
9 D = torch.tensor(n, dtype=torch.float32)
10 N_samples = D.size(0)
11
12 # The input is the full matrix without first and
13 # last column, plus the 1 column for the bias
14 X=D[:,1:-1]
15 X = torch.cat ((X, torch.ones((X.size(0),1))),1)
16
17 # The targets. Change all 2->0 and 4->1
18 T=D[:, -1:]
19 T[T==2]=0
20 T[T==4]=1
```

The model

```
1 class LogisticRegression(torch.nn.Module):
2     def __init__(self):
3         super(LogisticRegression, self).__init__()
4
5         # The linear layer (input dim, output dim)
6         # It also contains a weight matrix
7         # (here single output-> vector)
8         self.fc1 = torch.nn.Linear(10, 1)
9
10        # The forward pass of the network. x is the input
11    def forward(self, x):
12        return F.sigmoid(self.fc1(x))
```

Set up the environment

```
1 # Instantiate the model
2 model = LogisticRegression()
3
4 # The loss function: binary cross-entropy
5 criterion = torch.nn.BCELoss()
6
7 # Set up the optimizer: stochastic gradient descent
8 # with a learning rate of 0.01
9 optimizer = torch.optim.SGD(model.parameters(), lr
    =0.01)
```

Iterative training

```
1 # 1 epoch = 1 pass over the full dataset
2 for epoch in range(200):
3     print ("Starting epoch", epoch, " ", end='')
4     calcAccuracy()
5
6     for sample in range(N_samples):
7         # model -> train mode, clear gradients
8         model.train()
9         optimizer.zero_grad()
10
11         # Forward pass (stimulate model with inputs)
12         y = model(X[sample,:])
13
14         # Compute Loss
15         loss = criterion(y, T[sample])
16
17         # Backward pass: calculate the gradients
18         loss.backward()
19
20         # One step of stochastic gradient descent
21         optimizer.step()
```

Evaluation

Calculate the accuracy (in percent) at each epoch:

Proportion of correctly classified samples.

Random performance = 50% on a binary task.

```
1 def calcAccuracy():
2
3     # model -> eval mode
4     model.eval()
5     correct = 0.0
6     for sample in range(N_samples):
7
8         # threshold the output probability
9         y = 1 if model(X[sample,:]) > 0.5 else 0
10        correct += (y == T[sample]).numpy()
11
12    print ("Accuracy = ", 100.0*correct/N_samples)
```

Results

1	Starting epoch 0	Accuracy =	[34.69985359]
2	Starting epoch 1	Accuracy =	[91.21522694]
3	Starting epoch 2	Accuracy =	[93.99707174]
4	Starting epoch 3	Accuracy =	[95.16837482]
5	Starting epoch 4	Accuracy =	[95.75402635]
6	Starting epoch 5	Accuracy =	[95.60761347]
7	Starting epoch 6	Accuracy =	[95.75402635]
8	Starting epoch 7	Accuracy =	[95.75402635]
9	Starting epoch 8	Accuracy =	[96.33967789]
10	Starting epoch 9	Accuracy =	[96.48609078]
11	Starting epoch 10	Accuracy =	[96.63250366]
12	Starting epoch 11	Accuracy =	[96.63250366]
13	Starting epoch 12	Accuracy =	[96.63250366]
14	Starting epoch 13	Accuracy =	[96.63250366]
15	Starting epoch 14	Accuracy =	[96.63250366]

(...)

1	Starting epoch 197	Accuracy =	[97.07174231]
2	Starting epoch 198	Accuracy =	[97.07174231]
3	Starting epoch 199	Accuracy =	[97.07174231]

What is missing?


- The model is simpler than deep neural networks, but sufficient for the task.
- We did not use batch processing, i.e. using more than one sample for a given gradient update
- We calculated performance on the training set. We might overfit.

- 1 Neural Networks & learning
- 2 Optimization / gradient descent
- 3 A simple linear example in PyTorch
- 4 A multi-layer model in PyTorch
- 5 Convolutional networks

Example: the MNIST dataset

A dataset of handwritten digits introduced by Yann LeCun in 1999 with 60 000 training images and 10 000 test images.

One image is of size 28x28 pixels.



<http://yann.lecun.com/exdb/mnist/>

Writing Data Access

We subclass the Dataset class and implement the methods `__getitem__`, `__len__`() (and `__init__`!). This class will be used by actual PyTorch Dataloader, DataLoader.

```
1 import torch
2 from torch.utils.data import Dataset, DataLoader
3 from torchvision import transforms
4
5 class MNISTDataset(Dataset):
6     def __init__(self, dir, transform=None):
7         # Read in the files from the directory and
8         # store them in self.images and self.labels
9         # See MOODLE for the full code of this method
10        (...)
11
12        # The access is _NOT_ shuffled. The PyTorch
13        # Dataloader will need to do this.
14        def __getitem__(self, index):
15            return self.images[index], self.labels[index]
16        # Return the dataset size
17        def __len__(self):
18            return self.no_images
```

Writing Data Access

Let's recall training through gradient descent:

$$\theta^{[t+1]} = \theta^{[t]} + \nu \nabla \mathcal{L}(h(x, \theta), y^*)$$

The gradient is rarely (never?) taken over the whole dataset, but over a single sample, or batches (mini-batches) of a certain size. These batches are sample randomly from the dataset.

The actual shuffling and batching is performed by a built-in PyTorch DataLoader class, which uses an instance of our Dataset subclass:

```
1 dataset = MNISTDataset ("MNIST-png/testing" ,  
2     transforms.Compose([  
3     transforms.ToTensor() ,  
4     transforms.Normalize((0.1307 ,) , (0.3081 ,)) ]))  
5  
6 loader = torch.utils.data.DataLoader(dataset ,  
7     batch_size=50, shuffle=True)
```

We passed a set of image transforms to the Dataset class, which applies them to each image.

Tensor dimension conventions

PyTorch functions operate on multi-dimensional tensors and follow conventions on the order of dimensions.

- The first dimension is the batch dimension
 - ▶ Use 1 if you don't use batches (= batches of size 1).
 - ▶ Losses are reduced (sum or mean) over samples in a batch
- the second dimension is the channel dimension
 - ▶ Use 1 if you don't use channels (= single channels).
 - ▶ Channel arithmetics will be explained in detail in the section on convolutions.
- the following dimensions are application dependant, e.g. rows, columns in images.

The 2 layer model

The model is similar to our linear example. Differences:

- A hidden layer with 300 units and relu activation.
- A forward pass deals with a full batch.

```
1 class MLP(torch.nn.Module):
2     def __init__(self):
3         super(MLP, self).__init__()
4         # input size to 300 units
5         self.fc1 = torch.nn.Linear(28*28, 300)
6         # 300 units to 10 output classes
7         self.fc2 = torch.nn.Linear(300, 10)
8
9     def forward(self, x):
10        # Reshape from a 3D tensor (batchsize, 28, 28)
11        # to a flattened (batchsize, 28*28)
12        # 1 sample = 1 vector
13        x = x.view(-1, 28*28)
14        x = F.relu(self.fc1(x))
15        return self.fc2(x)
```

Setup model, loss, optimizer

This time we have more than 2 classes, we use the Cross-entropy loss.

```
1 # Instantiate the model
2 model = MLP()
3
4 # This criterion combines LogSoftMax and NLLLoss
5 # in one single class.
6 crossentropy = torch.nn.CrossEntropyLoss()
7
8 # Set up the optimizer: stochastic gradient descent
9 # with a learning rate of 0.01
10 optimizer = torch.optim.SGD(model.parameters(), lr
    =0.01)
11
12 # Init some statistics
13 running_loss = 0.0
14 running_correct = 0
15 running_count = 0
```

Cycling through batches and samples

```
1 # Cycle through epochs
2 for epoch in range(100):
3
4     # Cycle through batches. One batch is a set of
5     # images and a set of ground truth labels
6     for batch_idx, (data, labels) in enumerate(loader):
7
8         optimizer.zero_grad()
9
10        # We calculate a prediction on the full batch
11        y = model(data)
12
13        # Loss of the full batch, summed
14        loss = crossentropy(y, labels)
15
16        # Calculate gradients of the full batch
17        loss.backward()
18
19        # One gradient update
20        optimizer.step()
```

Track training error

```
1  # Calculate the winner class
2  _, predicted = torch.max(y.data, 1)
3
4  # How many correct samples?
5  running_correct += (predicted == labels).sum().item()
6  running_count += BATCHSIZE
7
8  # Every 100 batches, print statistics
9  if (batch_idx % 100) == 0:
10
11     train_err = 100.0*(1.0 - running_correct / running_count)
12     print ('Epoch: %d batch: %5d ' % (epoch + 1, batch_idx +
13         1), end="")
14     print ('train-loss: %.3f train-err: %.3f' % (
15         running_loss / 100, train_err))
16     running_loss = 0.0
17     running_correct = 0.0
18     running_count=0.0
```


Example output

```
1 Epoch: 1 batch:      1 train-loss: 0.024 train-err: 98.000
2 Epoch: 1 batch:    101 train-loss: 1.576 train-err: 36.140
3 Epoch: 1 batch:    201 train-loss: 0.747 train-err: 16.320
4 Epoch: 1 batch:    301 train-loss: 0.539 train-err: 12.880
5 Epoch: 1 batch:    401 train-loss: 0.481 train-err: 13.000
6 Epoch: 1 batch:    501 train-loss: 0.414 train-err: 11.280
7 Epoch: 1 batch:    601 train-loss: 0.386 train-err: 10.380
8 Epoch: 1 batch:    701 train-loss: 0.385 train-err: 10.900
9 Epoch: 1 batch:    801 train-loss: 0.363 train-err: 10.540
10 Epoch: 1 batch:    901 train-loss: 0.320 train-err: 9.120
11 Epoch: 1 batch:   1001 train-loss: 0.323 train-err: 8.920
12 Epoch: 1 batch:   1101 train-loss: 0.325 train-err: 9.400
13 Epoch: 2 batch:      1 train-loss: 0.304 train-err: 8.880
```

(...)


```
1 Epoch: 75 batch:    801 train-loss: 0.007 train-err: 0.000
2 Epoch: 75 batch:    901 train-loss: 0.007 train-err: 0.020
3 Epoch: 75 batch:   1001 train-loss: 0.008 train-err: 0.000
4 Epoch: 75 batch:   1101 train-loss: 0.009 train-err: 0.040
```

This is training error, not validation error, i.e. **NOT** representative of the performance of the model!

- 1 Neural Networks & learning
- 2 Optimization / gradient descent
- 3 A simple linear example in PyTorch
- 4 A multi-layer model in PyTorch

5 Convolutional networks

Fully connected layers might be harmful



Fully connected layers might be harmful

Parameters learned for one part of the image do not generalize to other parts of the image.



$\neq$



Shift-invariant linear operations

Classical MLPs are sequences of linear layers followed by point-wise non-linearities. What kind of linear and shift-invariant operators can there exist?

Let's consider an operator ϕ with following properties:

Linearity:

$$\phi(\alpha f + \beta f') = \alpha \phi(f) + \beta \phi(f')$$

Shift-invariance:

$$\phi({}^{m,n}S(f)) = {}^{m,n}S(\phi(f))$$


where ${}^{m,n}S(f)$ shifts a signal by m,n .

Impulse response h :



$$h = \phi({}^{0,0}p)$$

where ${}^{0,0}p$ is a Dirac impulse centered at 0, 0.

Convolutions




Convolutions




$$(3+2*5+3+2*4+5*5+2*7+3+2*7+1)/16 = 5$$


Convolutions




Convolutions



Convolutions



« LeNet »



"LeNet", by Yann LeCun, is a network composed of sequences of convolutions and spatial size reductions (« pooling »).

N (layer) input
channels


2	5	2	17	10	15
3	5	3	17	19	15
4	5	7	18	19	14
3	7	1	16	29	1
1	1	2	5	3	1
2	2	3	13	5	3
2	3	2	15	4	4

Each of the M output channels is a linear
combination of the N filtered input images.

We need $N \cdot M$ filters!

M (layer) output
channels

Trainable
filters



+bias

Template matching


Input image:



Filter kernel

- We normalize input and kernel (subtract mean, divide by standard deviation).
- We convolve the input with the kernel
- We threshold response and superimpose it on the red channel.

Template matching




Convolutions can match patterns:

- either in input space
- or in intermediate network layers (which still have a spatial organization).

LeNet in PyTorch

```
1 class LeNet(torch.nn.Module):
2     def __init__(self):
3         super(LeNet, self).__init__()
4         # 1 input channel, 20 output channels,
5         # 5x5 filters, stride=1, no padding
6         self.conv1 = torch.nn.Conv2d(1, 20, 5, 1, 0)
7         self.conv2 = torch.nn.Conv2d(20, 50, 5, 1, 0)
8         self.fc1 = torch.nn.Linear(4*4*50, 500)
9         self.fc2 = torch.nn.Linear(500, 10)
10
11     def forward(self, x):
12         x = F.relu(self.conv1(x))
13         # Max pooling with a filter size of 2x2
14         # and a stride of 2
15         x = F.max_pool2d(x, 2, 2)
16         x = F.relu(self.conv2(x))
17         x = F.max_pool2d(x, 2, 2)
18         x = x.view(-1, 4*4*50)
19         x = F.relu(self.fc1(x))
20         return self.fc2(x)
```


MLP vs LeNet on MNIST



Blue: training, Red: validation

Going deeper


2012 : AlexNet, 8 layers. New techniques: dropout, ReLU




2014 : GoogLeNet, 20 layers. New technique: intermediate supervision



2015 : Microsoft research, 150 layers. New technique: residual learning




Visualizing feature map activations




[Zeiler and Fergus,
ECCV 2014]

Visualizing feature map activations




[Zeiler and Fergus,
ECCV 2014]

Visualizing feature map activations



[Zeiler and Fergus,
ECCV 2014]

Conclusion



Problems to solve:

- Expressivity: model symmetries and invariances in your data
- Trainability
- Generalization

1960-2012: let's handcraft features

2012-2016: let's handcraft architectures (layers, units ...)

2016- ? : let's handcraft data-flow (geometry, causality, attention)