




6.S094: Deep Learning for Self-Driving Cars

# Lecture 1: Introduction to Deep Learning and Self-Driving Cars

[cars.mit.edu](http://cars.mit.edu)



## 最专业报告分享群：

- 每日分享5+科技行业报告
- 同行业匹配，覆盖人工智能、大数据、机器人、智慧医疗、智能家居、物联网等行业。
- 高质量用户，同频的人说同样的话

扫描右侧二维码，  
或直接搜索关注公众号：智东西（zhidxcom）  
回复“**报告群**”加入



# Administrative



- **Website:** [cars.mit.edu](https://cars.mit.edu)
- **Contact Email:** [deepcars@mit.edu](mailto:deepcars@mit.edu)
- **Required:**
  - Create an account on the website.
  - Follow the tutorial for each of the 2 projects.
- **Recommended:**
  - Ask questions
  - Win competition!



Lex Fridman  
Instructor



Benedikt Jenik  
TA



William Angell  
TA



Spencer Dodd  
TA



Dan Brown  
TA

# Target Audience

## You may be:

- New to programming
- New to machine learning
- New to robotics

## What you will learn:

- An overview of deep learning methods:
  - Deep Reinforcement Learning
  - Convolutional Neural Networks
  - Recurrent Neural Networks
- How deep learning can help improve each component of autonomous driving: perception, localization, mapping, control, planning, driver state

# Target Audience

Not many equation slides like the following:

**Lemma 0.1.** *Let  $\mathcal{C}$  be a set of the construction.*

*Let  $\mathcal{C}$  be a gerber covering. Let  $\mathcal{F}$  be a quasi-coherent sheaves of  $\mathcal{O}$ -modules. We have to show that*

$$\mathcal{O}_{\mathcal{O}_X} = \mathcal{O}_X(\mathcal{L})$$

.

*Proof.* This is an algebraic space with the composition of sheaves  $\mathcal{F}$  on  $X_{\acute{e}tale}$  we have

$$\mathcal{O}_X(\mathcal{F}) = \{morph_1 \times_{\mathcal{O}_X} (\mathcal{G}, \mathcal{F})\}$$

where  $\mathcal{G}$  defines an isomorphism  $\mathcal{F} \rightarrow \mathcal{F}$  of  $\mathcal{O}$ -modules. □

\* Though it would be more efficient, since the above is LaTeX code automatically generated character by character with Recurrent Neural Networks (RNNs)

[35] Andrej Karpathy. "The Unreasonable Effectiveness of Recurrent Neural Networks." (2015).



# Guest Speakers



## Mapping, Localization, and the Challenge of Autonomous Driving

[John Leonard](#)

Professor, MIT



## Past, Present, and Future of Motion Planning in a Complex World

[Sertac Karaman](#)

Professor, MIT



## From Research to Reality: Testing Self-Driving Cars on Boston Public Roads

[Karl Iagnemma](#)

CEO, nuTonomy and Research Scientist, MIT



**TBD**

[Chris Gerdes](#)

Professor, Stanford

# Project: DeepTraffic

## DeepTraffic

Americans spend 8 billion hours stuck in traffic every year.

Deep neural networks can help!

```
1
2 //
3 // a few things don't have var in front of them - they update already
  existing variables the game needs
4 lanesSide = 1; //1;
5 patchesAhead = 10; //13;
6 patchesBehind = 0; //7;
7 trainIterations = 100000;
8
9 // begin from convnetjs example
10 var num_inputs = (lanesSide * 2 + 1) * (patchesAhead + patchesBehind);
11 var num_actions = 5;
12 var temporal_window = 3; //1 // amount of temporal memory. 0 = agent lives
   in-the-moment :)
13 var network_size = num_inputs * temporal_window + num_actions *</pre></div><div data-bbox="202 506 297 522" data-label="Text"><p>Apply Code/Reset Net</p></div><div data-bbox="310 506 403 522" data-label="Text"><p>Save Code/Net to File</p></div><div data-bbox="414 506 518 522" data-label="Text"><p>Load Code/Net from File</p></div><div data-bbox="529 506 654 522" data-label="Text"><p>Submit Model to Competition</p></div><div data-bbox="193 542 409 625" data-label="Figure"><img alt="A line graph showing the performance of the DeepTraffic model over time. The x-axis represents time in thousands of iterations (0k to 1k), and the y-axis represents a performance metric ranging from -1.5 to 2.18. The curve starts at approximately -1.5 at 0k iterations, rises sharply to about 1.5 by 0.1k iterations, and then continues to rise more gradually, reaching a plateau of approximately 2.18 by 0.5k iterations, where it remains until 1k iterations."/></div><div data-bbox="193 636 389 657" data-label="Text"><p>Start Evaluation Run</p></div><div data-bbox="199 670 384 687" data-label="Text"><p>Value Function Approximating Neural Network:</p></div><div data-bbox="199 687 464 823" data-label="Diagram"><img alt="A diagram illustrating the architecture of the Value Function Approximating Neural Network. It shows a sequence of layers: an input layer with 135 nodes (labeled 'input(135)'), followed by a fully connected layer with 10 nodes (labeled 'fc(10)'), a ReLU activation layer with 10 nodes (labeled 'relu(10)'), another fully connected layer with 5 nodes (labeled 'fc(5)'), and finally a regression layer with 5 nodes (labeled 'regression(5)'). The input layer is represented by a grid of blue squares, while the subsequent layers are represented by vertical columns of blue squares."/></div><div data-bbox="8 149 173 739" data-label="Image"><img alt="A screenshot of the DeepTraffic game interface. It shows a top-down view of a multi-lane highway with several cars. A red car is positioned in the bottom lane, moving towards the top. The interface includes a speedometer on the left showing 'Speed: 80 mph' and 'Cars Passed: 290'. At the bottom, there are controls for 'Road Overlay' (set to 'None') and 'Simulation Speed' (set to 'Normal')."/></div><div data-bbox="96 747 162 765" data-label="Text"><p>Road Overlay:</p></div><div data-bbox="102 770 162 787" data-label="Text"><p>None</p></div><div data-bbox="87 795 170 812" data-label="Text"><p>Simulation Speed:</p></div><div data-bbox="113 819 151 836" data-label="Text"><p>Normal</p></div><div data-bbox="0 942 121 993" data-label="Page-Footer"><p><img alt="MIT Logo" data-bbox="0 942 50 993"/> Massachusetts<br/>Institute of<br/>Technology</p></div><div data-bbox="548 954 714 991" data-label="Page-Footer"><p>Course 6.S094:<br/>Deep Learning for Self-Driving Cars</p></div><div data-bbox="732 954 826 991" data-label="Page-Footer"><p>Lex Fridman:<br/>fridman@mit.edu</p></div><div data-bbox="849 954 920 991" data-label="Page-Footer"><p>Website:<br/>cars.mit.edu</p></div><div data-bbox="948 954 990 991" data-label="Page-Footer"><p>January<br/>2017</p></div>
```

# Project: DeepTesla



Tesla Control  
(by Autopilot)




(Ground Truth)

Learned Control  
(by Deep Neural Network)



Red = Disagree    Green = Agree





# Schedule

Mon, Jan 9	Introduction to Deep Learning and Self Driving Cars
Tue, Jan 10	<b>Learning to Move:</b> Reinforcement Learning for Motion Planning
	<b>DeepTraffic:</b> Solving Traffic with Deep Reinforcement Learning
Wed, Jan 11	<b>Learning to Drive:</b> End-to-End Learning for the Full Driving Task
	<b>DeepTesla:</b> End-to-End Learning from Human and Autopilot Driving
Thu, Jan 12	<b>Karl Iagnemma:</b> From Research to Reality: Testing Self-Driving Cars on Boston Public Roads
Fri, Jan 13	<b>John Leonard:</b> Mapping, Localization, and the Challenge of Autonomous Driving
Tue, Jan 17	<b>Chris Gerdes:</b> TBD
Wed, Jan 18	<b>Sertac Karaman:</b> Past, Present, and Future of Motion Planning in a Complex World
Thu, Jan 19	<b>Learning to Share:</b> Driver State Sensing and Shared Autonomy
Fri, Jan 20	<b>Eric Daimler:</b> The Future of Artificial Intelligence Research and Development
	<b>Learning to Think:</b> The Road Ahead for Human-Centered Artificial Intelligence

# Defining (Artificial) Intelligence

March 25, 1996





## Special Purpose:


Can it achieve a well-defined finite set of goals?

## General Purpose:


Can it achieve poorly-defined unconstrained set of goals?




(1995)



(2002)



(2009)



- **Formal tasks:** Playing board games, card games. Solving puzzles, mathematical and logic problems.
- **Expert tasks:** Medical diagnosis, engineering, scheduling, computer hardware design.
- **Mundane tasks:** Everyday speech, written language, perception, walking, object manipulation.

# How Hard is Driving?


## Open Question:

Is driving closer to **chess** or to **everyday conversation**?






# Chess Pieces: Self-Driving Car Sensors



## External

1. Radar
2. Visible-light camera
3. LIDAR
4. Infrared camera
5. Stereo vision
6. GPS/IMU
7. CAN
8. Audio




## Internal

1. Visible-light camera
2. Infrared camera
3. Audio



# Chess Pieces: Self-Driving Car Tasks

- **Localization and Mapping:**  
Where am I?
- **Scene Understanding:**  
Where is everyone else?
- **Movement Planning:**  
How do I get from A to B?
- **Driver State:**  
What's the driver up to?



# DARPA Grand Challenge II (2006)



**Result:** Stanford's Stanley wins

# DARPA Urban Challenge (2007)




**Result:** CMU's Boss (Tartan Racing) wins

# Industry Takes on the Challenge



Waymo / Google Self-Driving Car



Tesla Autopilot




Uber



nuTonomy



# If Driving is a Conversation: How Hard is it to Pass the Turing Test?



1. **Natural language processing** to enable it to communicate successfully
2. **Knowledge representation** to store information provided before or during the interrogation
3. **Automated reasoning** to use the stored information to answer questions and to draw new conclusions



## Turing Test:

Can a computer be mistaken for a human more than 30% of the time?






# Neuron: Biological Inspiration for Computation





- **Neuron:** computational building block for the brain
- Human brain:
  - ~100-1,000 trillion synapses
- **(Artificial) Neuron:** computational building block for the “neural network”
- (Artificial) neural network:
  - ~1-10 billion synapses

**Human brains have ~10,000 computational power than computer brains.**

# Perceptron: Forward Pass




# Perceptron Algorithm



Provide training set of (input, output) pairs and run:

1. Initialize perceptron with random weights
2. For the inputs of an example in the training set, compute the Perceptron's output
3. If the output of the Perceptron does not match the output that is known to be correct for the example:
  1. If the output should have been 0 but was 1, decrease the weights that had an input of 1.
  2. If the output should have been 1 but was 0, increase the weights that had an input of 1.
4. Go to the next example in the training set and repeat steps 2-4 until the Perceptron makes no more mistakes

# Neural Networks are Amazing

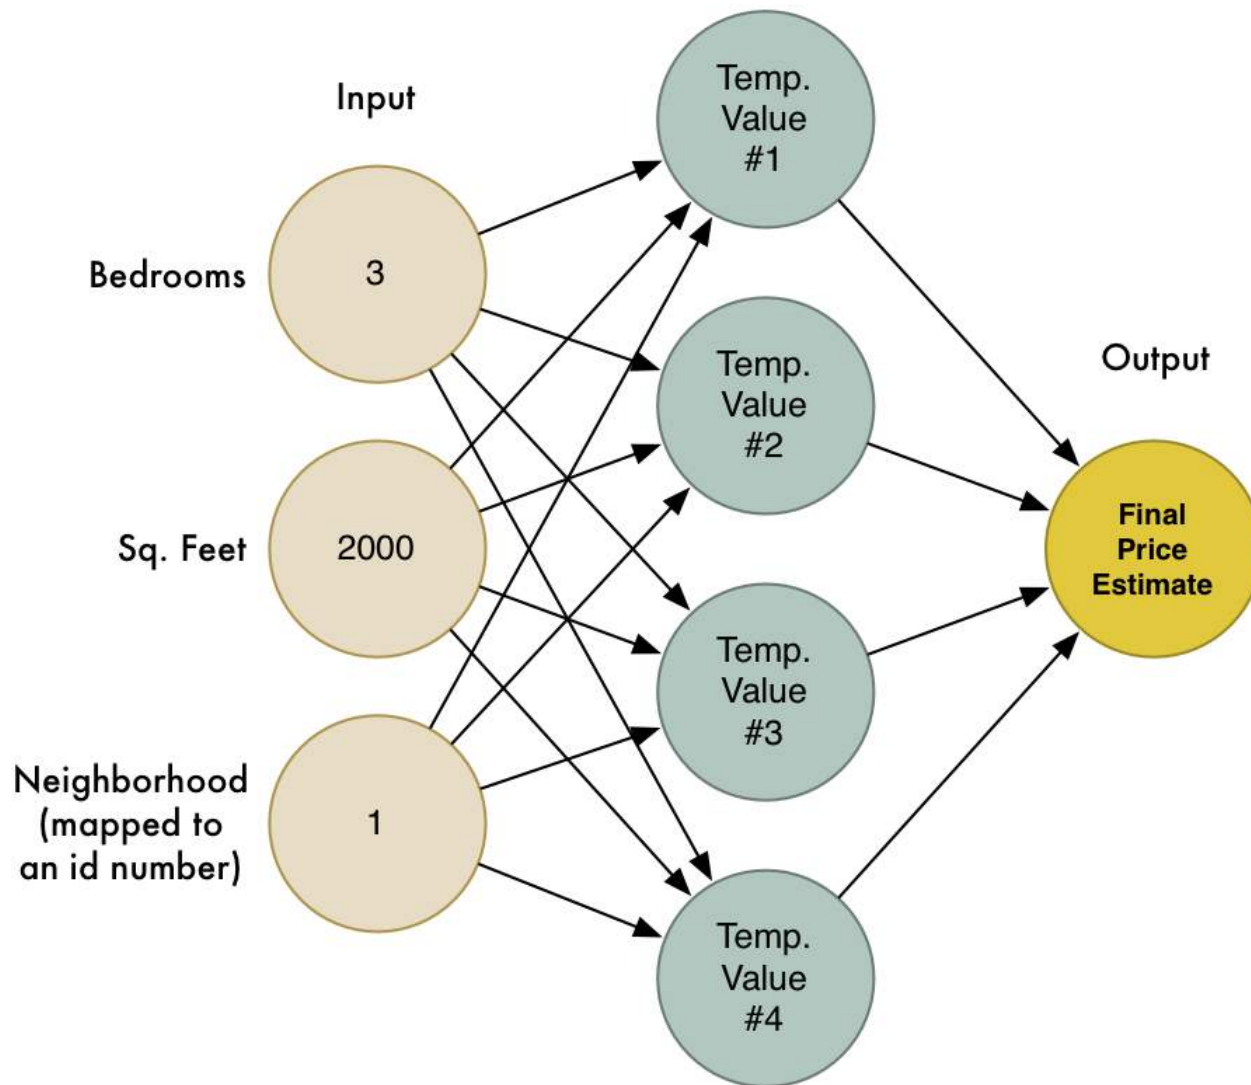


**Universality:** For any arbitrary function  $f(x)$ , there exists a neural network that closely approximate it for any input  $x$

Universality is an incredible property!\* And it holds for just 1 hidden layer.

\* Given that we have good algorithms for training these networks.

# Special Purpose Intelligence






# Neural Networks are Amazing: General Purpose Intelligence



Andrej Karpathy. “Deep Reinforcement Learning: Pong from Pixels.” 2016.



## Policy Network:



- 80x80 image (difference image)
- 2 actions: up or down
- 200,000 Pong games

**This is a step towards general purpose artificial intelligence!**

# Neural Networks are Amazing: General Purpose Intelligence



- Every (state, action) pair is **rewarded** when the final result is a **win**.
- Every (state, action) pair is **punished** when the final result is a **loss**.

## The fact that this works at all is amazing!

It could be called “general intelligence” but not yet “human-level” intelligence...

# Current Drawbacks

- Lacks Reasoning:
  - Humans only need simple instructions:  
“You’re in **control** of a paddle and you can move it up and down, and your task is to bounce the ball past the other player controlled by AI.”
- Requires **big** data: inefficient at learning from data
- Requires **supervised** data: costly to annotate real-world data
- Need to manually select network structure
- Needs hyperparameter tuning for training:
  - Learning rate
  - Loss function
  - Mini-batch size
  - Number of training iterations
  - Momentum: gradient update smoothing
  - Optimizer selection
- Defining a good reward function is difficult...

# Current Drawbacks


Defining a good reward function is difficult... **Coast Runners:** Discovers local pockets of high reward ignoring the “implied” bigger picture goal of finishing the race.



In addition, specifying a reward function for self-driving cars raises ethical questions...

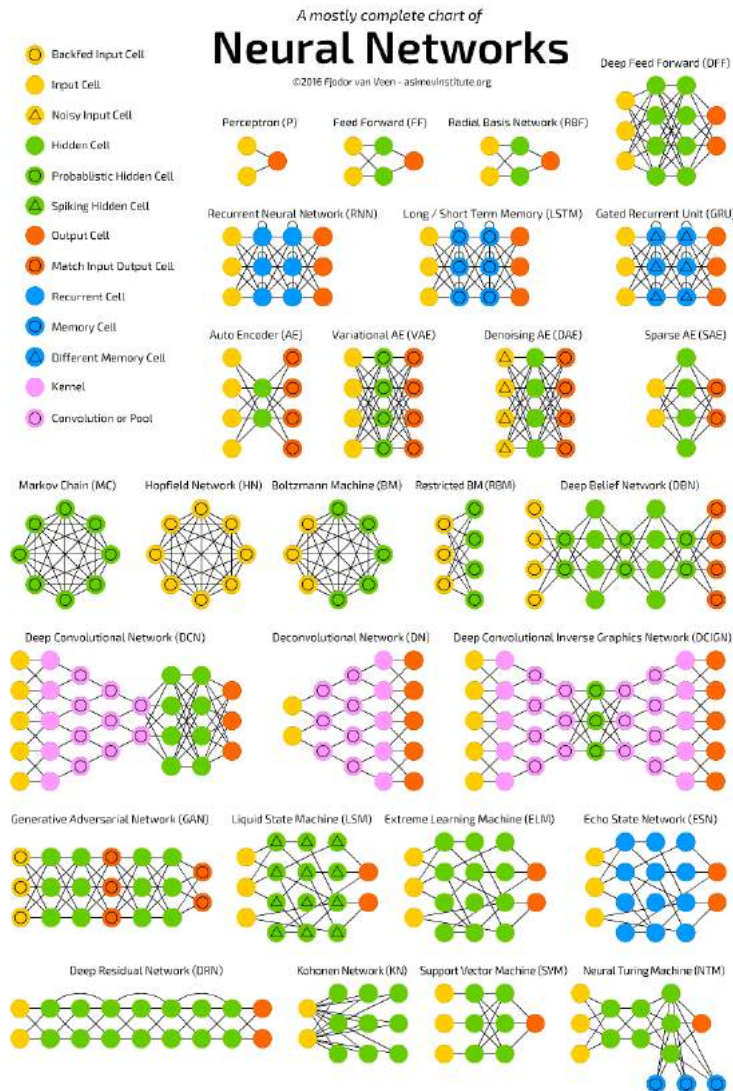
# Deep Learning Breakthroughs: What Changed?

- **Compute**  
CPUs, GPUs, ASICs
- **Organized large(-ish) datasets**  
Imagenet
- **Algorithms and research:**  
Backprop, CNN, LSTM
- **Software and Infrastructure**  
Git, ROS, PR2, AWS, Amazon Mechanical Turk, TensorFlow, ...
- **Financial backing of large companies**  
Google, Facebook, Amazon, ...





# Useful Deep Learning Terms



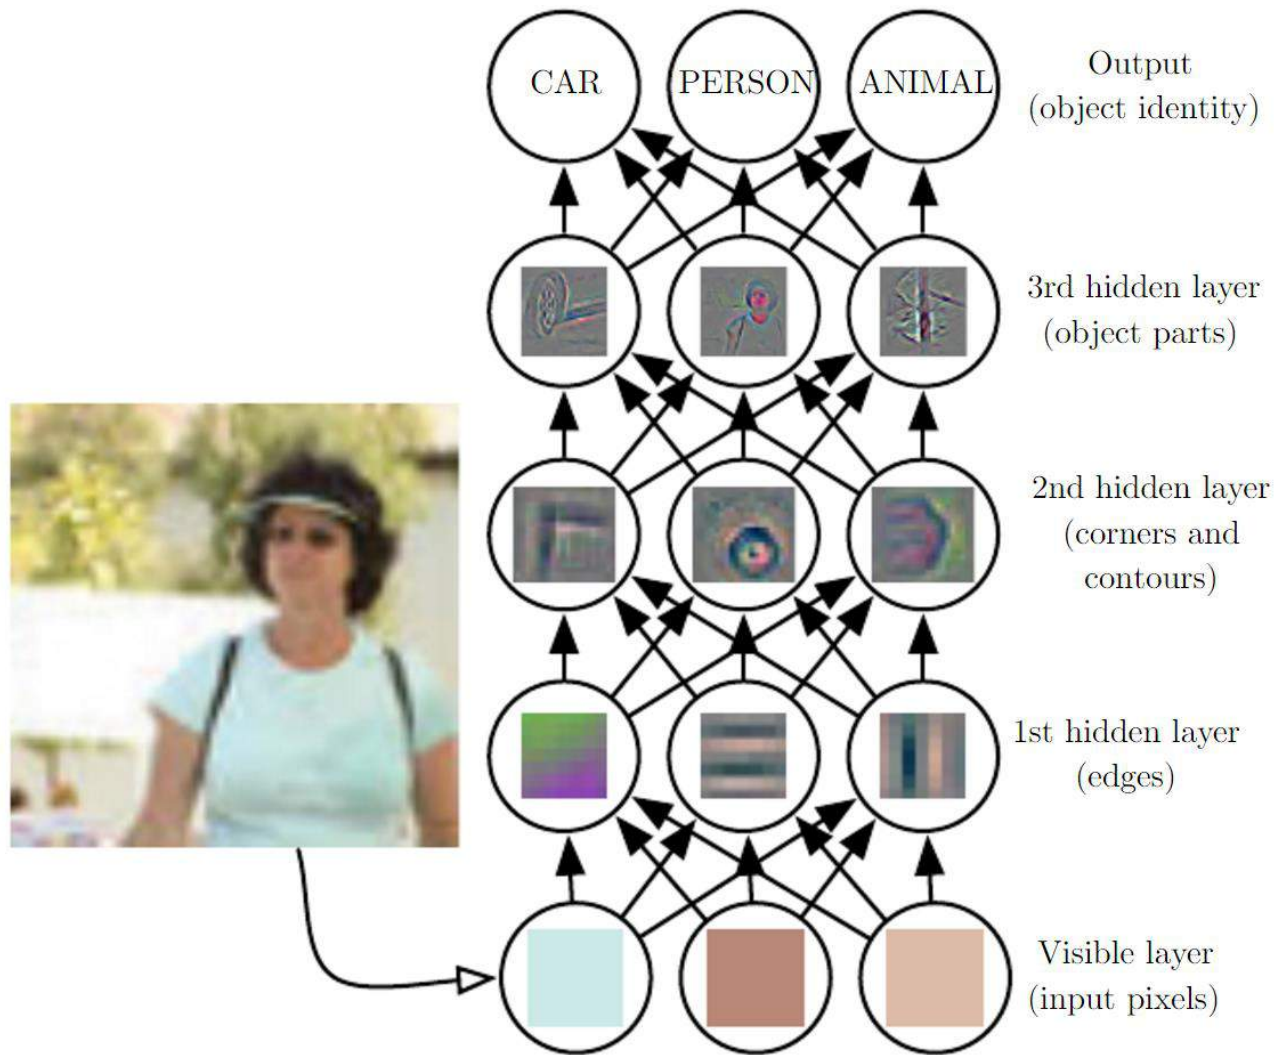
- Basic terms:
  - **Deep Learning = Neural Networks**
  - **Deep Learning** is a subset of **Machine Learning**
- Terms for neural networks:
  - **MLP**: Multilayer Perceptron
  - **DNN**: Deep neural networks
  - **RNN**: Recurrent neural networks
    - **LSTM**: Long Short-Term Memory
  - **CNN or ConvNet**: Convolutional neural networks
  - **DBN**: Deep Belief Networks
- Neural network operations:
  - Convolution
  - Pooling
  - Activation function
  - Backpropagation

Asimov Institute. "A mostly complete chart of neural networks."

# Neural Networks: Proceed with Caution



# Deep Learning is Representation Learning




[20] Goodfellow et al. "Deep learning." (2017).




# Representation Matters

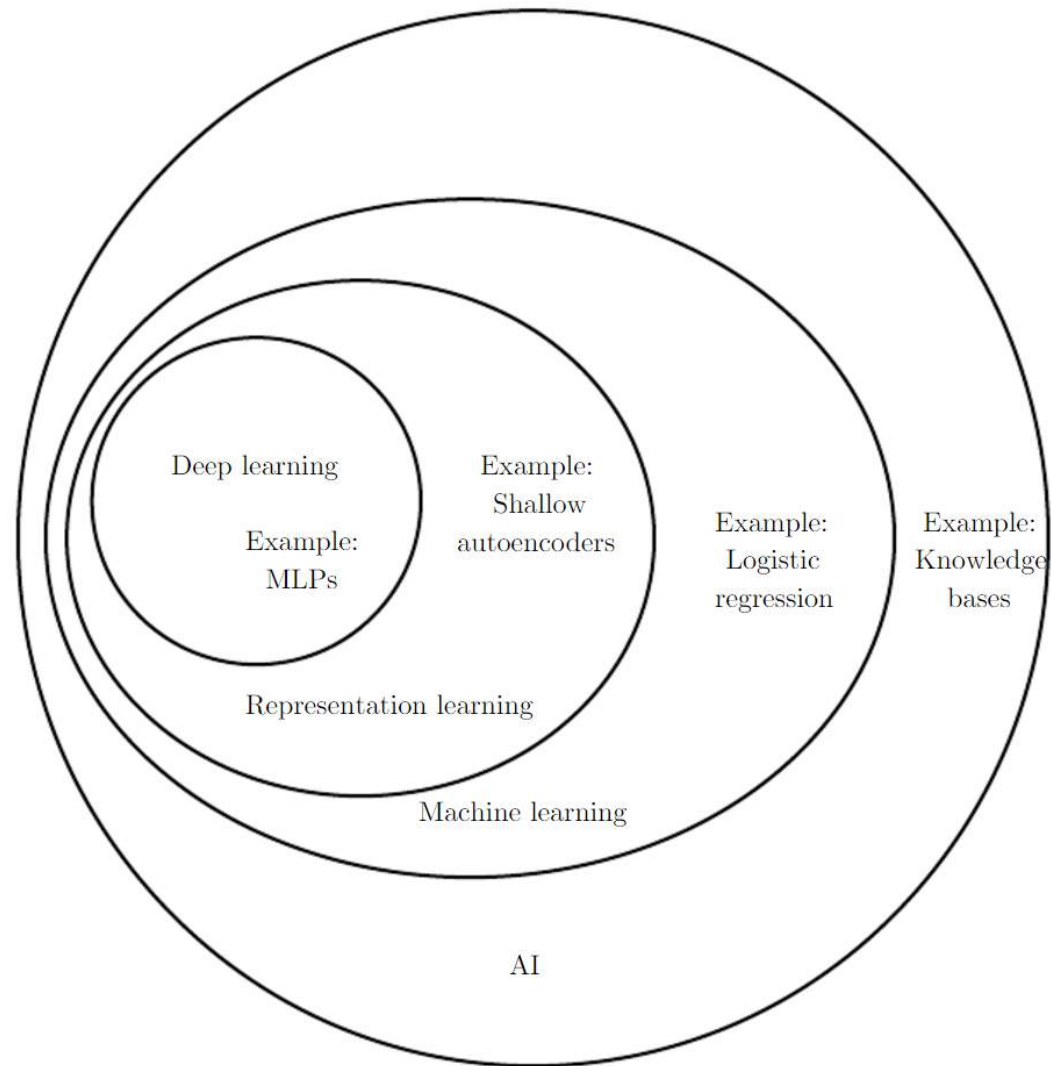
Cartesian coordinates



Polar coordinates

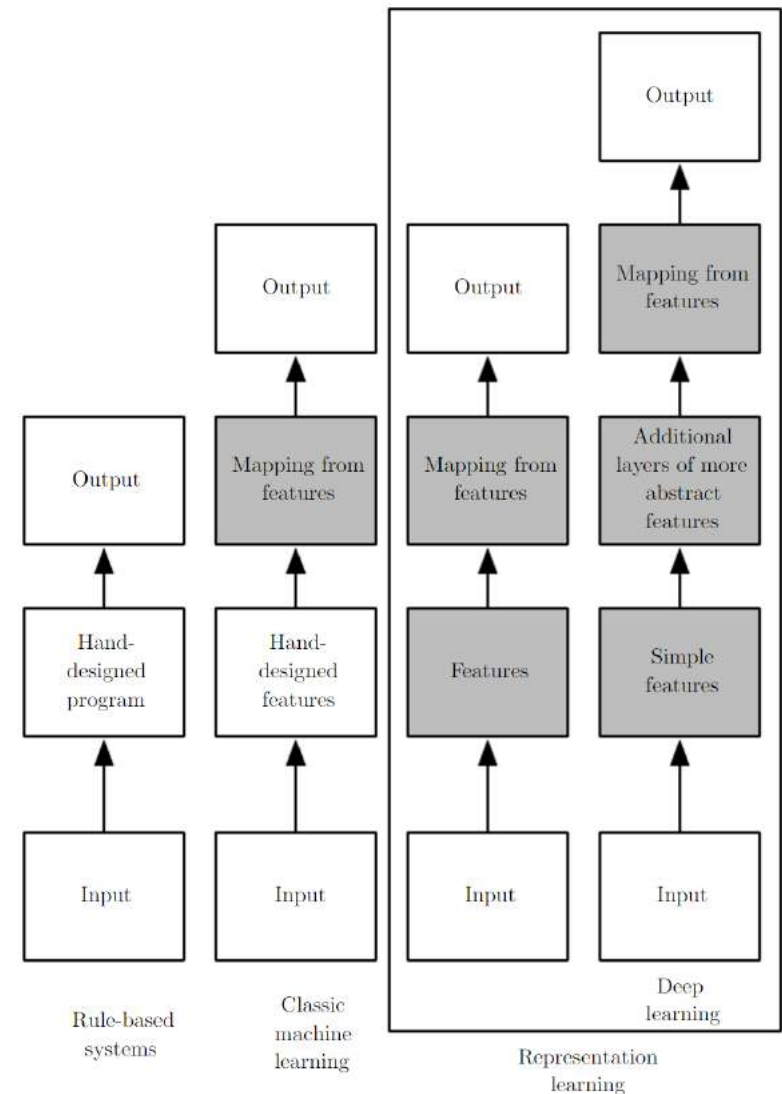



# Deep Learning is **Representation Learning**





# Deep Learning: **Scalable** Machine Learning



# Applications: Object Classification in Images




**mite**

**container ship**

**motor scooter**

**leopard**


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



Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. **"Imagenet classification with deep convolutional neural networks."** Advances in neural information processing systems. 2012.

# *Pause:* Progress on ImageNet

**ImageNet Classification Error (Top 5)**



# Computer Vision is Hard: Illumination Variability





# Computer Vision is Hard: Pose Variability and Occlusions

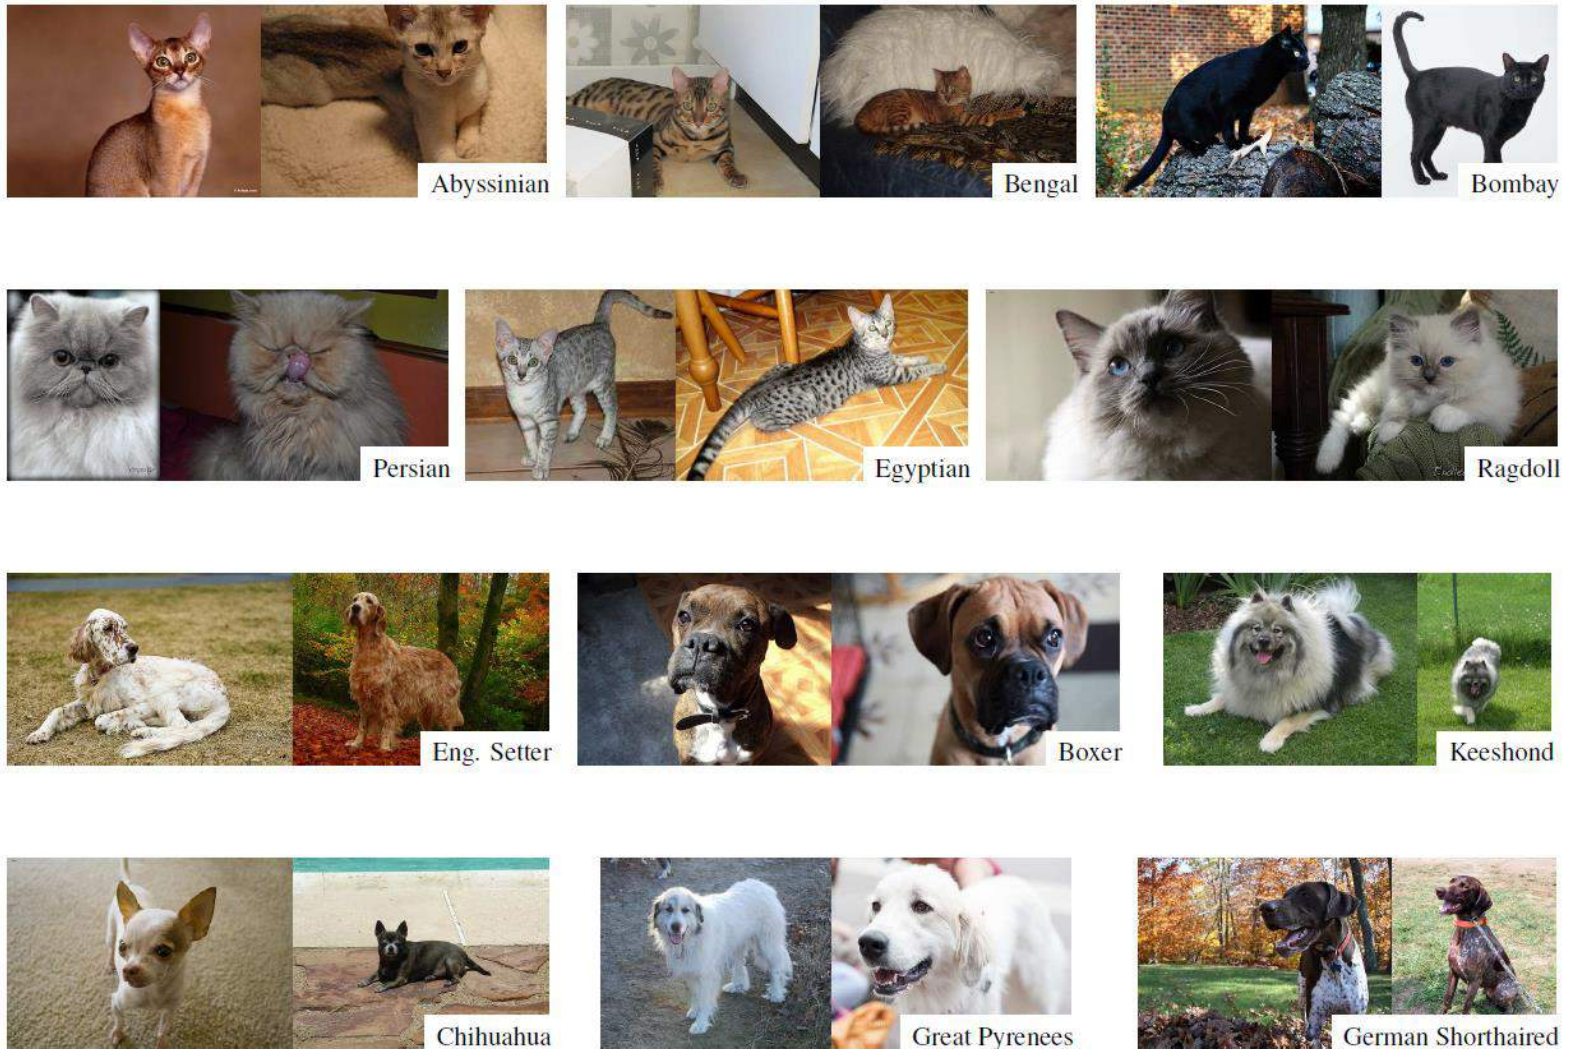


Figure 1. **The deformable and truncated cat.** Cats exhibit (al-

Parkhi et al. "The truth about cats and dogs." 2011.

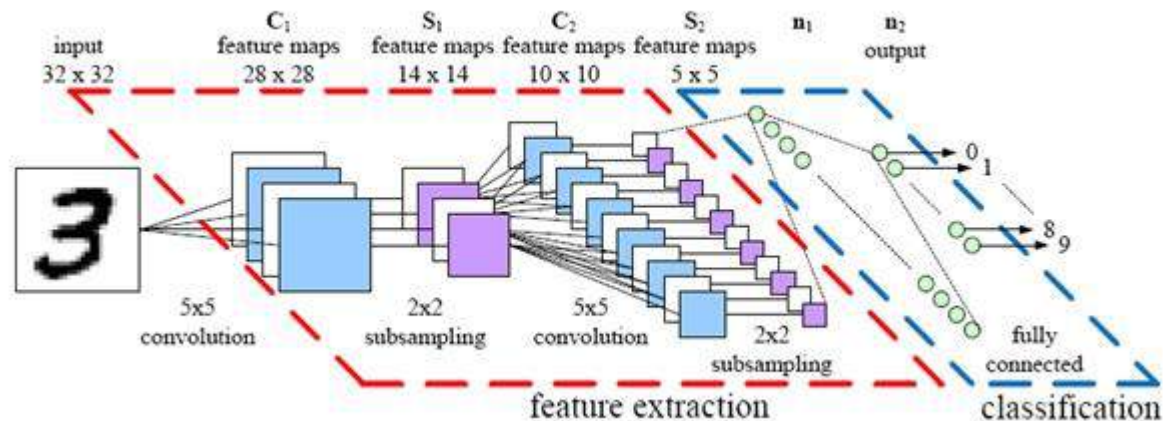


# Computer Vision is Hard: Intra-Class Variability

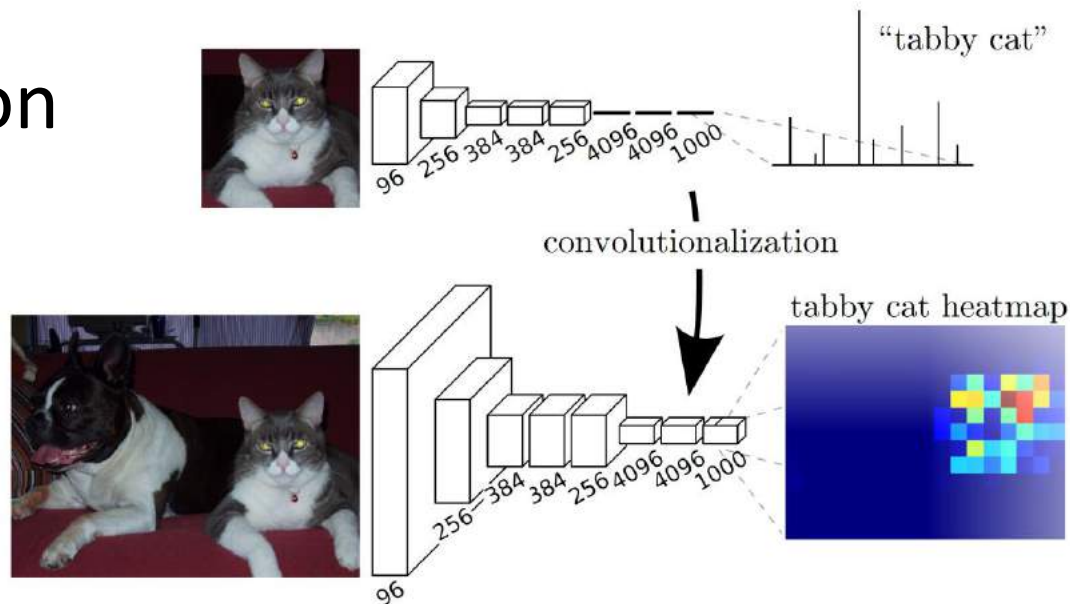


Parkhi et al. "Cats and dogs." 2012.

# *Pause:* Object Recognition / Classification




# *Pause:* Segmentation




Original



Ground Truth



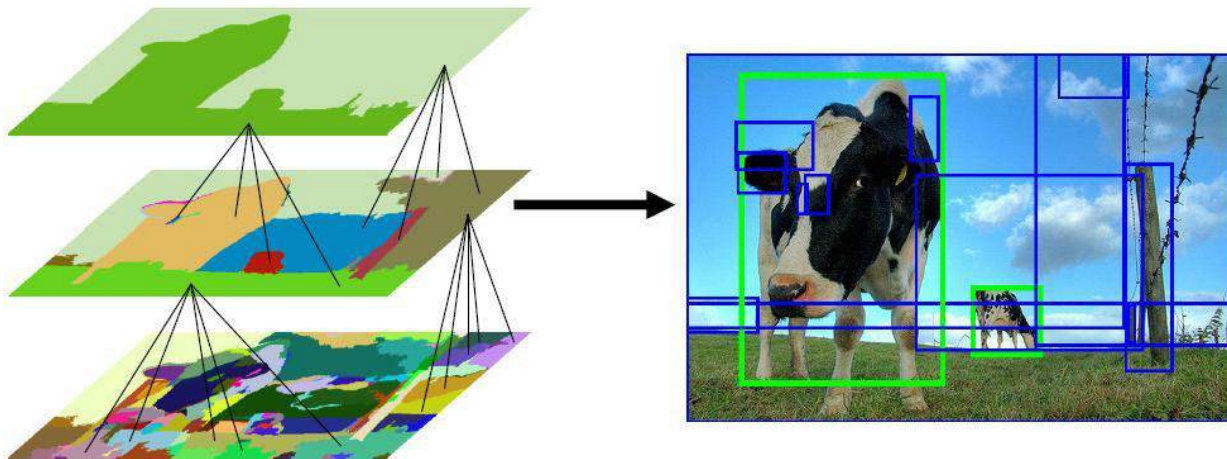
FCN-8




Source: Long et al. Fully Convolutional Networks for Semantic Segmentation. CVPR 2015.




# *Pause:* Object Detection



## **R-CNN: *Regions with CNN features***




1. Input image



2. Extract region proposals (~2k)

warped region



3. Compute CNN features

aeroplane? no.

⋮

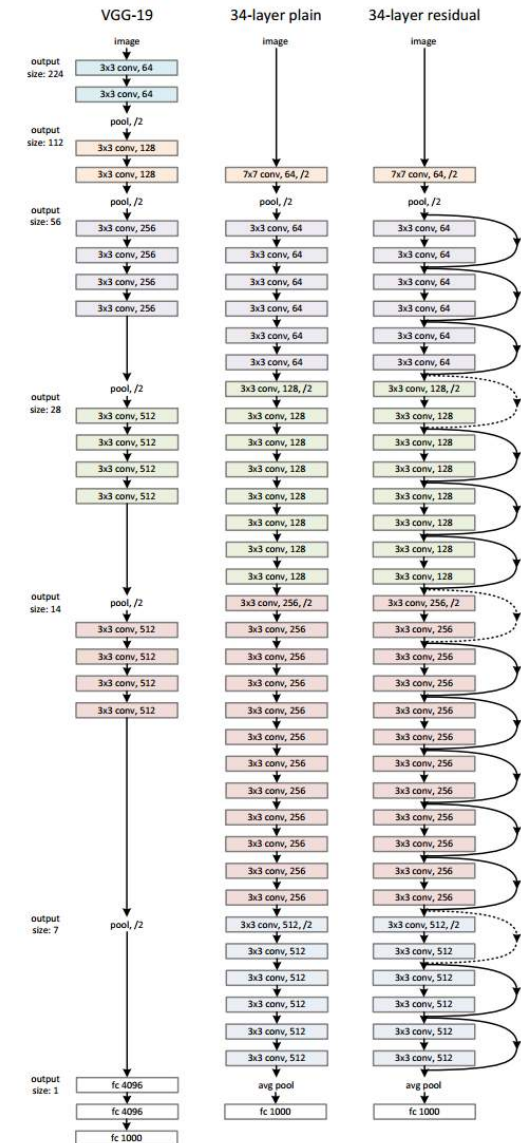

person? yes.

⋮

tvmonitor? no.

4. Classify regions


# Applications: Object Detection and Localization in Images



He Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Su.  
"Deep residual learning for image recognition." (2015).




# Applications: Colorization of Images



Zhang, Richard, Phillip Isola, and Alexei A. Efros. "Colorful Image Colorization." (2016).


# Applications: Automatic Translation of Text in Images




Google Translate

# (Pause...) Flavors of Neural Networks


one to one




one to many




many to one



many to many



many to many



↑  
“Vanilla”  
Neural  
Networks

Recurrent Neural Networks

Andrej Karpathy. **“The Unreasonable Effectiveness of Recurrent Neural Networks.”** (2015).

# Applications: Handwriting Generation from Text


**Input:**

**Text** --- up to 100 characters, lower case letters work best

Deep Learning for Self Driving Cars

**Output:**


Deep Learning  
for Self-Driving Cars



Alex Graves. **"Generating sequences with recurrent neural networks."** (2013).

# Applications: Character-Level Text Generation

Naturalism and decision for the majority of Arab countries' capitalide was grounded by the Irish language by [[John Clair]], [[An Imperial Japanese Revolt]], associated with Guangzham's sovereignty. His generals were the powerful ruler of the Portugal in the [[Protestant Immineners]], which could be said to be directly in Cantonese Communication, which followed a ceremony and set inspired prison, training.



Andrej Karpathy. "The Unreasonable Effectiveness of Recurrent Neural Networks." (2015).

Code: <https://github.com/karpathy/char-rnn>



# Applications: Character-Level Text Generation

Life Is About The Weather!

Life Is About The (Wild) Truth About Human-Rights

Life Is About The True Love Of Mr. Mom

Life Is About Where He Were Now

Life Is About Kids


Life Is About What It Takes If Being On The Spot Is Tough

Life Is About... An Eating Story

Life Is About The Truth Now

The meaning of life is literary recognition.

The meaning of life is the tradition of the ancient human reproduction



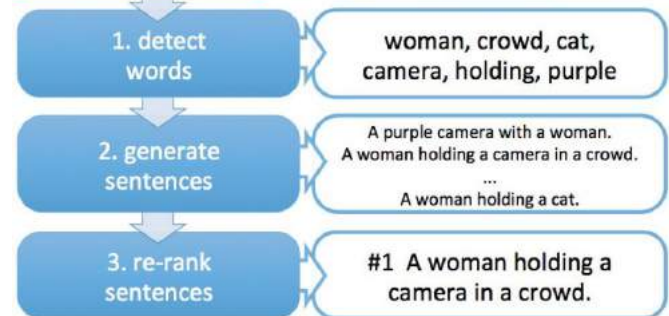

Andrej Karpathy. "The Unreasonable Effectiveness of Recurrent Neural Networks." (2015).

Code: <https://github.com/karpathy/char-rnn>

# Applications: Image Caption Generation



a man sitting on a couch with a dog  
a man sitting on a chair with a dog in his lap



# Applications: Image Question Answering



COCOQA 33827

**What is the color of the cat?**

Ground truth: black

IMG+BOW: **black (0.55)**

2-VIS+LSTM: **black (0.73)**

BOW: **gray (0.40)**

COCOQA 33827a


**What is the color of the couch?**

Ground truth: red

IMG+BOW: **red (0.65)**

2-VIS+LSTM: **black (0.44)**

BOW: **red (0.39)**



DAQUAR 1522

**How many chairs are there?**

Ground truth: two

IMG+BOW: **four (0.24)**

2-VIS+BLSTM: **one (0.29)**

LSTM: **four (0.19)**

DAQUAR 1520

**How many shelves are there?**

Ground truth: three

IMG+BOW: **three (0.25)**

2-VIS+BLSTM: **two (0.48)**

LSTM: **two (0.21)**



COCOQA 14855

**Where are the ripe bananas sitting?**

Ground truth: basket

IMG+BOW: **basket (0.97)**

2-VIS+BLSTM: **basket (0.58)**

BOW: **bowl (0.48)**

COCOQA 14855a

**What are in the basket?**

Ground truth: bananas

IMG+BOW: **bananas (0.98)**

2-VIS+BLSTM: **bananas (0.68)**

BOW: **bananas (0.14)**



DAQUAR 585

**What is the object on the chair?**

Ground truth: pillow

IMG+BOW: **clothes (0.37)**

2-VIS+BLSTM: **pillow (0.65)**

LSTM: **clothes (0.40)**

DAQUAR 585a

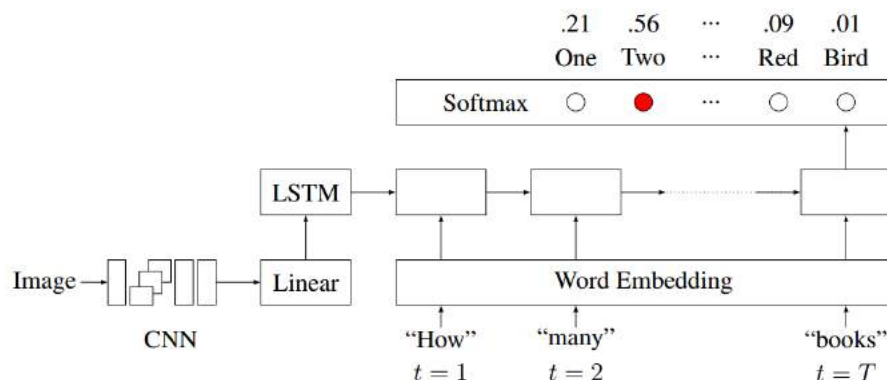
**Where is the pillow found?**

Ground truth: chair

IMG+BOW: **bed (0.13)**

2-VIS+BLSTM: **chair (0.17)**

LSTM: **cabinet (0.79)**



Ren et al. "Exploring models and data for image question answering." 2015.

Code: <https://github.com/renmengye/imageqa-public>




# Applications: Video Description Generation

## Correct descriptions.



S2VT: A man is doing stunts on his bike.




S2VT: A herd of zebras are walking in a field.

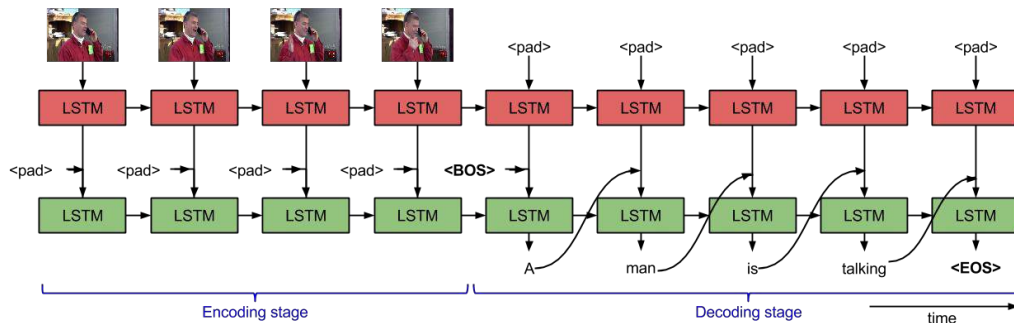
## Relevant but incorrect descriptions.



S2VT: A small bus is running into a building.



S2VT: A man is cutting a piece of a pair of a paper.





Venugopalan et al.

"Sequence to sequence-video to text." 2015.

Code: <https://vsubhashini.github.io/s2vt.html>

# Applications: Modeling Attention Steering




Jimmy Ba, Volodymyr Mnih, and Koray Kavukcuoglu. **"Multiple object recognition with visual attention."** (2014).




# Applications: Drawing with Selective Attention

## Reading



## Writing




Gregor et al. "DRAW: A recurrent neural network for image generation." (2015).

Code: <https://github.com/ericjang/draw>

# Applications: Adding Audio to Silent Film

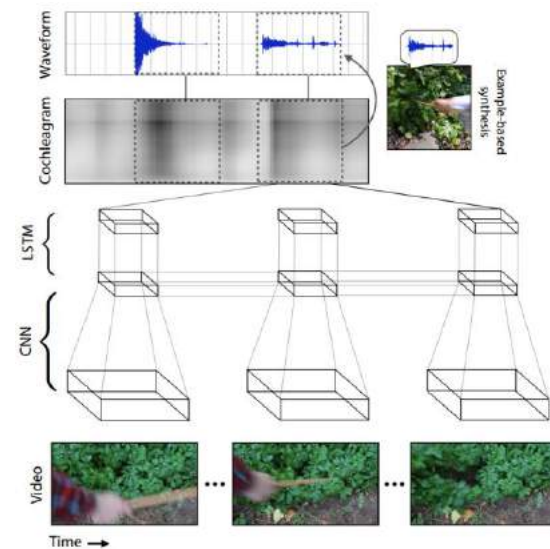


Silent video



Predicted soundtrack

Owens, Andrew, Phillip Isola, Josh McDermott, Antonio Torralba, Edward H. Adelson, and William T. Freeman. "**Visually Indicated Sounds.**" (2015).






# Moravec's Paradox: The “Easy” Problems are Hard



**Soccer is harder than Chess**



# Moravec's Paradox: The “Easy” Problems are Hard



# Question: Why?

# Answer: Data

**Visual perception:** 540 millions years of data

**Bipedal movement:** 230+ million years of data

**Abstract thought:** 100 thousand years of data


“Encoded in the large, highly evolved sensory and motor portions of the human brain is a **billion years of experience** about the nature of the world and how to survive in it.... Abstract thought, though, is a new trick, perhaps less than **100 thousand years** old. We have not yet mastered it. It is not all that intrinsically difficult; it just seems so when we do it.”  
- *Hans Moravec, Mind Children (1988)*



Hans Moravec (CMU)



Rodney Brooks (MIT)



Marvin Minsky (MIT)



# Walking is Hard. How Hard is Driving?

**Human performance:** 1 fatality per 100,000,000 miles

**Error rate for AI to improve on:** **0.000001%**

## Challenges:

- Snow
- Heavy rain
- Big open parking lots
- Parking garages
- Any pedestrian behaving irresponsibly or just unpredictably
- Reflections, dynamics blinding ones
- Merging into a high-speed stream of oncoming traffic

# Google Self-Driving Car: Driver Disengagements

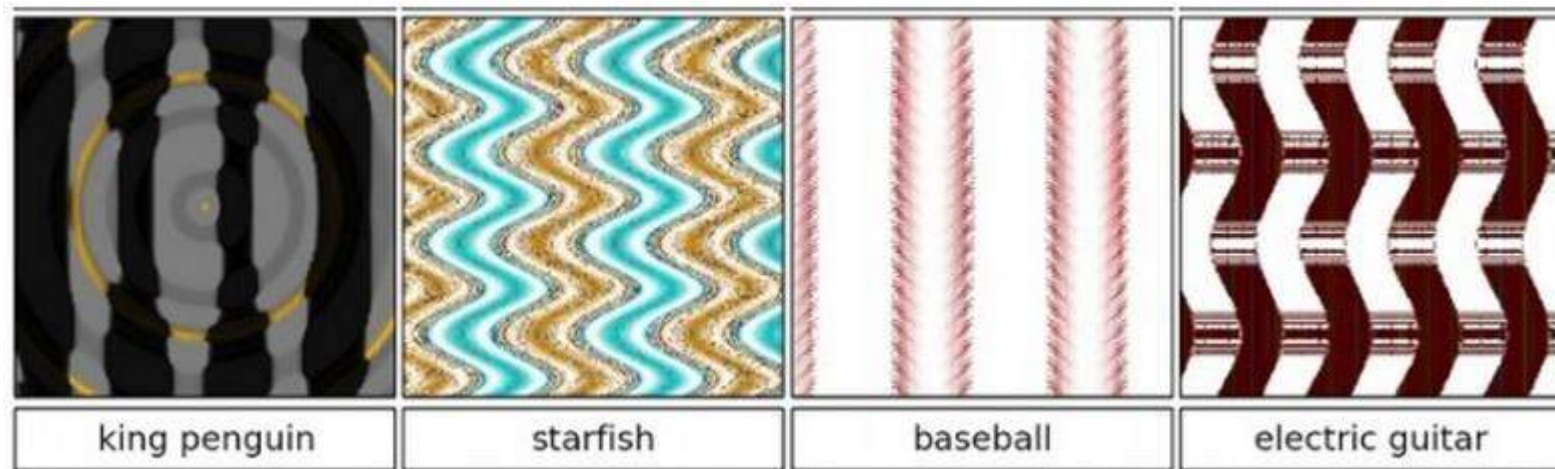

Month	Number Disengages	Autonomous miles on public roads
2014/09	2	4207.2
2014/10	19	23971.1
2014/11	21	15836.6
2014/12	43	9413.1
2015/01	53	18192.1
2015/02	14	18745.1
2015/03	30	22204.2
2015/04	51	31927.3
2015/05	13	38016.8
2015/06	11	42046.6
2015/07	29	34805.1
2015/08	7	38219.8
2015/09	16	36326.6
2015/10	16	47143.5
2015/11	16	43275.9
<b>Total</b>	341	424331

# Google Self-Driving Car: Driver Disengagements

Cause	Sep 2014	Oct 2014	Nov 2014	Dec 2014	Jan 2015	Feb 2015	Mar 2015	Apr 2015	May 2015	Jun 2015	Jul 2015	Aug 2015	Sep 2015	Oct 2015	Nov 2015	Total
disengage for weather conditions during testing	0	0	0	0	1	5	0	6	0	0	0	0	0	0	1	13
disengage for a recklessly behaving road user	1	0	1	1	1	3	3	7	0	0	0	2	1	0	3	23
disengage for hardware discrepancy	0	1	0	0	2	1	0	1	0	5	8	1	8	8	4	39
disengage for unwanted maneuver of the vehicle	0	3	6	14	15	1	3	2	1	0	3	2	0	3	2	55
disengage for a perception discrepancy	1	2	3	18	19	2	20	30	4	4	8	0	4	3	1	119
disengage for incorrect behavior prediction of other traffic participants	0	2	2	0	1	0	2	0	0	0	0	0	0	1	0	8
disengage for a software discrepancy	0	11	9	9	14	2	1	5	8	2	9	2	3	1	4	80
disengage for construction zone during testing	0	0	0	1	0	0	1	0	0	0	1	0	0	0	0	3
disengage for emergency vehicle during testing	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1
<b>Total</b>	<b>2</b>	<b>19</b>	<b>21</b>	<b>43</b>	<b>53</b>	<b>14</b>	<b>30</b>	<b>51</b>	<b>13</b>	<b>11</b>	<b>29</b>	<b>7</b>	<b>16</b>	<b>16</b>	<b>16</b>	<b>341</b>

# Robustness:

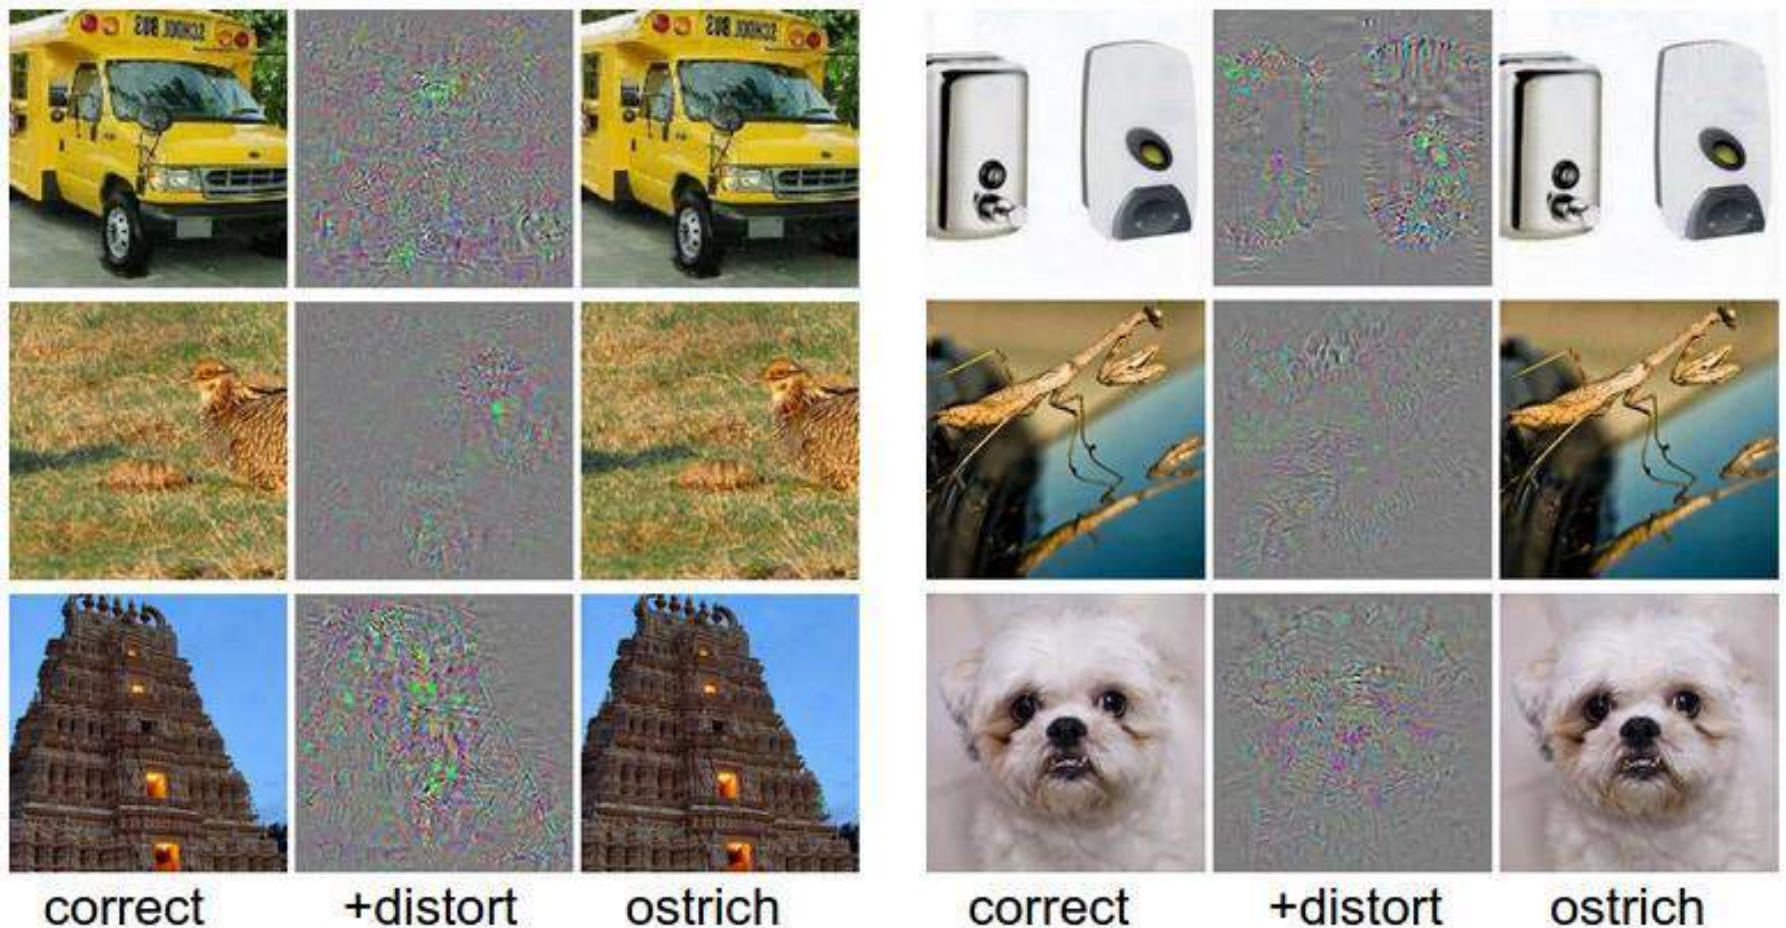
## >99.6% Confidence in the Wrong Answer



Nguyen et al. "Deep neural networks are easily fooled: High confidence predictions for unrecognizable images." 2015.




# Robustness: Fooled by a Little Distortion



Szegedy et al. "Intriguing properties of neural networks." 2013.



# Mark I Perceptron



- Frank Rosenblatt
- 400 pixel image input
- Weights encoded in potentiometers
- Weight updated by electric motors

## The New York Times

### ***NEW NAVY DEVICE LEARNS BY DOING***

*July 8, 1958*

“The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence... Dr. Frank Rosenblatt, a research psychologist at the Cornell Aeronautical Laboratory, Buffalo, said Perceptrons might be fired to the planets as mechanical space explorers”

# AI Winters

Two major episodes:

- 1974-80
- 1987-93


Smaller episodes:

- 1966: the failure of machine translation
- 1970: the abandonment of connectionism
- 1971-75: DARPA's frustration with the Speech Understanding Research program
- 1973: the large decrease in AI research in the UK in response to the Lighthill report.
- 1973–74: DARPA's cutbacks to academic AI research in general
- 1987: the collapse of the Lisp machine market
- 1988: the cancellation of new spending on AI by the Strategic Computing Initiative
- 1993: expert systems slowly reaching the bottom
- 1990s: the quiet disappearance of the fifth-generation computer project's original goals.

“In no part of the field have discoveries made so far produced the major impact that was then promised.”


# The Seasons of Deep Learning

- 1940s-1960s: Cybernetics
  - Biological learning (1943)
  - Perceptron (1958)
- 1980s-1990s: Connectionism
  - Backpropagation
- 2006-: Deep Learning



# 3<sup>rd</sup> Summer of Deep Learning

Interest over time ?



Google trends: “Deep Learning”

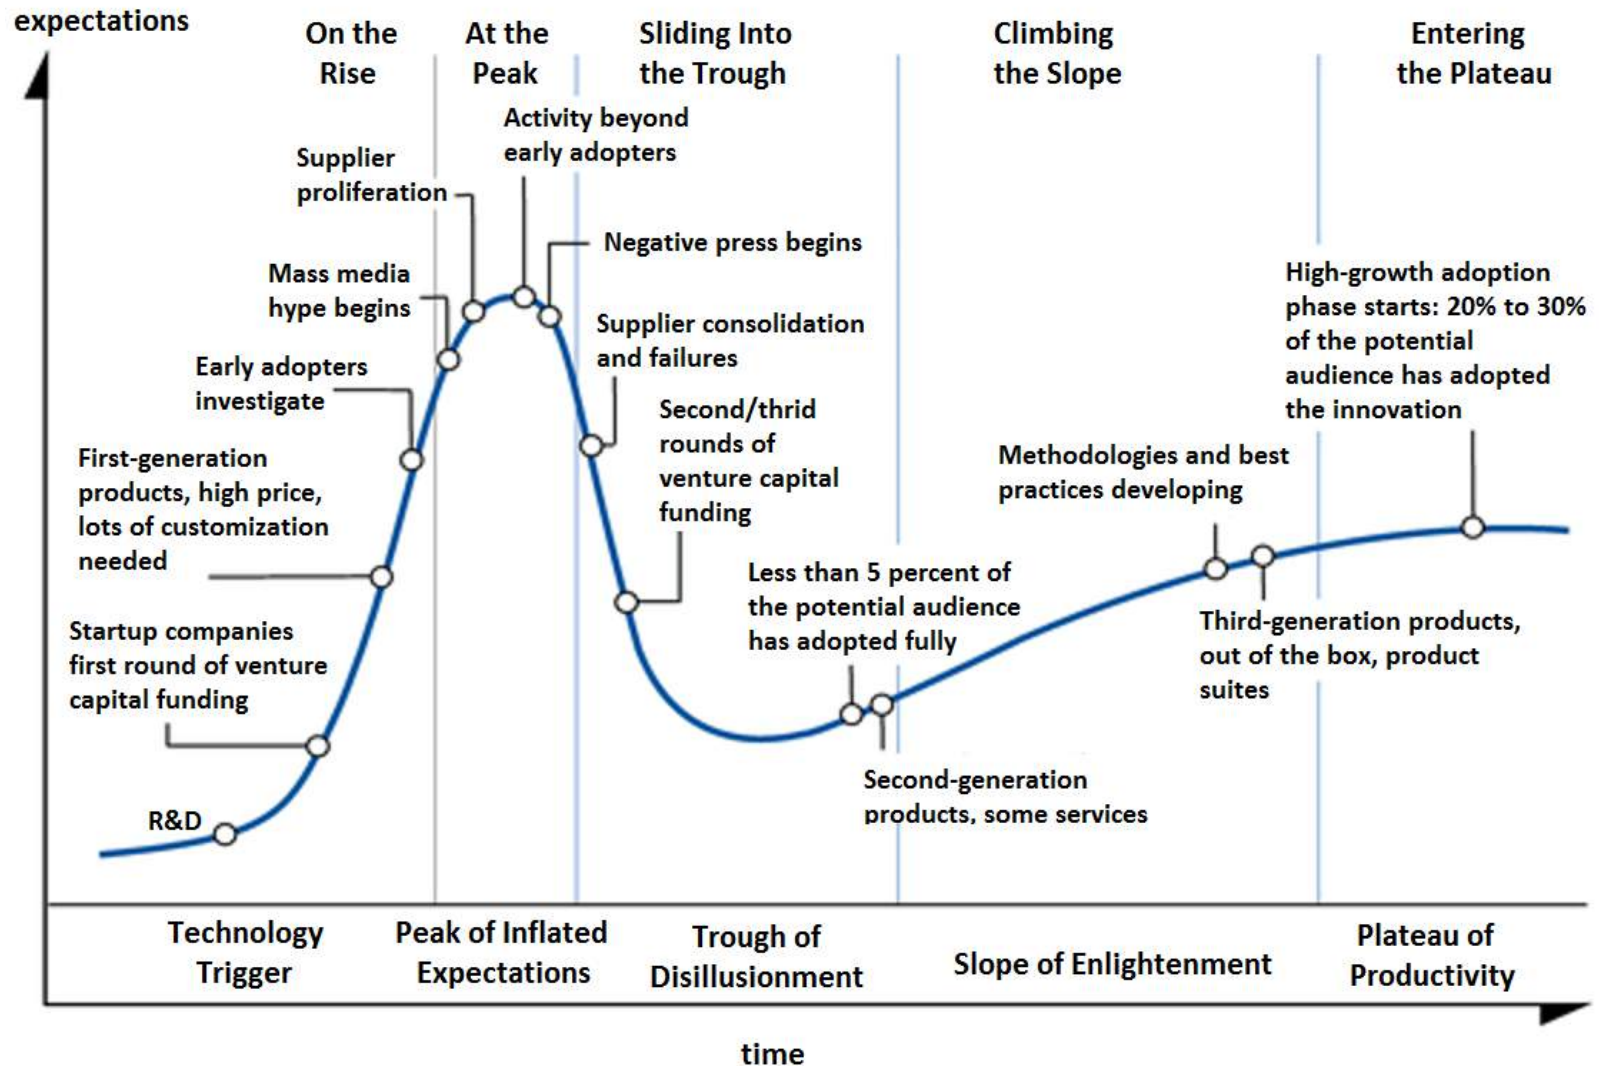


# Proceed with Caution: What's Next for Deep Learning?


(5 year vision)

- **Ilya Sutskever, Research Director of OpenAI:**  
**Deeper models**, models that need fewer examples for training.
- **Christian Szegedy, Senior Research Scientist at Google:**  
Become so efficient that they will be able to **run on cheap mobile devices**.
- **Pieter Abbeel, Associate Professor in Computer Science at UC Berkeley:**  
Significant advances in deep unsupervised learning and deep **reinforcement learning**.
- **Ian Goodfellow, Senior Research Scientist at Google:**  
Neural networks that **can summarize what happens in a video clip**, and will be able to generate short videos. Neural networks that model the behavior of genes, drugs, and proteins and then used to design new medicines.
- **Koray Kavukcuoglu & Alex Graves, Research Scientists at Google DeepMind:**  
An increase in **multimodal learning**, and a stronger focus on learning that persists **beyond individual datasets**.
- **Charlie Tang, Machine Learning group, University of Toronto:**  
Deep learning algorithms ported to **commercial products**, much like how the face detector was incorporated into consumer cameras in the past 10 years.

# Gartner Hype Cycle



## New-vehicle market share of autonomous vehicles, %



### Factors in disruption scenarios

Regulatory challenges  
Safe, reliable technical solutions  
Consumer acceptance, willingness to pay


### High disruption

Fast  
Comprehensive  
Enthusiastic

### Low disruption

Gradual  
Incomplete  
Limited


# Sensor modules market value for autonomous cars from 2015 to 2030 (in \$B)





# Attention to (AI) Drivers: Proceed with Caution

## Camera Spoofing



correct

+distort

ostrich

## LIDAR Spoofing






- Interface: Python, (C++)
- Automatic Differentiation
- Multi GPU, Cluster Support
- Currently most popular



# Keras



- On top of Tensorflow (and Theano)
  - Interface: Python
  - Goal: provide a simplified interface
- 
- **Also:** TF Learn, TF Slim



# Torch

- Used by researchers doing lower level (closer to the details) neural net work
- Interface: Lua
- Fragmented across different plugins

# facebook



# theano

- Interface: Python (tight NumPy integration)
- One of the earlier frameworks with GPU support
- Encourages low-level tinkering



# cuDNN



- The library that most frameworks use for doing the actual computation
- Implements primitive neural network functions in CUDA on the GPU




- Multi GPU Support (scales well)
- Interface: Python, R, Julia, Scala, Go, Javascript ...





- Interface: Python
- Often best on benchmarks
- Nervana was working on a neural network chip
- Bought by Intel





# Caffe

- Interface: C++, Python
- One of the earliest GPU supported
- Initial focus on computer vision (and CNNs)

Berkeley

Artificial Intelligence Research Laboratory

# Microsoft Cognitive Toolkit (CNTK)

- Interface: Custom Language (BrainScript), Python, C++, C#
- Multi GPU Support (scales very well)
- Mostly used at MS Research



# In the Browser

- Keras.js
  - GPU Support
  - Full sized networks
  - Can use trained Keras models
- ConvNetJS
  - Built by a Andrej Karpathy
  - Good for explaining neural network concepts
    - Fun to play around with
    - Very few requirements
  - Full CNN, RNN, Deep Q Learning

# References

All references cited in this presentation are listed in the following Google Sheets file:

<https://goo.gl/9Xhp2t>