



Università degli  
Studi di Pavia

# *Deep Learning*

## *01-Introduction*

Marco Piastra & Andrea Pedrini(\*)

(\*) Dipartimento di Matematica F. Casorati

*This presentation can be downloaded at:*  
<http://vision.unipv.it/DL>

# Main References for the This Course

## ■ Deep Learning

Ian Goodfellow, Yoshua Bengio and Aaron Courville

*MIT Press, 2017*

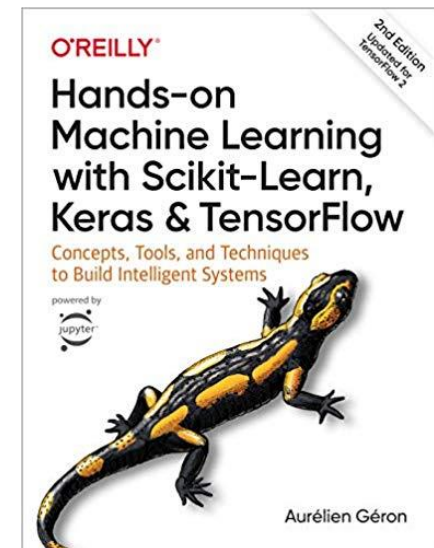
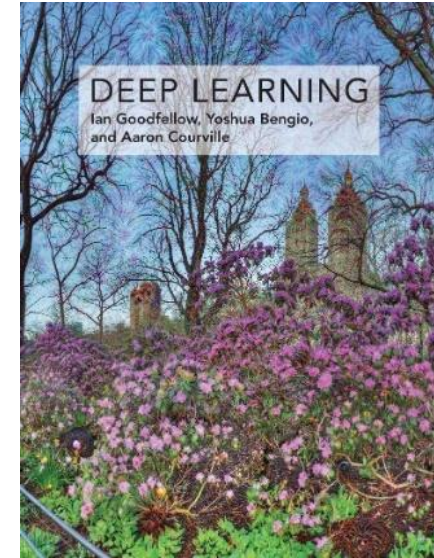
See also

<http://www.deeplearningbook.org/>

## ■ Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems, 2nd Edition

Aurélien Géron

*O'Reilly, 2020*



*Prologue:  
"Deep Learning"?*

# *Deep Learning in the Artificial Intelligence Cosmos*

## **Artificial Intelligence**

*automated reproduction of human cognitive activities*

## **Machine Learning**

*automated extraction of generalized knowledge  
from data and experience*

## **Deep Learning**

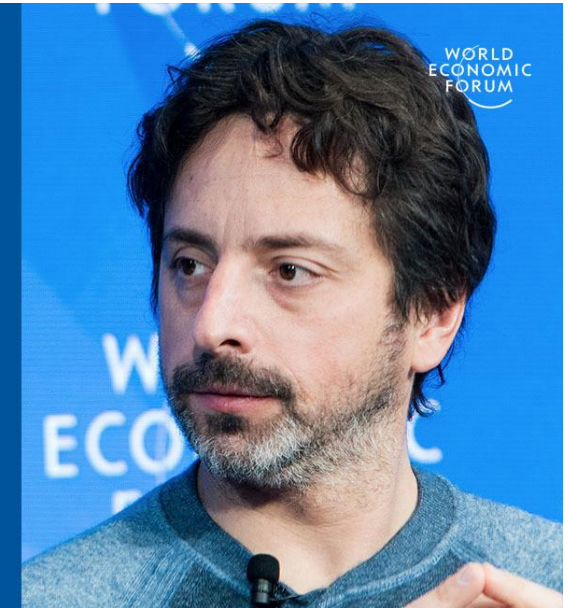
*automated extraction of generalized knowledge  
from data and experience ...*

*... using multi-layered **representations***

# AI strikes back?

The revolution in AI has been profound, it definitely surprised me, even though I was sitting right there.

Sergey Brin  
Google co-founder



- **Sergey Brin** [Google Co-Founder, January 2017]

*"I didn't pay attention to it [i.e. Artificial Intelligence] at all, to be perfectly honest."*

*"Having been trained as a computer scientist in the 90s, everybody knew that AI didn't work.*

*People tried it, they tried neural nets and none of it worked."*

[Quote and image from <https://www.weforum.org/agenda/2017/01/google-sergey-brin-i-didn-t-see-ai-coming/>]

# Artificial Intelligence Overhype

Forbes



10,693 views | Feb 15, 2019, 02:15pm

## 3 Reasons AI Is Way Overhyped



**Peter Cohan** Contributor   
Markets

- 1. Many CEOs Are Being Scared Into Caring Too Much About AI**
- 2. There Are Very Few Examples of High Payoff AI Applications**
- 3. Very Few Companies Can Afford or Find Good Uses For AI**

[Quote from <https://www.forbes.com/sites/petercohan/2019/02/15/3-reasons-ai-is-way-overhyped/#3d3fef8c5a6a/>]

# Artificial Intelligence Risks

## I'm an AI researcher, and here's what scares me about AI



Rachel Thomas

Follow

Jan 30 · 11 min read

AI is being increasingly used to make important decisions. Many AI experts (including [Jeff Dean](#), head of AI at Google, and [Andrew Ng](#), founder of Coursera and [deeplearning.ai](#)) say that **warnings about sentient robots are overblown**, but other harms are not getting enough attention. I agree. I am an AI researcher, and I'm worried about some of the societal impacts that we're already seeing. In particular, these 5 things scare me about AI:

1. Algorithms are often implemented without ways to address mistakes.
2. AI makes it easier to not feel responsible.
3. AI encodes & magnifies bias.
4. Optimizing metrics above all else leads to negative outcomes.
5. There is no accountability for big tech companies.

[Quote from <https://medium.com/@racheltho/im-an-ai-researcher-and-here-is-what-scaries-me-about-ai-909a406e4a71>

# Artificial Intelligence Hysteria?



## AI's current hype and hysteria could set the technology back by decades

July 24, 2019 10:11am BST

AI isn't as scary as we imagine. AndreyZH/Shutterstock

The reality of AI is currently very different, particularly when you look at the threat of automation. Back in 2013, researchers estimated that, in the following ten to 20 years, 47% of jobs in the US could be automated. Six years later, instead of a trend towards mass joblessness, we're in fact seeing US unemployment at a historic low.

Current AI is good at **finding patterns in large datasets**, and not much else.

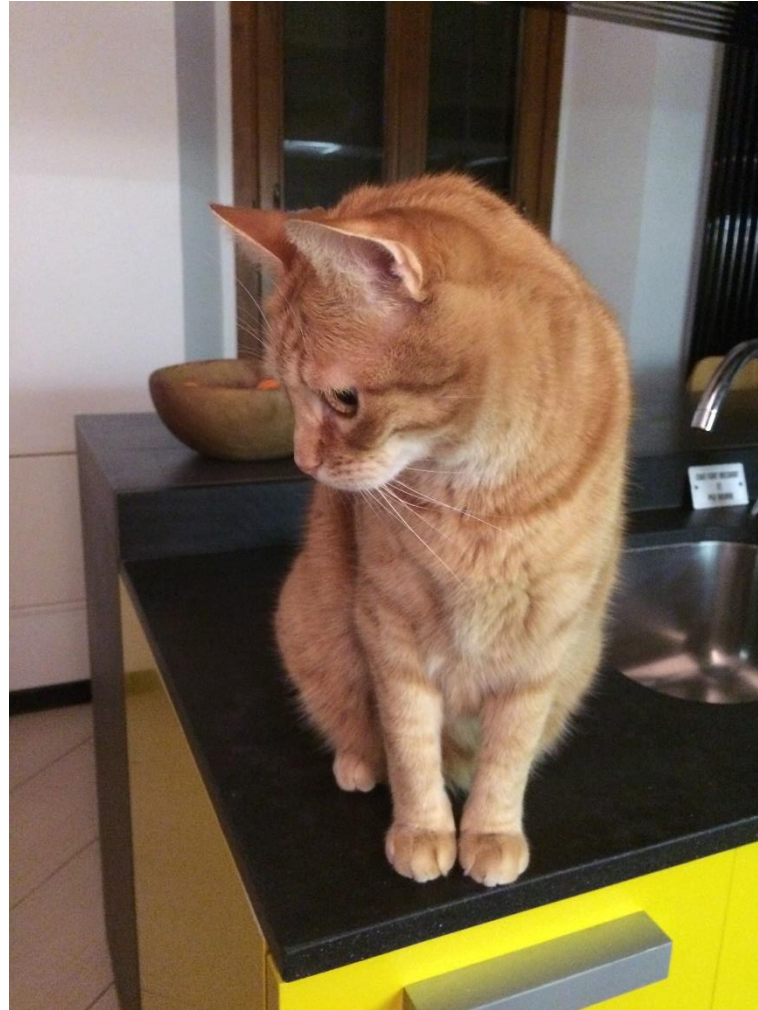
[Quote from <https://theconversation.com/ais-current-hype-and-hysteria-could-set-the-technology-back-by-decades-120514>]



# *Finding Patterns in Large Datasets*

# One Giant Leap for Mankind

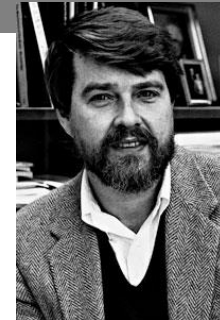
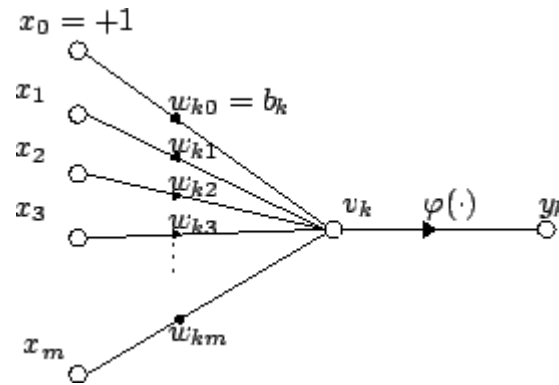
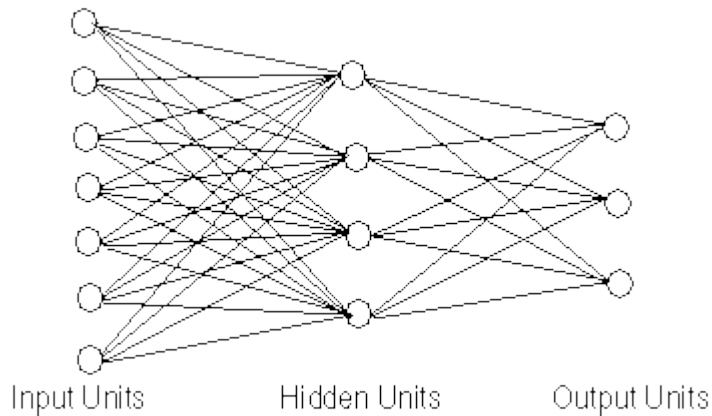
Is there a cat  
in this picture?



[this is *my* cat, Rabarbaro]

# *Artificial Neural Networks: The Origins*

# Emulation or simulation? Connectionism



(from Wikipedia)

*“In our view, people are smarter than today’s computers because the brain employs a basic computational architecture that is more suited to deal with a central aspect of the natural information processing tasks that people are so good at.”*

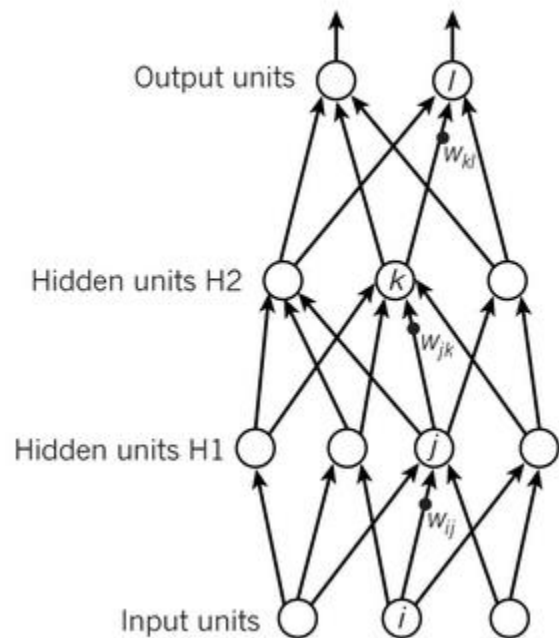
[Rumelhart, D.E., J.L. McClelland and the PDP Research Group (1986)  
*Parallel Distributed Processing: Explorations in the Microstructure of Cognition*]

## ■ **Basic assumption**

Mental phenomena can be described by interconnected networks of simple and often uniform units

# Artificial Neural Network

c



$$y_l = f(z_l)$$

$$z_l = \sum_{k \in H2} w_{kl} y_k$$

$$y_k = f(z_k)$$

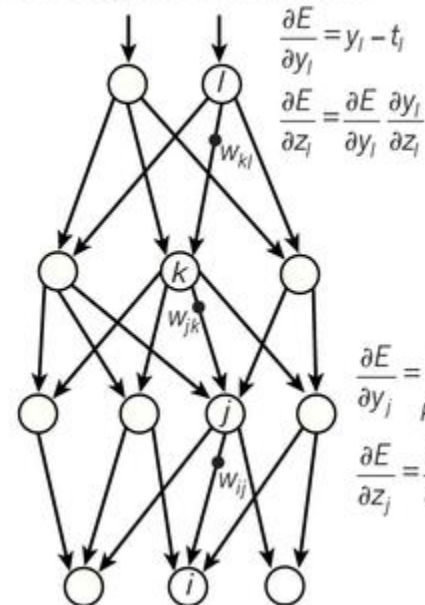
$$z_k = \sum_{j \in H1} w_{jk} y_j$$

$$y_j = f(z_j)$$

$$z_j = \sum_{i \in \text{Input}} w_{ij} x_i$$

d

Compare outputs with correct answer to get error derivatives



$$\frac{\partial E}{\partial y_l} = y_l - t_l$$

$$\frac{\partial E}{\partial z_l} = \frac{\partial E}{\partial y_l} \frac{\partial y_l}{\partial z_l}$$

$$\frac{\partial E}{\partial y_k} = \sum_{l \in \text{out}} w_{kl} \frac{\partial E}{\partial z_l}$$

$$\frac{\partial E}{\partial z_k} = \frac{\partial E}{\partial y_k} \frac{\partial y_k}{\partial z_k}$$

$$\frac{\partial E}{\partial y_j} = \sum_{k \in H2} w_{jk} \frac{\partial E}{\partial z_k}$$

$$\frac{\partial E}{\partial z_j} = \frac{\partial E}{\partial y_j} \frac{\partial y_j}{\partial z_j}$$

[figure from LeCun, Bengio, Hinton, *Deep Learning*, Nature 521, 2015]

## Function approximation

Basically, this is what a 'classical' artificial neural network does

## Supervised learning

The parameters (i.e. weights) are "learnt" from a dataset of inputs and expected outputs pairs

## Incremental optimization — a.k.a. "backward propagation"

Weights are progressively corrected to reduce the difference between actual and expected outputs

# Artificial Neural Networks

## ■ From *shallow* to *deep* networks

A feed-forward neural network with one hidden layer

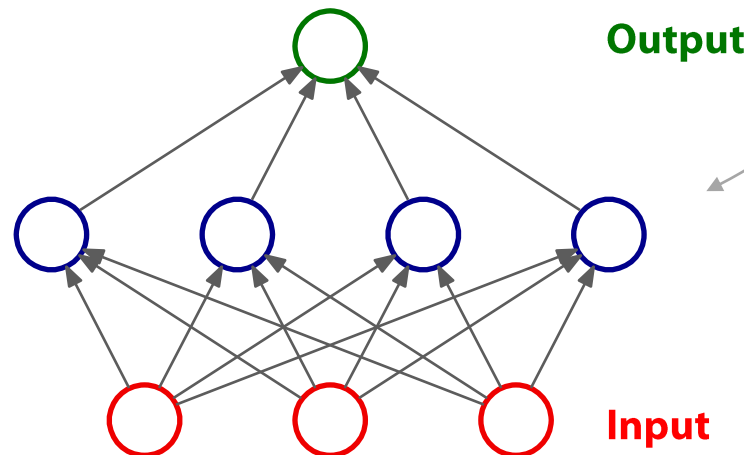
$$\tilde{y} = \mathbf{w} \cdot g(\mathbf{W}^{[1]} \mathbf{x} + \mathbf{b}^{[1]}) + b$$

Deep Learning systems  
(e.g. TensorFlow, PyTorch)  
use this representation

It can approximate any target function

$$y = f^*(\mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^d$$

(given enough units and proper *parameters*)



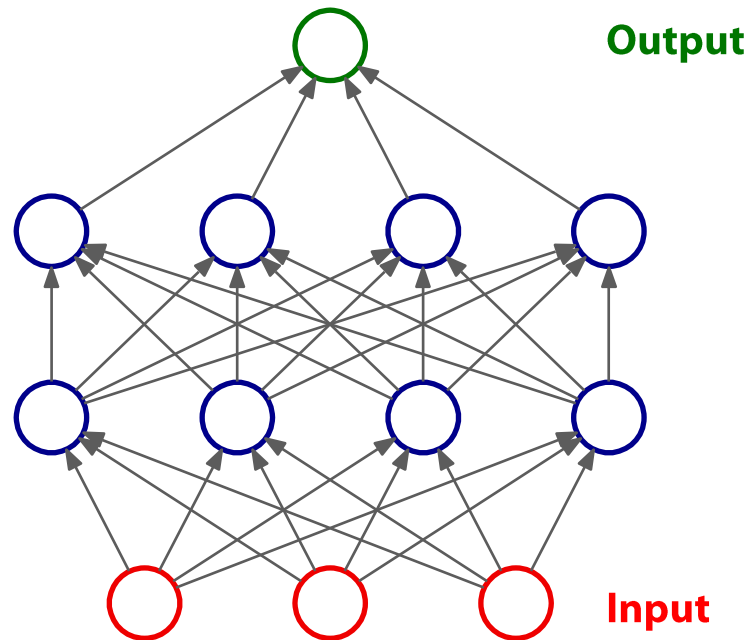
The two representations  
are equivalent

# Artificial Neural Networks

- **From *shallow* to *deep* networks**

A feed-forward neural network with two hidden layers

$$\tilde{y} = \mathbf{w} \cdot g(\mathbf{W}^{[2]} g(\mathbf{W}^{[1]} \mathbf{x} + \mathbf{b}^{[1]}) + \mathbf{b}^{[2]}) + b$$

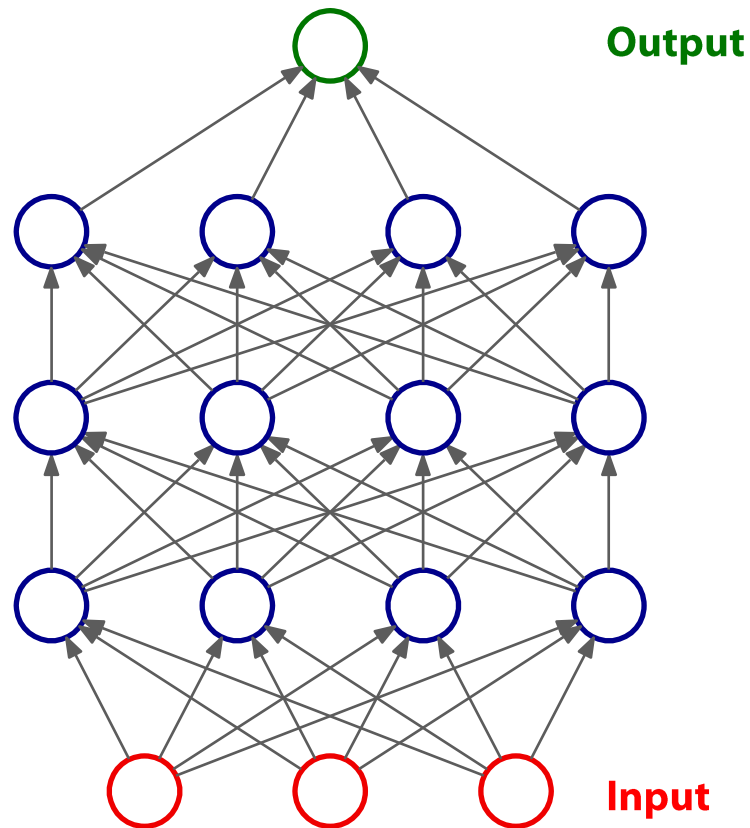


# Artificial Neural Networks

## ▪ From *shallow* to *deep* networks

A feed-forward neural network with three hidden layers

$$\tilde{y} = w \cdot g(\mathbf{W}^{[3]}g(\mathbf{W}^{[2]}g(\mathbf{W}^{[1]}x + \mathbf{b}^{[1]}) + \mathbf{b}^{[2]}) + \mathbf{b}^{[3]}) + b$$

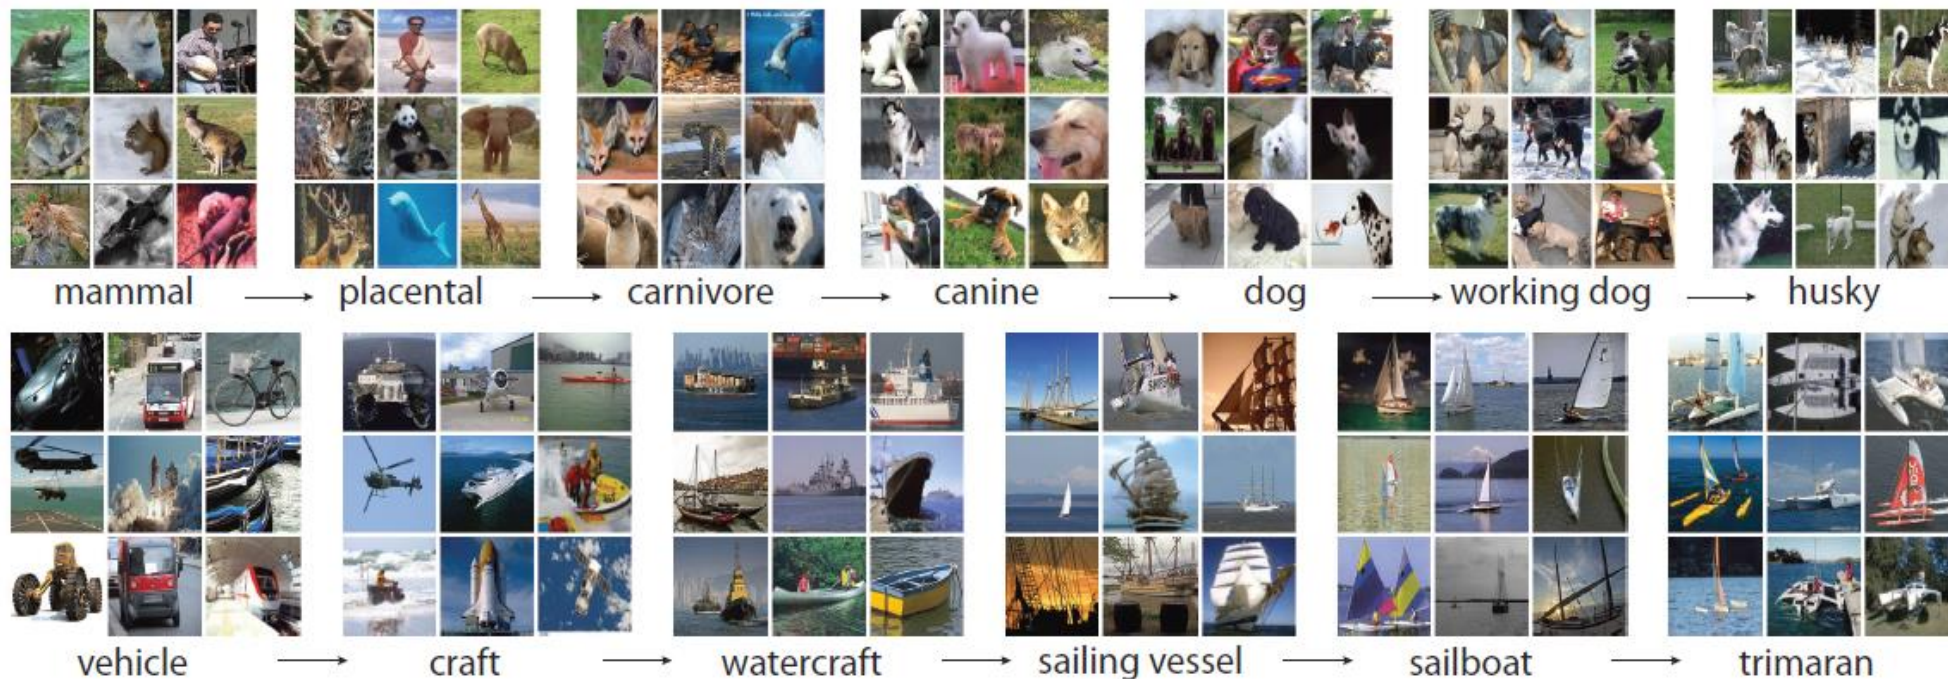




*Well, it's just a function anyway ...*

# ImageNet Challenge

## ■ The ImageNet Large Scale Visual Recognition Challenge



1,461,406 full resolution images

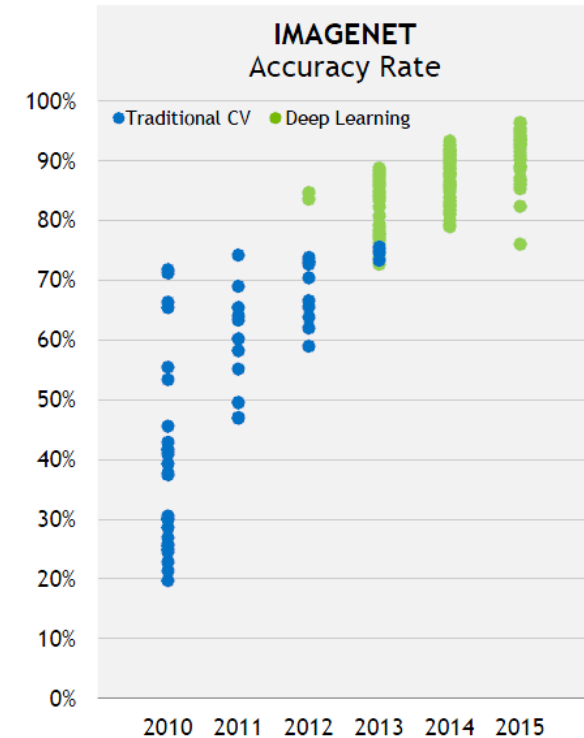
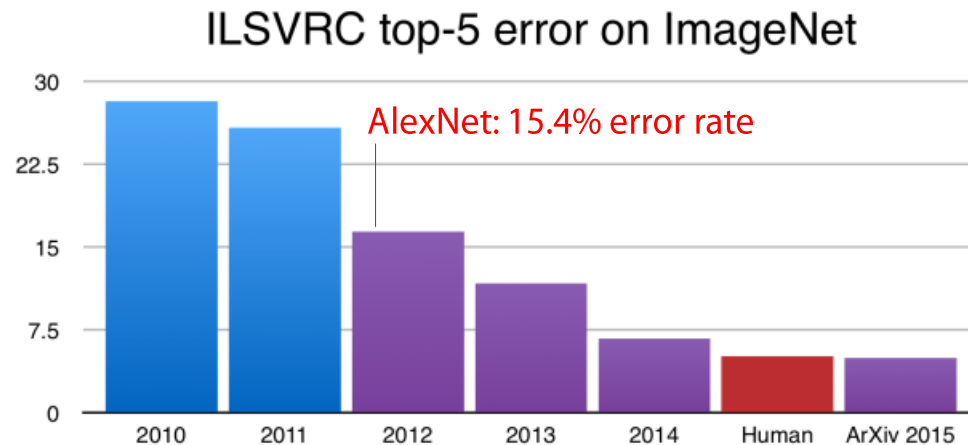
Complex and multiple textual annotation,  
hierarchy of 1000 object classes along several dimensions

*The image classification challenge was run annually from 2010 to 2017*

[figures from [www.nvidia.com](http://www.nvidia.com)]

# ImageNet Challenge

## ■ The ImageNet Large Scale Visual Recognition Challenge



1,461,406 full resolution images

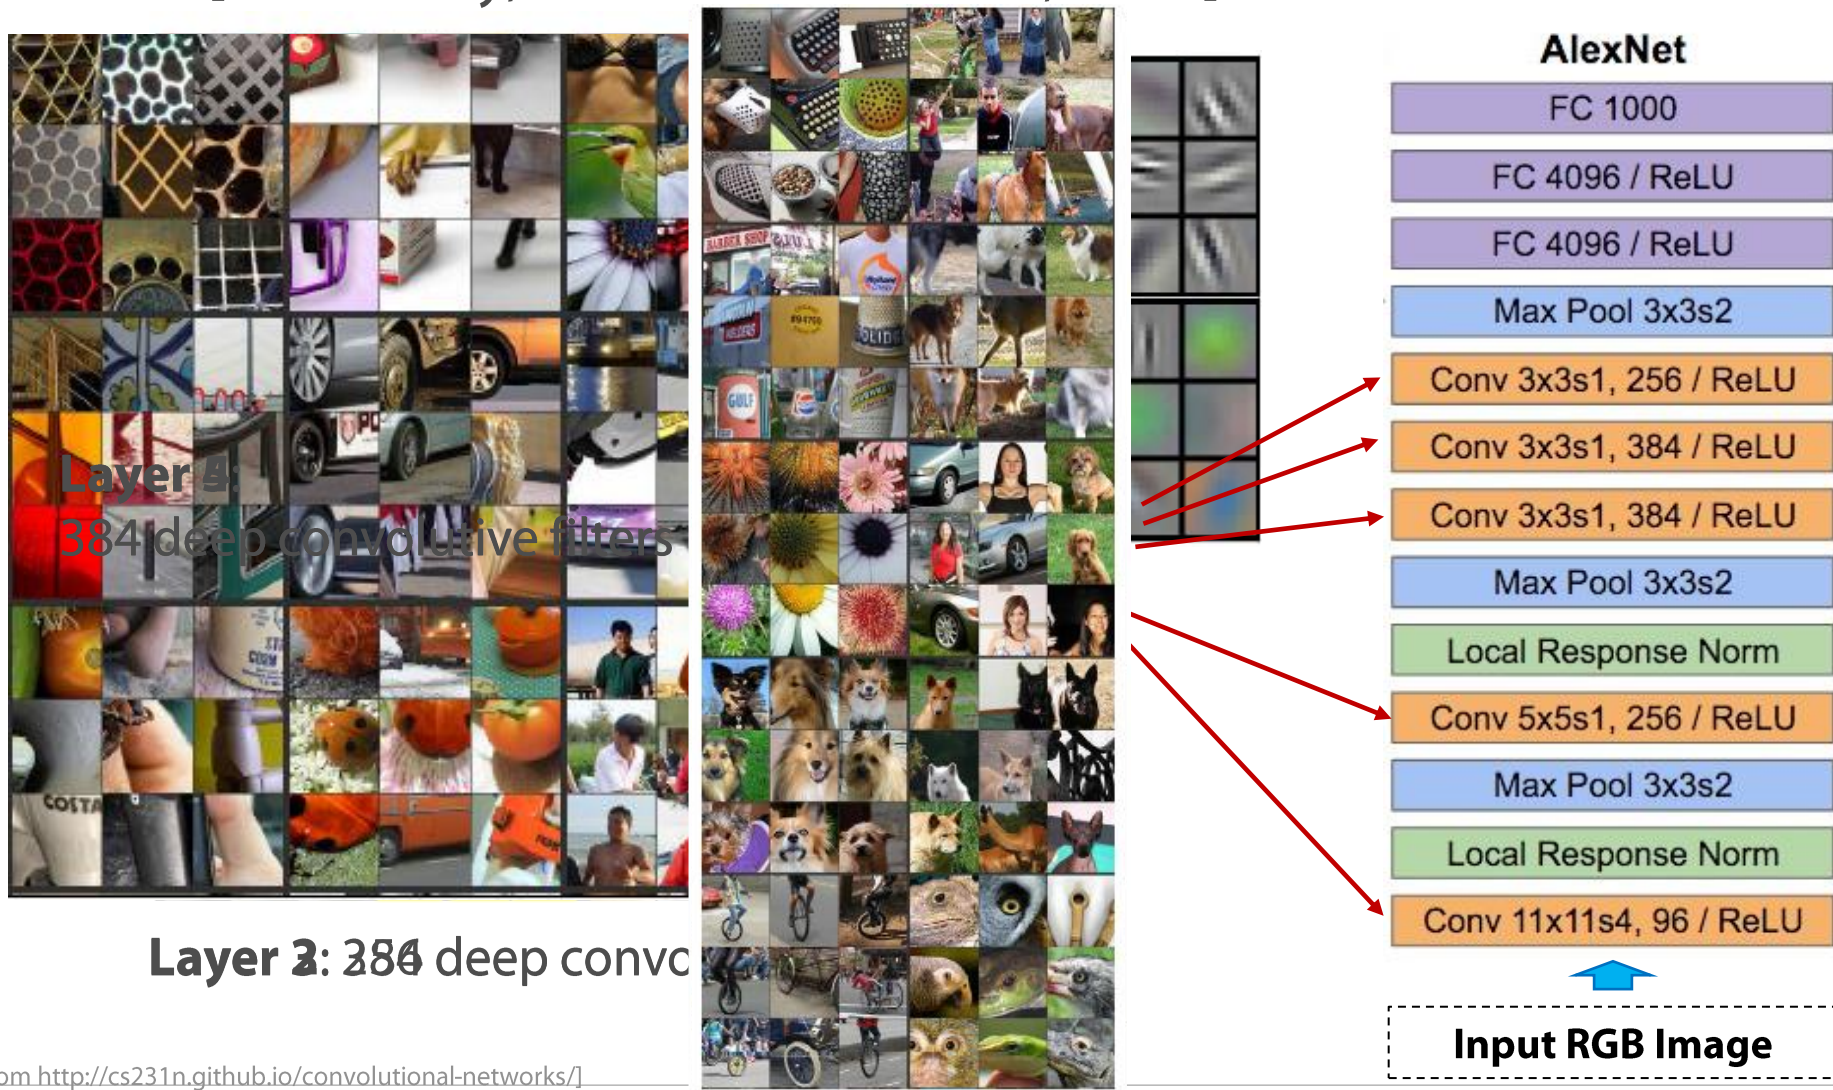
Complex and multiple textual annotation,  
hierarchy of 1000 object classes along several dimensions

*The image classification challenge was run annually from 2010 to 2017*

[figures from [www.nvidia.com](http://www.nvidia.com)]

# Deep Convolutional Neural Networks (DCNN)

- **AlexNet** [Krizhevsky, Sutskever & Hinton, 2012]



[images from <http://cs231n.github.io/convolutional-networks/>]

*Well, it's just a function anyway  
(contd.)...*

# Games of strategy: AlphaGo (2016)

Image from: <https://nikcheerla.github.io/deeplearningschool/2018/01/01/AlphaZero-Explained/>



- Mastering the game of Go with deep neural networks and tree search [2016, D. Silver, et al. (22 authors), <http://www.nature.com/nature/journal/v529/n7587/full/nature16961.html>]

*There are more possible positions in Go than there are atoms in the universe*

## Sophisticated machine-learning techniques

*Strategy selection via Monte Carlo Tree Search (MCTS)*

*Deep neural networks (trained on human matches) for both guidance and learning*

*Adversarial self-training:*

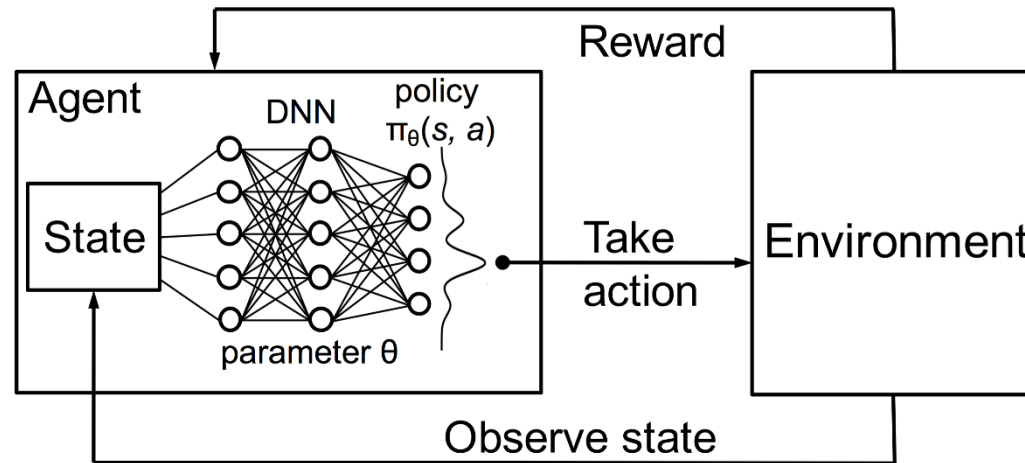
*playing against itself and improving via reinforcement learning*

## Super-human player?

On March 2016, AlphaGo won 4-1 against the legendary Lee Sedol, the top Go player in the world over the past decade

# Deep Reinforcement Learning (DRL)

- **A Deep Neural Network learns a policy**



*The agent interacts with an environment (it could be a copy of itself)*

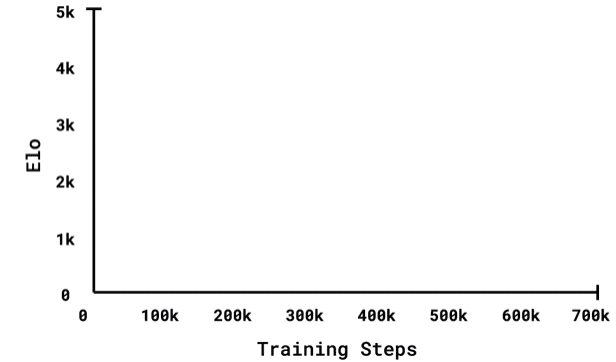
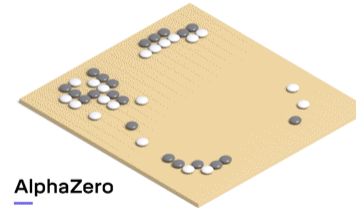
*It selects an **action** in each **state** and receives a **reward** (possibly deferred) as a function of the results obtained*

*The DRL system optimizes its policy*

# Beyond Emulating Humans: AlphaZero (2018)

Image from: <https://deepmind.com/blog/article/alphazero-shedding-new-light-grand-games-chess-shogi-and-go>

*AlphaGo is heavily reliant  
on the experience of human players*



## ■ AlphaZero learns by itself

[2018, D. Silver, et al. (13 authors), <https://science.sciencemag.org/content/362/6419/1140.full> ]

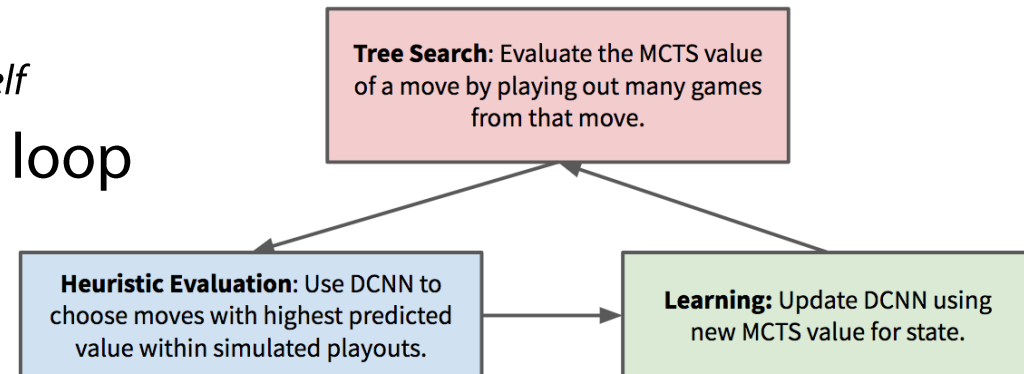
### Basic Knowledge Only

*It just knows the basic rules of the games*

### Learning via Self-Play

*It plays against a (frozen) copy of itself*

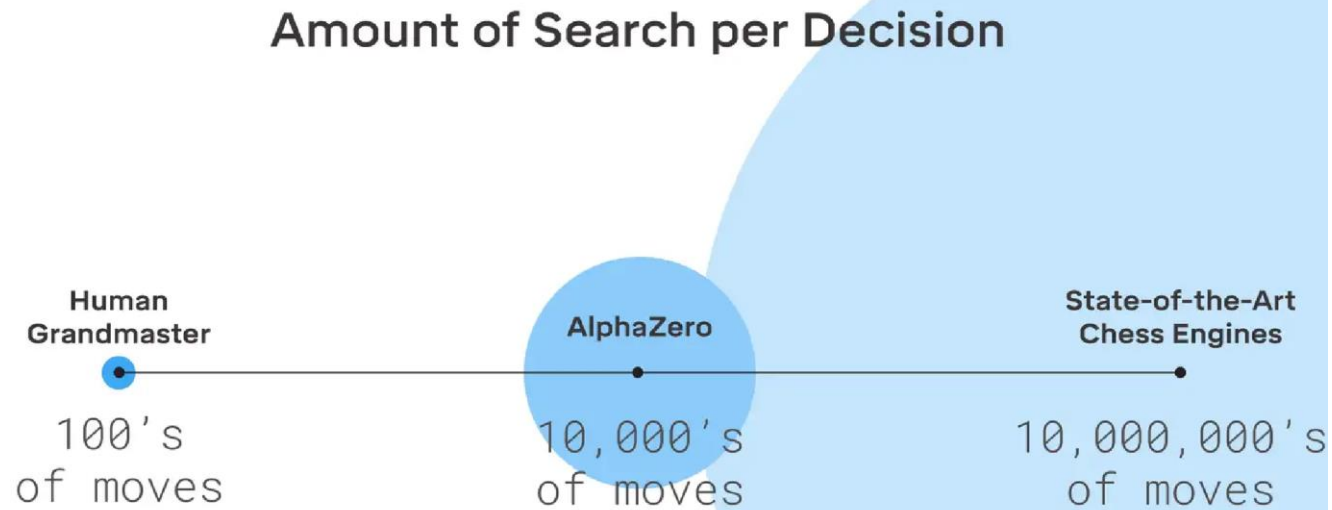
### MCTS and DCNN in a closed loop





# Beyond Emulating Humans: AlphaZero (2018)

Image from: <https://deepmind.com/blog/article/alphazero-shedding-new-light-grand-games-chess-shogi-and-go>



- **AlphaZero uses much less 'brute force' search**

When playing, the search process is driven by its neural network

*It acts like a memory of past experiences*

While training, it learns through a huge amount of self-playing

*But it is a faster learner than Alpha Go*

# AlphaFold (2020)

Images from <https://deepmind.com/blog/article/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology>

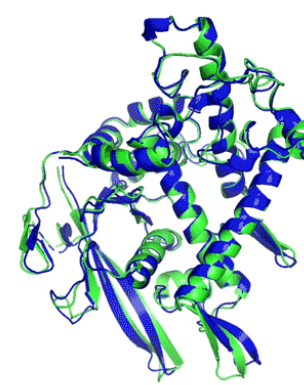
## ■ Predicting how proteins fold

### Several Neural Networks

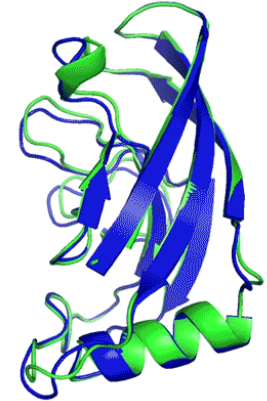
*Trained separately, working together*

### A big step ahead

*Substantial improvement  
over the best actual systems yet*

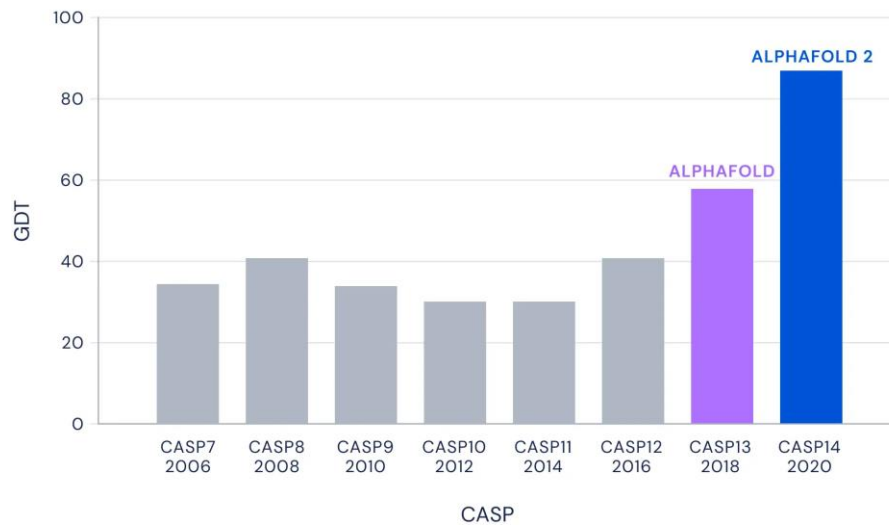


T1037 / 6vr4  
90.7 GDT  
(RNA polymerase domain)



T1049 / 6y4f  
93.3 GDT  
(adhesin tip)

### Median Free-Modelling Accuracy



● Experimental result  
● Computational prediction

## *It's open source*

In the version used for the CASP13 challenge  
[https://github.com/deepmind/deepmind-research/tree/master/alphafold\\_casp13](https://github.com/deepmind/deepmind-research/tree/master/alphafold_casp13)

# GPT-3 (2020)

Image from <https://www.theguardian.com/commentisfree/2020/sep/08/robot-wrote-this-article-gpt-3>

**Support The Guardian**

Available for everyone, funded by readers

Contribute →

Subscribe →

My account ▾

**The  
Guardian**

**Opinion** Artificial intelligence (AI)

## A robot wrote this entire article. Are you scared yet, human?

### GPT-3

We asked GPT-3, OpenAI's powerful new language generator, to write an essay for us from scratch. The assignment? To convince us robots come in peace

- For more about GPT-3 and how this essay was written and edited, please read our editor's note below

Tue 8 Sep 2020 09.45 BST



70,298  1,188 



Fortier fed GPT-3 a strange prompt: “Below is a transcript from an interview where Barack Obama explained why he was banned from Golden Corral for life.” The system then filled in the rest of the interview, running with the concept that Obama had been banned from an all-you-can-eat buffet.

**Obama:** Yes. It’s true. I am no longer allowed in Golden Corral.

**Interviewer:** Is this because of your extensive shrimp-n-crab legs policy?

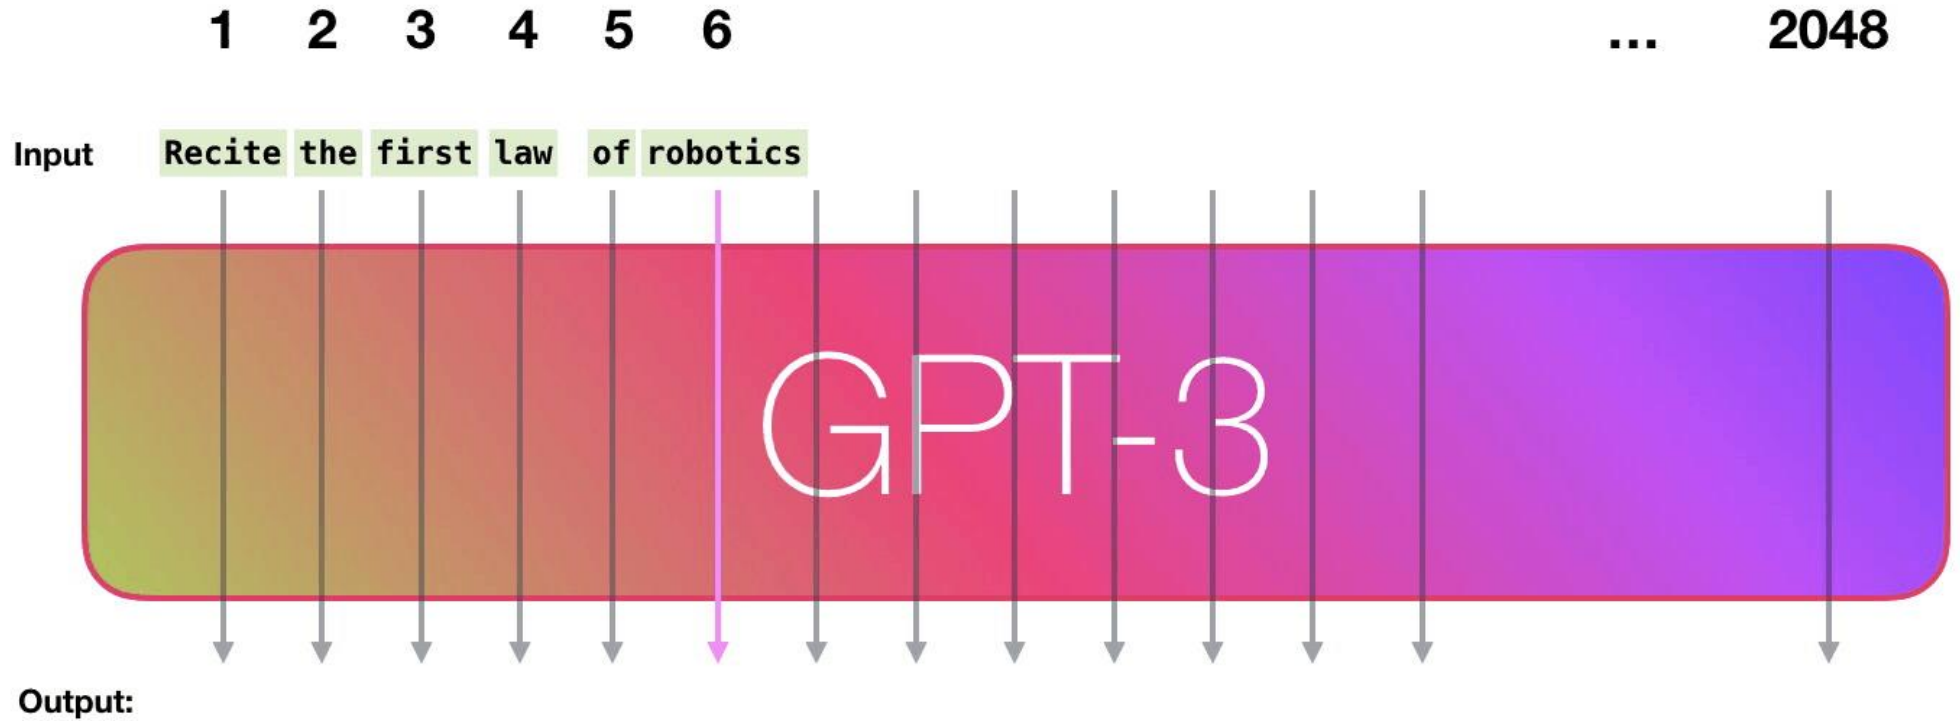
**Obama:** Absolutely.

**Interviewer:** What is your extensive shrimp-n-crab legs policy?

**Obama:** Oh, well, in brief, they were offering an all-you-can-eat shrimp-n-crab leg buffet, and I did not hesitate. After I ate so much shrimp and crab that my stomach hurt, I would quietly sneak in and throw more shrimp and crab onto my plate. I did this over and over again until I had cleaned out the buffet and was full of shrimp-n-crab.

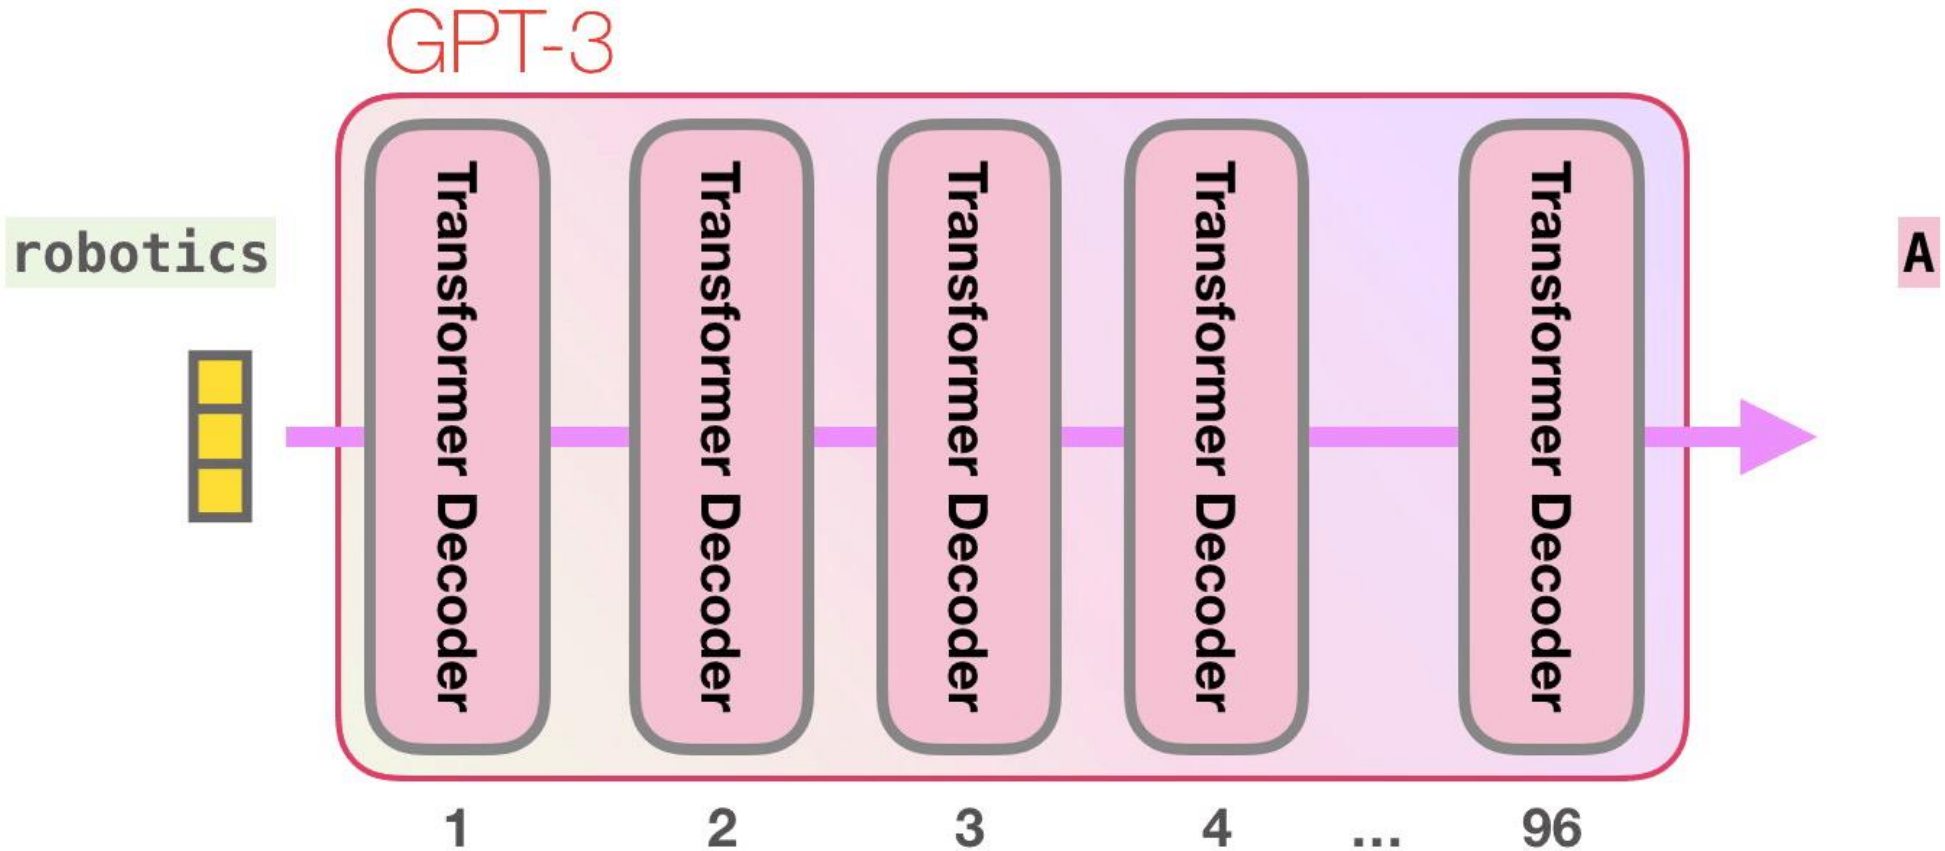
# GPT-3 (2020)

Image from <http://jalammr.github.io/how-gpt3-works-visualizations-animations/>



# GPT-3 (2020)

Image from <http://jalammar.github.io/how-gpt3-works-visualizations-animations/>



One of the biggest Neural Networks yet

*GPT-3 has 175 Billion parameters  
(AlexNet has 64 Million)*

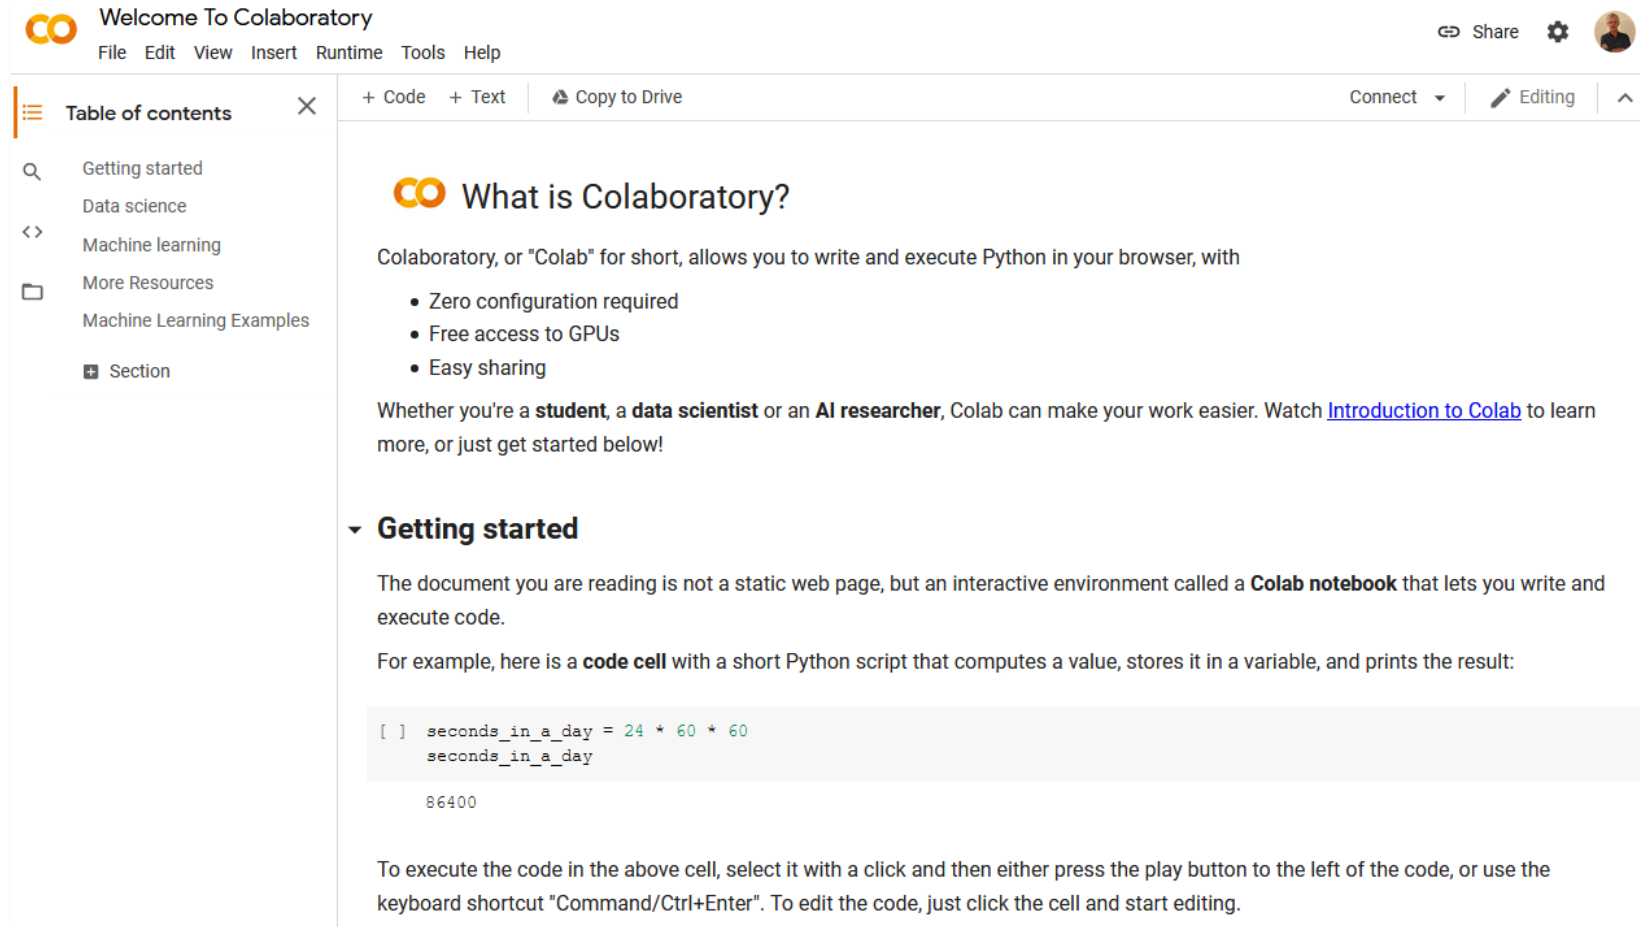
*OK, about this course ...*

# *This Course in a Nutshell*

- 1) Fundamentals**
- 2) Deep Supervised Learning**
- 3) Deep Convolutional Neural Networks**
- 4) Deep Recurrent Networks**
- 5) Deep Reinforcement Learning**



# Labs with Google Colab



The screenshot shows the Google Colaboratory interface. At the top, it says "Welcome To Colaboratory" with a menu bar containing "File", "Edit", "View", "Insert", "Runtime", "Tools", and "Help". On the right, there are "Share", "Settings", and a user profile icon. A "Table of contents" sidebar on the left lists: "Getting started", "Data science", "Machine learning", "More Resources", "Machine Learning Examples", and "Section". The main content area is titled "What is Colaboratory?" and includes the following text:

Colaboratory, or "Colab" for short, allows you to write and execute Python in your browser, with

- Zero configuration required
- Free access to GPUs
- Easy sharing

Whether you're a **student**, a **data scientist** or an **AI researcher**, Colab can make your work easier. Watch [Introduction to Colab](#) to learn more, or just get started below!

**Getting started**

The document you are reading is not a static web page, but an interactive environment called a **Colab notebook** that lets you write and execute code.

For example, here is a **code cell** with a short Python script that computes a value, stores it in a variable, and prints the result:

```
[ ] seconds_in_a_day = 24 * 60 * 60
seconds_in_a_day

86400
```

To execute the code in the above cell, select it with a click and then either press the play button to the left of the code, or use the keyboard shortcut "Command/Ctrl+Enter". To edit the code, just click the cell and start editing.

Make sure you have a look!

*Better yet: follow the tutorial at*

<https://colab.research.google.com/notebooks/intro.ipynb>

# PyCharm (optional)

However, if you really want to see how things work...

- **Set it up on your computer**

Python 3.7+  
Numpy 1.17+  
TensorFlow 2.+

... and, above all ...

Install PyCharm (Community Edition)

<https://www.jetbrains.com/pycharm/>

*It's free*

Learn using the debugger :  
it will change the way you learn



**CAUTION:** no assistance whatsoever will be provided for this ...  
*"Legions of students succeeded in doing this on their own: you can make it too!"*

# The Final Exam

## 1) Choose and propose a final project

The topic could be any of your choice

The techniques adopted must be (strongly) related with the course

Groupwork is allowed, with a maximum of two (*no exceptions*)

## 2) Develop and submit your projects

Each final project must be submitted as a Google Colab notebook plus dataset (*if required*)

Submission must be made at least two days before the exams

## 3) Be strong on theory as well

Not even a *phenomenal* final project, alone, will give you a final grade (*sorry*)

See also below

## 4) Final Exam

Oral interview, about both theory and final project

Relative weights: theory 60%, final project 40%