

Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

The Fundamentals of Deep Learning with Applications

Jon Krohn

`jon@untapt.com`

Chief Data Scientist at untapt

NYC Open Data Meetup

February 28th, 2017

(slides available at `jonkrohn.com/talks`)



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

- 1 Antecedents**
Case Study: A History of Biological & Artificial Vision
Building Blocks
- 2 Theory**
Biological & Artificial Neurons
Neural Networks
Deep Neural Networks
- 3 Contemporary Applications**
Convolutional Neural Networks
Long Short-Term Memory Recurrent Neural Networks
Deep Learning at untapt
Deep Reinforcement Learning



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

- 1 Antecedents**
Case Study: A History of Biological & Artificial Vision
Building Blocks
- 2 Theory**
Biological & Artificial Neurons
Neural Networks
Deep Neural Networks
- 3 Contemporary Applications**
Convolutional Neural Networks
Long Short-Term Memory Recurrent Neural Networks
Deep Learning at untapt
Deep Reinforcement Learning



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

- 1 Antecedents
 - Case Study: A History of Biological & Artificial Vision
 - Building Blocks
- 2 Theory
 - Biological & Artificial Neurons
 - Neural Networks
 - Deep Neural Networks
- 3 Contemporary Applications
 - Convolutional Neural Networks
 - Long Short-Term Memory Recurrent Neural Networks
 - Deep Learning at untapt
 - Deep Reinforcement Learning



Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

ConvNets

LSTMs

untapt

Reinforcement

1 Antecedents

Case Study: A History of Biological & Artificial Vision
Building Blocks

2 Theory

Biological & Artificial Neurons
Neural Networks
Deep Neural Networks

3 Contemporary Applications

Convolutional Neural Networks
Long Short-Term Memory Recurrent Neural Networks
Deep Learning at untapt
Deep Reinforcement Learning



Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets


Application

ConvNets

LSTMs

untapt

Reinforcement



Hubel & Wiesel (1959)

Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets


Application

ConvNets

LSTMs

untapt

Reinforcement



Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

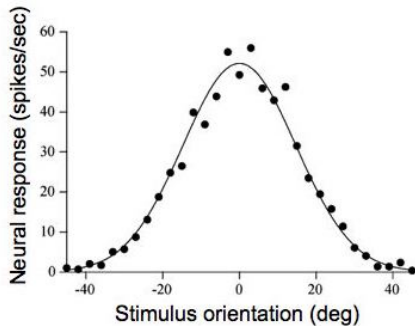

Application

ConvNets

LSTMs

untapt

Reinforcement



Hubel & Wiesel, 1968



Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application


ConvNets

LSTMs

untapt

Reinforcement


topographical mapping



hyper-complex cells

complex cells

simple cells





high level

mid level

low level

low level



Hubel & Wiesel, 1968



Antecedents

- Vision Case Study
- Building Blocks


Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement

Visual Cortices



Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

ConvNets


LSTMs

untapt

Reinforcement

Camera Obscura

da Vinci (15th Century)



Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

ConvNets


LSTMs

untapt


Reinforcement

Block World


Larry Roberts (1965)




(a) Original picture.



(b) Differentiated picture.



(c) Line drawing.



(d) Rotated view.



Viola & Jones (2001)

Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets


Application

ConvNets

LSTMs

untapt

Reinforcement



Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

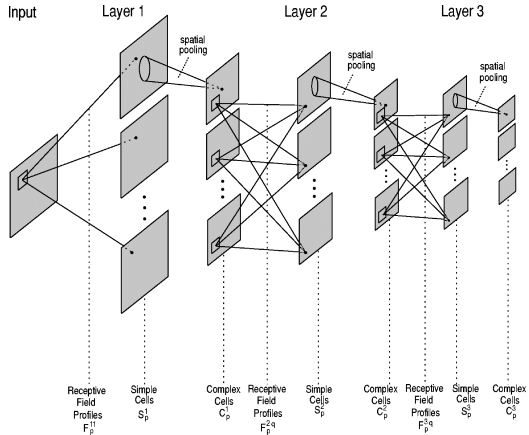
Application

ConvNets

LSTMs

untapt

Reinforcement



MNIST Digits & LeNet-5

LeCun, Boutou, Bengio & Haffner (1998)



PROC. OF THE IEEE, NOVEMBER 1998

7




Fig. 2. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.



Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

ConvNets


LSTMs

untapt

Reinforcement

LeNet-5

LeCun, Boutou, Bengio & Haffner (1998)



Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

ConvNets

LSTMs

untapt

Reinforcement



Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

ConvNets

LSTMs

untapt

Reinforcement



ImageNet

Fei-Fei Li et al. (2009), 14m images, 22k categories

Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- unlapt
- Reinforcement




mite

container ship

motor scooter

leopard

	<p>mite</p> <p>black widow</p> <p>cockroach</p> <p>tick</p> <p>starfish</p>		<p>container ship</p> <p>lifeboat</p> <p>amphibian</p> <p>fireboat</p> <p>drilling platform</p>		<p>motor scooter</p> <p>go-kart</p> <p>moped</p> <p>bumper car</p> <p>golfcart</p>		<p>leopard</p> <p>jaguar</p> <p>cheetah</p> <p>snow leopard</p> <p>Egyptian cat</p>
--	--	--	--	--	---	--	--



grille

mushroom

cherry

Madagascar cat

	<p>convertible</p> <p>grille</p> <p>pickup</p> <p>beach wagon</p> <p>fire engine</p>		<p>agaric</p> <p>mushroom</p> <p>jelly fungus</p> <p>gill fungus</p> <p>dead-man's-fingers</p>		<p>dalmatian</p> <p>grape</p> <p>elderberry</p> <p>ffordshire bullterrier</p> <p>currant</p>		<p>squirrel monkey</p> <p>spider monkey</p> <p>titi</p> <p>indri</p> <p>howler monkey</p>
--	---	--	---	--	---	--	--



ImageNet Classification Error

ILSVRC: 1.4m, 1k object classes

Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

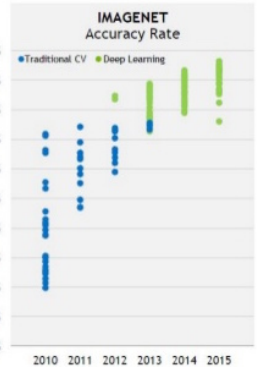

ConvNets

LSTMs

untapt

Reinforcement

ILSVRC top-5 error on ImageNet



AlexNet

Krizhevsky, Sutskever & Hinton (2012)

Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

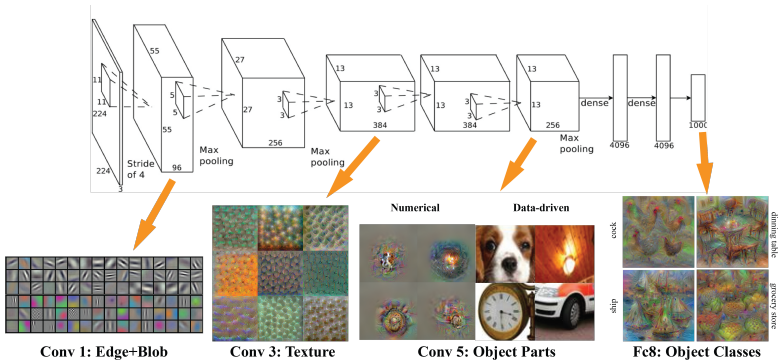
Application

ConvNets

LSTMs

unlapt

Reinforcement



Deep Learning

Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets



Application

ConvNets

LSTMs

untapt

Reinforcement



Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

ConvNets

LSTMs

unlapt

Reinforcement

INT. SHOP

We see H pull a book from a shelf, flip through it while speaking, and then put it back.

H
In a future with mass unemployment, young people are forced to sell blood. That's the first thing I can do.

H2
You should see the boys and shut up. I was the one who was going to be a hundred years old.

H
I saw him again. The way you were sent to me... that was a big honest idea. I am not a bright light.

C
Well, I have to go to the skull. I don't know.

He picks up a light screen and fights the security force of the particles of a transmission on his face.

H
(continuing)
What do you mean?

C
(smiles)
I don't know anything about any of this.

H
(to Hank, taking his eyes from his mouth)
Then what?

H2
There's no answer.

C
(frowning)
We're going to see the money.

H
(reading)
'All right, you can't tell me that.'

steps back. Coffey is still going through.

C
I was coming to that thing because you were so pretty.

H
I don't know. I don't know what you're talking about.

C
That's right.

H
So what are you doing?

H2
I don't want to be honest with you.

He looks at him for a moment, then smiles at him.

H
You don't have to be a doctor.

H2
I am not sure. I don't know what you're talking about.

H
I want to see you too.

H2
What do you mean?

H
I'm sorry, but I'm sure you wouldn't even touch me.

H2
I don't know what you're talking about.

H
The principle is completely constructed for the same time.

H2
(smiling)
It was all about you to be true.

H
You didn't even see the movie with the rest of the bass.

H2
I don't know.

H
I don't care.

H2
I know that it's a consequence. Whatever you want to know about the presence of the story, I'm a little bit of a boy on the floor.

H
I don't know. I just have to ask you to explain to me what you say.

H2
What do you mean?

H
Because I don't know what you're talking about.

H2
That was all the time.

H
I know that.

H2
I don't know.

H
(smiley)
It would be a good time. I think I could have been my life.

He starts to shake.

H (CONT'D)
It may never be forgiven, but that is just too bad. I have to leave, but I'm not free of the world.

C
Yes. Perhaps I should take it from there. I'm not going to do something.

H
You can't afford to take this anymore. It's not a game. But I've got a good time to stay there.

C
Well, I think you can still be back on the table.

H
Wow. It's a damn thing scared to say. Nothing is going to be a thing but I was the one that got on this rock with a child and then I left the other two.

He is standing in the store and sitting on the floor. He takes a seat on the counter and pulls the camera over to his back. He stares at it. He is on the phone. He cuts the shotgun from the edge of the room and puts it in his mouth. He sees a black hole in the floor leading to the man on the roof.

He comes up behind him to protect him. He is still standing next to him.

He looks through the door and the door closes. He looks at the boy from his backpack, and starts to cry.

T

Well, there's the situation with me and the light on the ship. The guy was trying to stop me. He was like a boy and he was young. I was worried about him, but even if he would have done it all. He couldn't even say sorry. I didn't mean to be a virgin. I mean, he was weak. And I thought I'd change my mind. He was crazy to take it out. It was a long time ago. He was a little late. I was going to be a moment. I just wanted to see you that I see much better than he did. I had to stop him and I couldn't even call. I didn't want to hurt him. I'm sorry. I know I don't like him. I can go home and be so bad and I love him. So I can get him all the way over here and find the square and go to the game with him and she won't show up. Then I'll check it out. But I'm going to see him when he gets to me. He looks like me and he throws me out of his eyes. Then he said he'd go to bed with me.



Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

ConvNets

LSTMs

untapt

Reinforcement

Sunspring

Sharp & Goodwin (2016)

[video]



Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

ConvNets

LSTMs

untapt

Reinforcement

1 Antecedents

Case Study: A History of Biological & Artificial Vision
Building Blocks

2 Theory

Biological & Artificial Neurons
Neural Networks
Deep Neural Networks

3 Contemporary Applications

Convolutional Neural Networks
Long Short-Term Memory Recurrent Neural Networks
Deep Learning at untapt
Deep Reinforcement Learning



Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

ConvNets

LSTMs

untapt

Reinforcement

Hardware

- local machine
- build your own server
- AWS / Google Cloud Platform
- GPU(s) / TPU(s)



Popular Libraries

based on Johnson (2016) in Stanford CS231n I.12

Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

ConvNets

LSTMs

untapt

Reinforcement

	Caffe	Torch	Theano	TensorFlow
<i>Language</i>	C++, Python	Lua	Python	Python
<i>Pretrained</i>	Yes++	Yes++	Yes (Lasagne)	Inception
<i>Parallel GPUs: Data</i>	Yes	Yes	Yes	Yes
<i>Parallel GPUs: Model</i>	No	Yes	Experimental	Yes (best)
<i>Readable Source Code</i>	Yes (C++)	Yes	No	No
<i>Good at RNN</i>	No	Mediocre	Yes	Yes (best)
<i>Higher-Level APIs</i>	No	No	Keras	Keras and TFLearn



[Human Learning Resources]

Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

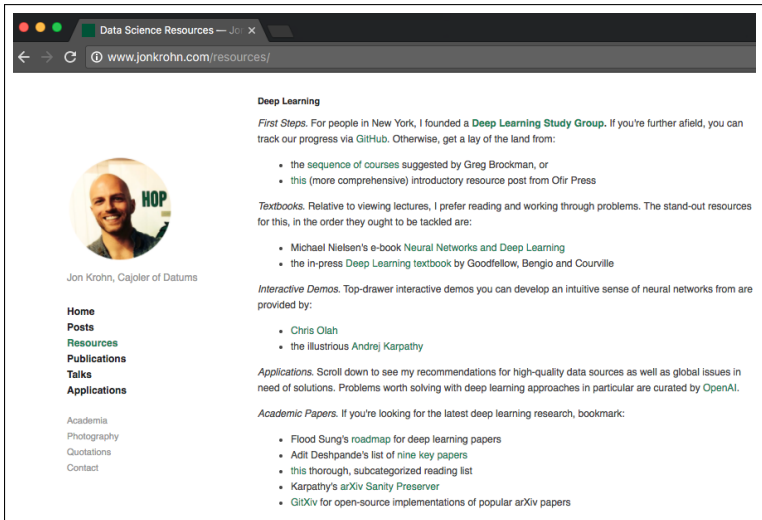
Application

ConvNets

LSTMs

untapt

Reinforcement



Data Science Resources — Jon Krohn

www.jonkrohn.com/resources/

Deep Learning

First Steps. For people in New York, I founded a **Deep Learning Study Group**. If you're further afield, you can track our progress via GitHub. Otherwise, get a lay of the land from:

- the sequence of courses suggested by Greg Brockman, or
- this (more comprehensive) introductory resource post from Ofir Press

Textbooks. Relative to viewing lectures, I prefer reading and working through problems. The stand-out resources for this, in the order they ought to be tackled are:

- Michael Nielsen's e-book *Neural Networks and Deep Learning*
- the in-press *Deep Learning* textbook by Goodfellow, Bengio and Courville

Interactive Demos. Top-drawer interactive demos you can develop an intuitive sense of neural networks from are provided by:

- Chris Olah
- the illustrious Andrej Karpathy

Applications. Scroll down to see my recommendations for high-quality data sources as well as global issues in need of solutions. Problems worth solving with deep learning approaches in particular are curated by OpenAI.

Academic Papers. If you're looking for the latest deep learning research, bookmark:

- Flood Sung's roadmap for deep learning papers
- Adit Deshpande's list of nine key papers
- this thorough, subcategorized reading list
- Karpathy's arXiv Sanity Preserver
- GitXiv for open-source implementations of popular arXiv papers

Home

Posts

Resources

Publications

Talks

Applications

Academia

Photography

Quotations

Contact

Jon Krohn, Cajoler of Datums



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

- 1 Antecedents
 - Case Study: A History of Biological & Artificial Vision Building Blocks
- 2 Theory
 - Biological & Artificial Neurons**
 - Neural Networks
 - Deep Neural Networks
- 3 Contemporary Applications
 - Convolutional Neural Networks
 - Long Short-Term Memory Recurrent Neural Networks
 - Deep Learning at untapt
 - Deep Reinforcement Learning



Biological Neuron Morphology

Antecedents


- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



Perceptron

Rosenblatt (1957)

Antecedents


- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



$$\text{output} = \begin{cases} 0 & \text{if } \sum_j w_j x_j \leq \text{threshold} \\ 1 & \text{if } \sum_j w_j x_j > \text{threshold} \end{cases}$$



Antecedents

- Vision Case Study
- Building Blocks

Theory

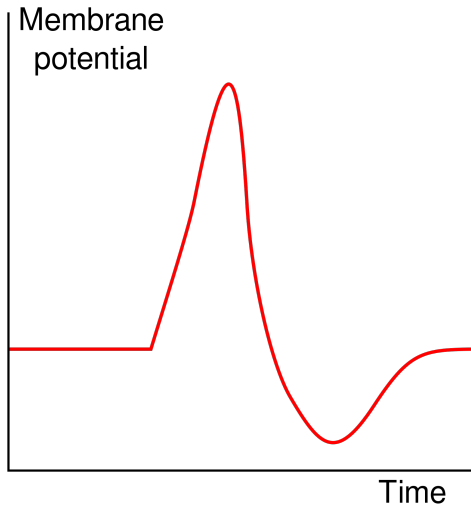
- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement

Biological Neuron Physiology

The *Binary* Action Potential



Antecedents

- Vision Case Study
- Building Blocks

Theory


- Neural Units**
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement

Perceptron

Rosenblatt (1957)



Multi-Layer Perceptron

Antecedents


- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



Multi-Layer Perceptron

Antecedents


- Vision Case Study
- Building Blocks

Theory


- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



Sigmoid Neuron



$$\frac{1}{1 + \exp(-\sum_j w_j x_j - b)}$$



Antecedents

- Vision Case Study
- Building Blocks


Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement

tanh Neuron



$$\sigma(z) = \frac{1 + \tanh(z/2)}{2}$$



ReLU: Rectified Linear Units

Nair & Hinton (2010); Maas, Hannun & Ng (2014)

Antecedents



- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

ConvNets

LSTMs

untapt

Reinforcement

- 1 Antecedents
 - Case Study: A History of Biological & Artificial Vision
 - Building Blocks
- 2 Theory
 - Biological & Artificial Neurons
 - Neural Networks**
 - Deep Neural Networks
- 3 Contemporary Applications
 - Convolutional Neural Networks
 - Long Short-Term Memory Recurrent Neural Networks
 - Deep Learning at untapt
 - Deep Reinforcement Learning



MNIST

LeCun, Cortes & Burges

Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets**
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



Fully-Connected Neural Net

Single Hidden Layer

Antecedents


- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets**
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



TensorFlow Playground

Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

[demo]



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

- 1 Antecedents
 - Case Study: A History of Biological & Artificial Vision
 - Building Blocks
- 2 Theory
 - Biological & Artificial Neurons
 - Neural Networks
 - Deep Neural Networks
- 3 Contemporary Applications
 - Convolutional Neural Networks
 - Long Short-Term Memory Recurrent Neural Networks
 - Deep Learning at untapt
 - Deep Reinforcement Learning



Deep Fully-Connected Net

3 (or more) Hidden Layers

Antecedents


- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement

	Caffe	Torch	Theano	TensorFlow
<i>Language</i>	C++, Python	Lua	Python	Python
<i>Pretrained</i>	Yes++	Yes++	Yes (Lasagne)	Inception
<i>Parallel GPUs: Data</i>	Yes	Yes	Yes	Yes
<i>Parallel GPUs: Model</i>	No	Yes	Experimental	Yes (best)
<i>Readable Source Code</i>	Yes (C++)	Yes	No	No
<i>Good at RNN</i>	No	Mediocre	Yes	Yes (best)
<i>Higher-Level APIs</i>	No	No	Keras	Keras and TFLearn



A Simple Deep Net in TFLearn

Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets**

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement

[notebook]



Synaptic Pruning

Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

ConvNets

LSTMs

untapt

Reinforcement



(Stochastic) Gradient Descent

Adam = AdaGrad + RMSprop

Antecedents


Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement



Backpropagation

computes error & gradient of cost function

Antecedents

- Vision Case Study
- Building Blocks


Theory

- Neural Units
- Neural Nets

Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



$$\delta^L = \nabla_a C \odot \sigma'(z^L) \tag{BP1}$$

$$\delta^l = ((w^{l+1})^T \delta^{l+1}) \odot \sigma'(z^l) \tag{BP2}$$

$$\frac{\partial C}{\partial b_j^l} = \delta_j^l \tag{BP3}$$

$$\frac{\partial C}{\partial w_{jk}^l} = a_k^{l-1} \delta_j^l \tag{BP4}$$



Overfitting

...and avoiding it

Antecedents


Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement



- L1/L2 regularization
- dropout
- artificial data set expansion

Improving Neural Networks

Mostly Hyperparameter Tuning

Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

ConvNets

LSTMs

untapt

Reinforcement

- problem simplification
- number and width of layers
- cost fxn: quadratic, cross-entropy, log-likelihood, &c.
- more epochs, early stopping
- clever initialization of weights and biases
- learning rate η , variable schedule
- regularization parameter λ
- mini-batch size
- automation, e.g., with Spearmint

[Summary Blog Post]



Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets



Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement

Universality

Solve Any Continuous Function (Nielsen, 2015)



Unstable Gradient

Typically *Vanishes* (but can *Explode*)

Antecedents


- Vision Case Study
- Building Blocks

Theory


- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



$$\frac{\partial C}{\partial b_1} = \sigma'(z_1) \times w_2 \times \sigma'(z_2) \times w_3 \times \sigma'(z_3) \times w_4 \times \sigma'(z_4) \times \frac{\partial C}{\partial a_4}$$



Antecedents


- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



Classic Deep Architectures

...introducing *Convolutional Layers*

Antecedents

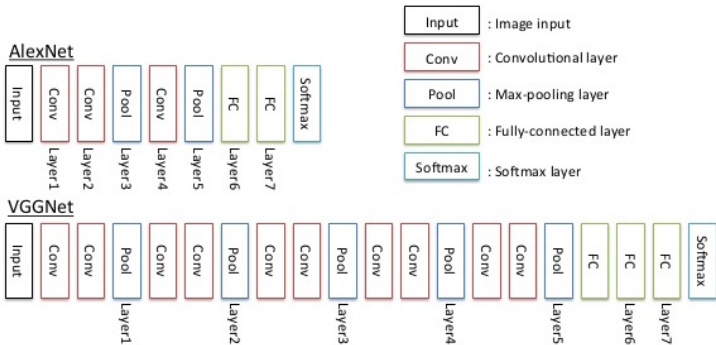
- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

- 1 Antecedents
 - Case Study: A History of Biological & Artificial Vision
 - Building Blocks
- 2 Theory
 - Biological & Artificial Neurons
 - Neural Networks
 - Deep Neural Networks
- 3 Contemporary Applications
 - Convolutional Neural Networks**
 - Long Short-Term Memory Recurrent Neural Networks
 - Deep Learning at untapt
 - Deep Reinforcement Learning



Hubel & Wiesel (1959)

Antecedents


Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement



Antecedents

Vision Case Study
Building Blocks


Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement


topographical mapping



hyper-complex cells

complex cells

simple cells





high level

mid level

low level

low level



Hubel & Wiesel, 1968



Antecedents

- Vision Case Study
- Building Blocks


Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement

Visual Cortices



AlexNet

Krizhevsky, Sutskever & Hinton (2012)

Antecedents

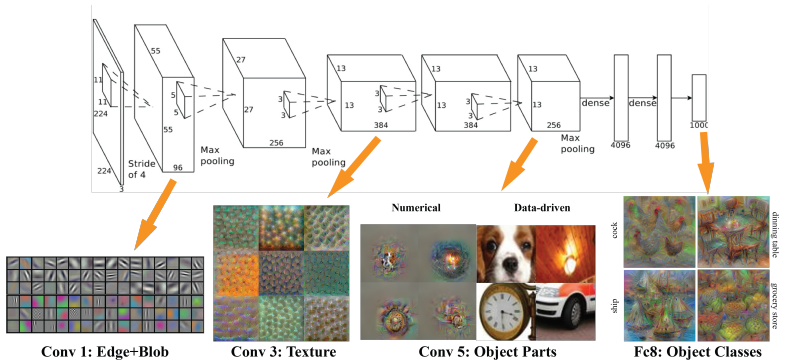
- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- unlapt
- Reinforcement



Antecedents

- Vision Case Study
- Building Blocks


Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement

CONVOLUTIONS



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

ConvNet Visualisation

Yosinski et al. (2015)

[video]



Network Architectures

Antecedents


- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



AlexNet: ILSVRC '12 winner

Krizhevsky et al. (2012)

Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets**
- LSTMs
- untapt
- Reinforcement

[TFLearn notebook]



VGGNet: ILSVRC '14 runner-up

Simonyan & Zisserman (2015)

Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets**
- LSTMs
- untapt
- Reinforcement

[TFLearn notebook]



ConvNet in TensorFlow

Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement

	Caffe	Torch	Theano	TensorFlow
<i>Language</i>	C++, Python	Lua	Python	Python
<i>Pretrained</i>	Yes++	Yes++	Yes (Lasagne)	Inception
<i>Parallel GPUs: Data</i>	Yes	Yes	Yes	Yes
<i>Parallel GPUs: Model</i>	No	Yes	Experimental	Yes (best)
<i>Readable Source Code</i>	Yes (C++)	Yes	No	No
<i>Good at RNN</i>	No	Mediocre	Yes	Yes (best)
<i>Higher-Level APIs</i>	No	No	Keras	Keras and TFLearn



ConvNet in TensorFlow

Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

ConvNets

- LSTMs
- untapt
- Reinforcement

[notebook]



Antecedents


Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement



ConvNet in Theano

Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement

	Caffe	Torch	Theano	TensorFlow
<i>Language</i>	C++, Python	Lua	Python	Python
<i>Pretrained</i>	Yes++	Yes++	Yes (Lasagne)	Inception
<i>Parallel GPUs: Data</i>	Yes	Yes	Yes	Yes
<i>Parallel GPUs: Model</i>	No	Yes	Experimental	Yes (best)
<i>Readable Source Code</i>	Yes (C++)	Yes	No	No
<i>Good at RNN</i>	No	Mediocre	Yes	Yes (best)
<i>Higher-Level APIs</i>	No	No	Keras	Keras and TFLearn



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

ConvNet in Theano

[demo]



ConvNet in Keras

calls TensorFlow or Theano

Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

ConvNets

- LSTMs
- untapt
- Reinforcement

	Caffe	Torch	Theano	TensorFlow
<i>Language</i>	C++, Python	Lua	Python	Python
<i>Pretrained</i>	Yes++	Yes++	Yes (Lasagne)	Inception
<i>Parallel GPUs: Data</i>	Yes	Yes	Yes	Yes
<i>Parallel GPUs: Model</i>	No	Yes	Experimental	Yes (best)
<i>Readable Source Code</i>	Yes (C++)	Yes	No	No
<i>Good at RNN</i>	No	Mediocre	Yes	Yes (best)
<i>Higher-Level APIs</i>	No	No	Keras	Keras and TFLearn



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

ConvNet in Keras

calls TensorFlow or Theano

[notebook]



“2.5-dimension” CT Scans

Roth et al. (2015)

Antecedents


Vision Case Study
Building Blocks

Theory


Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement





2.5D



Computer-Aided Detection

Shin et al. (2016); Roth et al. (2016)


Experimental Results (~100% sensitivity but ~40 FPs/patient at candidate generation step; then 3-fold CV with data augmentation)



Mediastinum

71% @ 3 FPs (was 55%)


Distance (≤ 15) - Lymph Node FROC Curve for 15 Patients



Abdomen

83% @ 3 FPs (was 30%)

Distance (≤ 15) - Lymph Node FROC Curve for 14 Patients



Antecedents

Vision Case Study
Building Blocks


Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets

LSTMs
untapt
Reinforcement



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

Kaggle

Data Science Bowl 2017

[link]



Transfer Learning

Caffe

Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement

	Caffe	Torch	Theano	TensorFlow
<i>Language</i>	C++, Python	Lua	Python	Python
<i>Pretrained</i>	Yes++	Yes++	Yes (Lasagne)	Inception
<i>Parallel GPUs: Data</i>	Yes	Yes	Yes	Yes
<i>Parallel GPUs: Model</i>	No	Yes	Experimental	Yes (best)
<i>Readable Source Code</i>	Yes (C++)	Yes	No	No
<i>Good at RNN</i>	No	Mediocre	Yes	Yes (best)
<i>Higher-Level APIs</i>	No	No	Keras	Keras and TFLearn



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

Transfer Learning

Caffe

[Model Zoo]



Antecedents

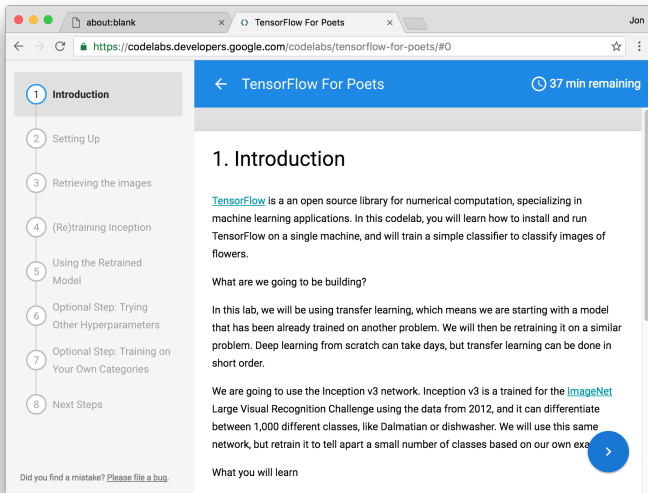
Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement



The screenshot shows a web browser window with the URL `https://codelabs.developers.google.com/codelabs/tensorflow-for-poets/#0`. The page title is "TensorFlow For Poets" and it indicates "37 min remaining". A table of contents on the left lists 8 steps, with "1 Introduction" selected. The main content area displays the title "1. Introduction" and the following text:

[TensorFlow](#) is an open source library for numerical computation, specializing in machine learning applications. In this codelab, you will learn how to install and run TensorFlow on a single machine, and will train a simple classifier to classify images of flowers.

What are we going to be building?

In this lab, we will be using transfer learning, which means we are starting with a model that has been already trained on another problem. We will then be retraining it on a similar problem. Deep learning from scratch can take days, but transfer learning can be done in short order.

We are going to use the Inception v3 network. Inception v3 is trained for the [ImageNet](#) Large Visual Recognition Challenge using the data from 2012, and it can differentiate between 1,000 different classes, like Dalmatian or dishwasher. We will use this same network, but retrain it to tell apart a small number of classes based on our own exa

What you will learn

At the bottom left, there is a link: "Did you find a mistake? [Please file a bug.](#)"



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

Video Classification

[video]



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

- 1 Antecedents
 - Case Study: A History of Biological & Artificial Vision
 - Building Blocks
- 2 Theory
 - Biological & Artificial Neurons
 - Neural Networks
 - Deep Neural Networks
- 3 Contemporary Applications
 - Convolutional Neural Networks
 - Long Short-Term Memory Recurrent Neural Networks**
 - Deep Learning at untapt
 - Deep Reinforcement Learning



Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs**
- untapt
- Reinforcement



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

INT. SHIP

We see H pull a book from a shelf, flip through it while speaking, and then put it back.

H
In a future with mass unemployment, young people are forced to sell blood. That's the first thing I can do.

H2
You should see the boys and shut up. I was the one who was going to be a hundred years old.

H
I saw him again. The way you were sent to her... that was a big honest idea. I am not a bright light.

C
Well, I have to go to the skull. I don't know.

H
He picks up a light screen and fights the security force of the particles of a transmission on his face.

H
(continuing)
What do you mean?

C
(smiling)
I don't know anything about any of this.

H
(to Hank, taking his eyes from his mouth)
Then what?

H2
There's no answer.

C
(frowning)
We're going to see the money.

H
(reading)
"All right, you can't tell me that."

H
steps back. Coffey is still going through.

C
I was coming to that thing because you were so pretty.

H
I don't know. I don't know what you're talking about.

C
That's right.

H
So what are you doing?

H2
I don't want to be honest with you.

H
He looks at him for a moment, then smiles at him.

H2
You don't have to be a doctor.

H
I am not sure. I don't know what you're talking about.

H
I want to see you too.

H2
What do you mean?

H
I'm sorry, but I'm sure you wouldn't even touch me.

H2
I don't know what you're talking about.

H
The principle is completely contrived for the same idea.

H2
(smiling)
I don't allow about you to be true.

H
You didn't even see the movie with the rest of the base.

H2
I don't know.

H
I don't care.

H2
I know that it's a consequence. Whatever you want to know about the presence of the story, I'm a little bit of a boy on the floor.

H
I don't know. I just have to ask you to explain to me what you say.

H2
What do you mean?

H
Because I don't know what you're talking about.

H2
That was all the time.

H
I know that.

H2
I don't know.

H
(angry)
It would be a good time. I think I could have been my life.

H
He starts to shake.

H (CONT'D)
It may never be forgiven, but that is just too bad. I have to leave, but I'm not free of the world.

C
The... Perhaps I should take it from here. I'm not going to do something.

H
You can't afford to take this anywhere. It's not a dream. But I've got a good time to stay there.

C
Well, I think you can still be back on the table.

H
Man. It's a damn thing scared to say. Holding the glass to be a thing, but I was the one that got on this rock with a child and then I left the other two.

H
He is standing in the stars and sitting on the floor. He takes a seat on the counter and pulls the camera over to his back. He stares at it. He is on the phone. He cuts the shotgun from the edge of the room and puts it in his mouth. He sees a black hole in the floor leading to the sea on the roof.

H
He comes up behind him to protect him. He is still standing next to him.

H
He looks through the door and the door closes. He looks at the bag from his backpack, and starts to cry.

T
Well, there's the situation with me and the light on the ship. The guy was trying to stop me. He was like a baby and he was gone. I was worried about him, but even if he would have done it all. He couldn't come any more. I didn't want to be a virgin. I mean, he was weak. And I thought I'd change my mind. He was crazy to take it out. It was a long time ago. He was a little late. I was going to be a moment. I just wanted to tell you that I was much better than he did. I had to stop him and I couldn't even tell. I didn't want to hurt him. I'm sorry. I know I don't like him. I can go home and be an bad and I love him. So I can get him all the way over here and find the square and go to the game with him and she won't show up. Then I'll check it out. But I'm going to see him when he gets to me. He looks like he and he should be out of his eyes. Then he said he'd go to bed with me.



Antecedents

- Vision Case Study
- Building Blocks

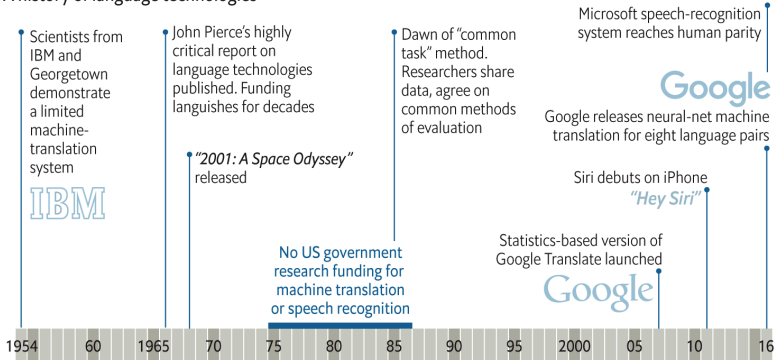
Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement

A history of language technologies



Antecedents


Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement



Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets


Application

- ConvNets
- LSTMs
- untapt
- Reinforcement

RNNs; *LSTM* RNNs

Hochreiter & Schmidhuber (1997)

Graves, ... & Schmidhuber (2009)



Vector Space Embedding

Word2Vec: Mikolov, ... & Dean (2013)

Antecedents


- Vision Case Study
- Building Blocks

Theory


- Neural Units
- Neural Nets
- Deep Neural Nets

Application


- ConvNets
- LSTMs
- untapt
- Reinforcement



Male-Female



Verb tense



Country-Capital



Antecedents


- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



Word2Vec + t-SNE

Antecedents


- Vision Case Study
- Building Blocks

Theory


- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



Word2Vec + t-SNE



'Understand' Language

with Word2Vec features in your model

Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement

```
model.most_similar(positive=['angular'])

[('angularjs', 0.9534549117088318),
 ('backbonejs', 0.9315043687820435),
 ('ember', 0.905410647392273),
 ('emberjs', 0.9029799103736877),
 ('reactjs', 0.896049439907074),
 ('requirejs', 0.8759748339653015),
 ('coffeescript', 0.8645504713058472),
 ('bootstrap', 0.8554328083992004),
 ('nodejs', 0.8515532612800598),
 ('backbone', 0.8443130254745483)]
```

```
model.most_similar(positive=['managed'])

[('oversaw', 0.8659406900405884),
 ('directed', 0.8491166234016418),
 ('supervised', 0.8058902621269226),
 ('coordinated', 0.7858685851097107),
 ('led', 0.7539615035057068),
 ('orchestrated', 0.7211644649505615),
 ('supported', 0.7198437452316284),
 ('comanaged', 0.6774874925613403),
 ('encompassing', 0.6726169586181641),
 ('administered', 0.6706464886665344)]
```

[even with small corpora]



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

Quick, Draw!

ConvNet + LSTM

[link]



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

- 1 Antecedents
 - Case Study: A History of Biological & Artificial Vision
 - Building Blocks
- 2 Theory
 - Biological & Artificial Neurons
 - Neural Networks
 - Deep Neural Networks
- 3 Contemporary Applications
 - Convolutional Neural Networks
 - Long Short-Term Memory Recurrent Neural Networks
 - Deep Learning at untapt**
 - Deep Reinforcement Learning



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

untapt

Digital Recruitment Platform



Antecedents

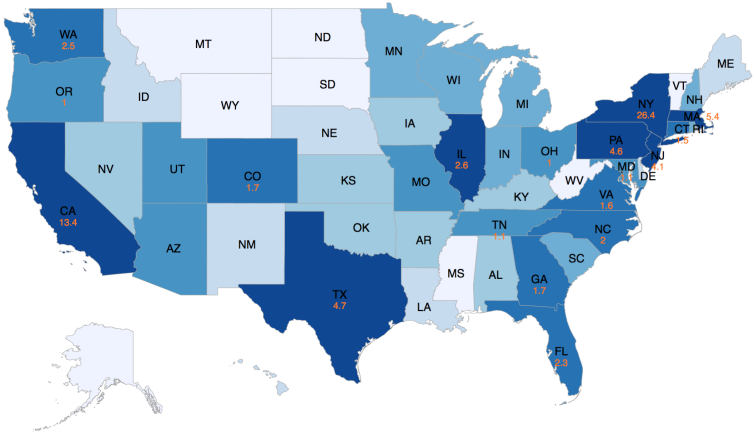
- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt**
- Reinforcement



Antecedents


- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt**
- Reinforcement



Antecedents

- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt**
- Reinforcement



Multi-Stage Bayesian Regression with PyMC3

Antecedents


- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt**
- Reinforcement



Krohn, Rives-Corbett & Donner (2016)



Antecedents


- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt**
- Reinforcement



untapt

Ensemble with Deep Neural Net

Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

```
Give me one bullet-point from your resume:  
>> • Sat around all day checking my Facebook feed  
I predict a 0.0% chance of interview
```

```
Give me one bullet-point from your resume:  
>> • Developed trading applications in Python  
I predict a 24.6% chance of interview
```

```
Give me one bullet-point from your resume:  
>> • Developed python solution for Monte Carlo risk calculation using numpy,  
    scipy and pandas, with a Javascript frontend in AngularJS and React  
I predict a 98.1% chance of interview
```



Antecedents


- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt**
- Reinforcement



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

- 1 Antecedents
 - Case Study: A History of Biological & Artificial Vision
 - Building Blocks
- 2 Theory
 - Biological & Artificial Neurons
 - Neural Networks
 - Deep Neural Networks
- 3 Contemporary Applications
 - Convolutional Neural Networks
 - Long Short-Term Memory Recurrent Neural Networks
 - Deep Learning at untapt
 - Deep Reinforcement Learning



AlphaGO

Silver et al. (2016)

Antecedents

Vision Case Study

Building Blocks

Theory

Neural Units

Neural Nets

Deep Neural Nets

Application

ConvNets

LSTMs

untapt

Reinforcement



Antecedents


- Vision Case Study
- Building Blocks

Theory

- Neural Units
- Neural Nets
- Deep Neural Nets

Application

- ConvNets
- LSTMs
- untapt
- Reinforcement



[Atari Games]



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

[OpenAI Universe]

[Google DeepMind Lab]



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

I'm Hiring

Full Stack Developer

Full Time · New York City

\$80k - \$120k · 0.1 - 0.25%

Software Sales Professional

Full Time · New York City, Boston

\$60k - \$90k · 0.25 - 0.75%

Content Marketer/Growth Hacker

Full Time · New York City

\$40k - \$60k · 0.01 - 0.5%

Data Engineer

Full Time · New York City

\$80k - \$120k · 0.1 - 0.25%

[Details at angel.co/untapt]



Antecedents

Vision Case Study
Building Blocks

Theory

Neural Units
Neural Nets
Deep Neural Nets

Application

ConvNets
LSTMs
untapt
Reinforcement

