

Artificial Intelligence and Machine Learning: The Quiet Revolution

Nikolaos Passalis

Dept. of Informatics

Aristotle University of Thessaloniki, Greece

What is Intelligence?

Intelligence

- Property that characterizes human behavior!

Intelligence

- Property that characterizes human behavior!
- Are your choices really intelligent?



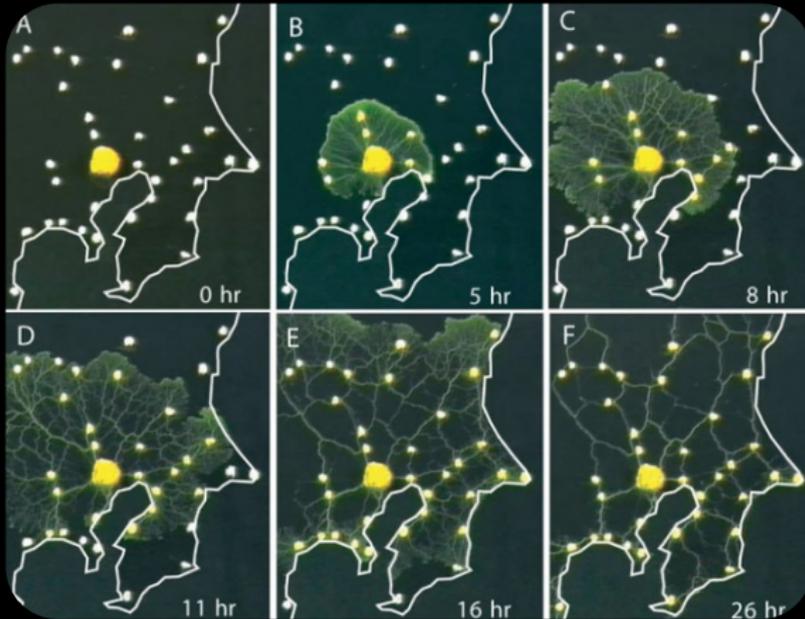
Intelligence

- Closely associated with human behavior, but not so distinctive feature!
- Many other animals, or even unicellular organisms, exhibit intelligent behavior!



Intelligence

- A slime mold can create a network as efficient as the current Tokyo Railway network!



Intelligence

- Not easy to define!

'A very general mental capability that, among other things, involves the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, learn quickly and learn from experience.'

Mainstream Science on Intelligence (Wall Street Journal, 1994)

Intelligence

- Many other definitions also exist!

'Intelligence measures an agents ability to achieve goals in a wide range of environments.' Legg & Hutter

'Intelligence is a force, F , that acts so as to maximize future freedom of action. It acts to maximize future freedom of action, or keep options open, with some strength T , with the diversity of possible accessible futures, S , up to some future time horizon, τ . In short, intelligence doesn't like to get trapped.' Alexander Wissner-Gross

Why do we need artificial intelligence?

- So, are we trying to “recreate” something that we do not even know what it is?
- Why is intelligence useful?

Why do we need artificial intelligence?

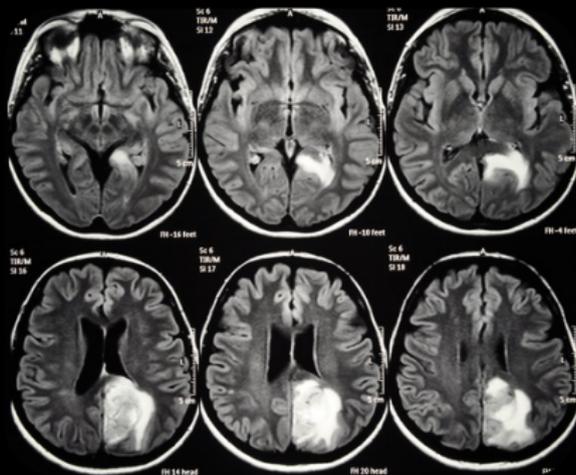
- So, are we trying to “recreate” something that we do not even know what it is?
- Why is intelligence useful?
- Automate laborious and tedious tasks!
- A new “Industrial Revolution”?

History of A.I.

- Myth of **Talos**: Bronze Age (3200-600 BCE)
 - Talos was a giant “robot” made of bronze to protect Crete from invaders.
- Many developments followed (mathematics, logic, programmable computers)
- In 1950, Alan Turing proposed the **Turing Test**
 - A machine is considered *intelligent* if a human cannot distinguish between the machine and a human.

History of A.I.

- Failed on less structured tasks...
 - Image understanding, speech recognition, etc.
 - Diagnose diseases from raw medical data instead of just encoding the existing knowledge!



History of A.I.

- In 1958 a computational model of a neuron that can be used to solve simple problems was proposed!
 - **Perceptrons!**
 - Inspired by the McCulloch-Pitts model (1943)!
- A model of a neuron is useless!
- We have to **train** it to do some useful stuff!

History of A.I.

- In 1958 a computational model of a neuron that can be used to solve simple problems was proposed!
 - **Perceptrons!**
 - Inspired by the McCulloch-Pitts model (1943)!
- A model of a neuron is useless!
- We have to **train** it to do some useful stuff!
- No practical way to train networks composed of perceptrons.
- Limited practical value!

History of A.I.

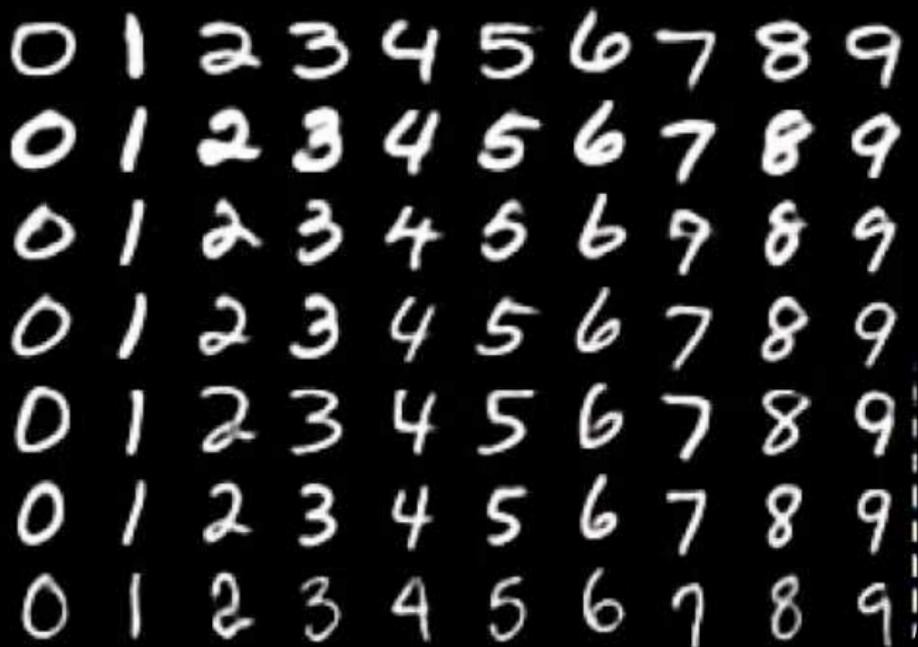
- In 1975 a method for training networks with many such neurons was proposed!
 - The algorithm is called “back-propagation”
- There was some excitement about neural networks!
- After all, they work similarly to the human brain (to some extent)
- But the scientific community put them aside!
 - No solid theory that explains their behavior (this is still true!)
 - Many resources (and data) are needed to train large network!
 - Prone to several poorly understood phenomena (more on this later!)

The dawn of Deep Learning

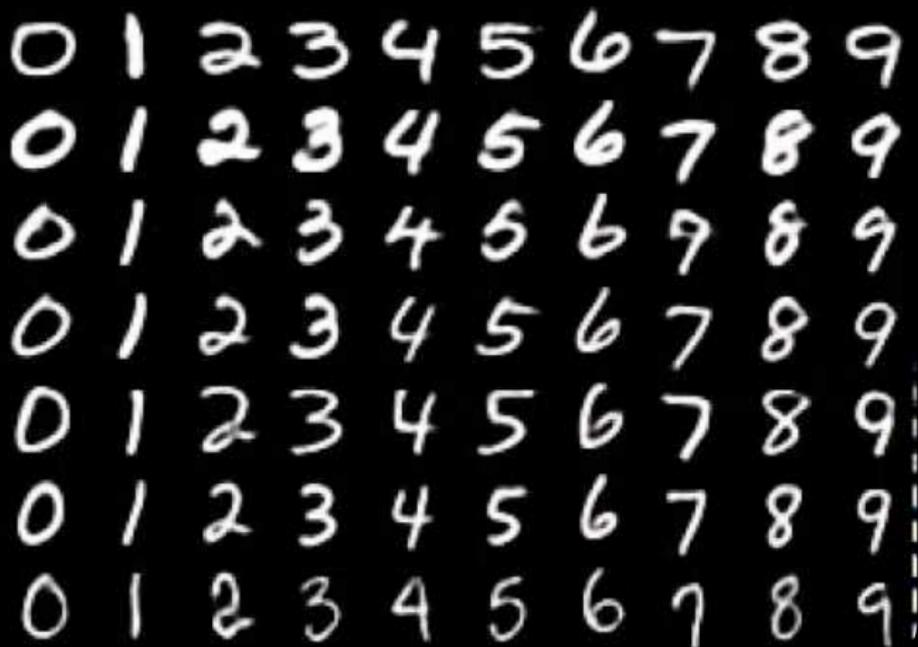
- After 2000 there were several breakthroughs that allowed for the creation of powerful neural networks.
 - **Powerful computing hardware** (dedicated hardware specialized for training neural networks)
 - **Data are now everywhere!**
 - New architectures and methods for training large networks!
- **Deep** networks composed of *many layers* of neurons → **deep learning**.
- Solved many tasks with great success!!!
- Achieve state-of-the-art performance on almost every task they are used for!

So, what can a neural network do?

- Recognize a handwritten digit!

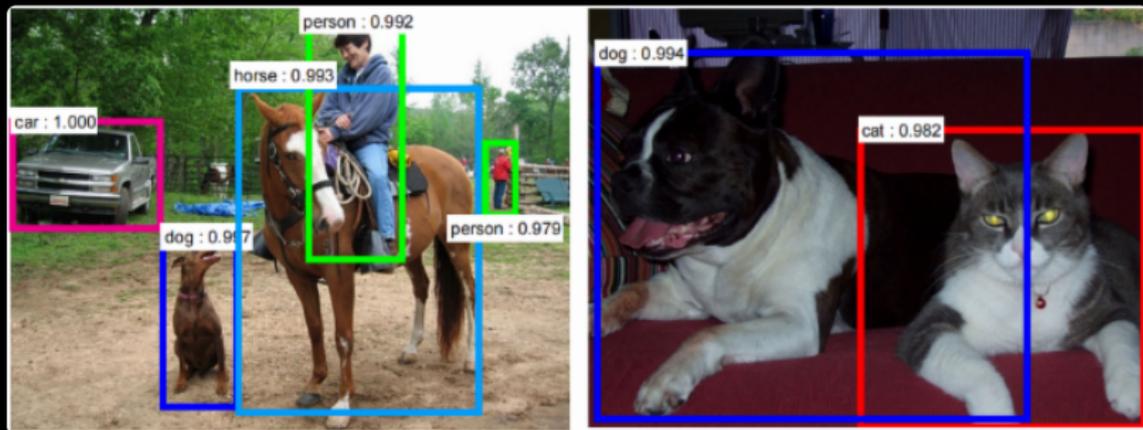


- Recognize a handwritten digit!

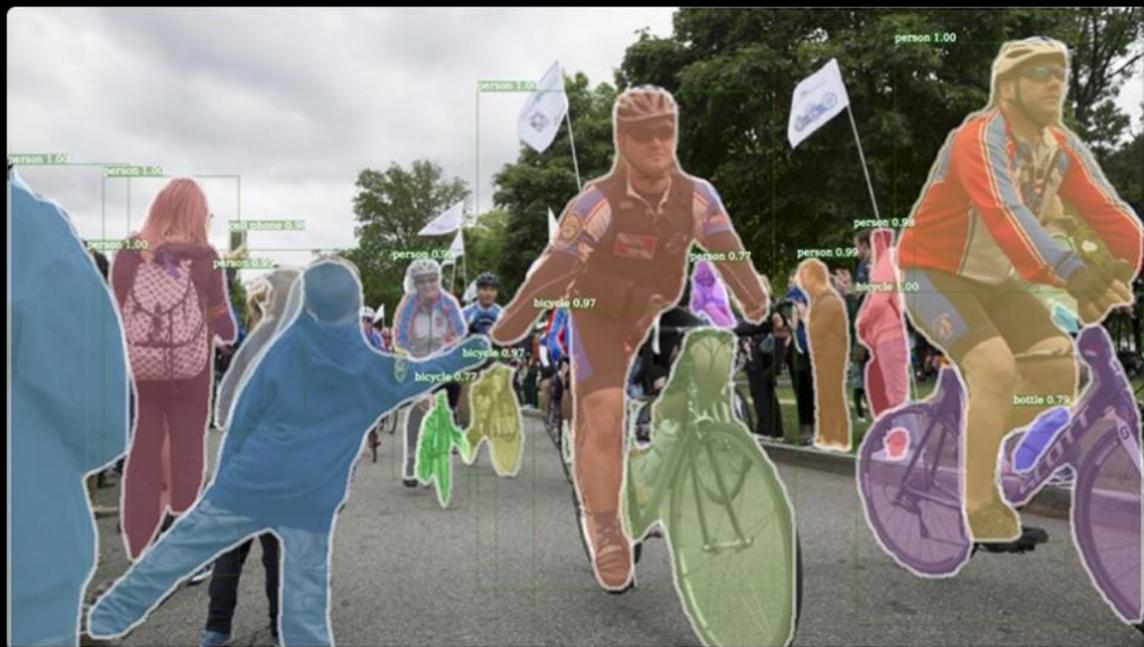


- Recognition accuracy $> 99.80\%$

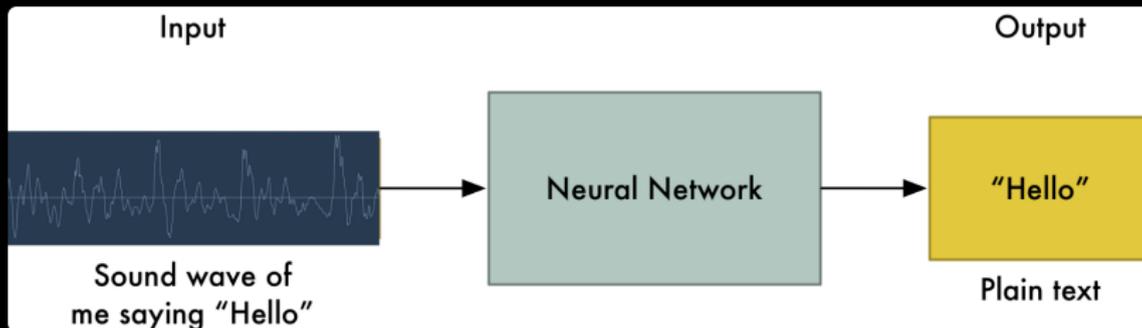
- Detect objects!



- Fine-grained image segmentation!



- Speech recognition!



State-of-the-art accuracy on many Machine Learning tasks...

Hm... Is that A.I.?

Let's do some serious stuff!

- Image captioning!



"man in black shirt is playing guitar."



"construction worker in orange safety vest is working on road."



"two young girls are playing with lego toy."

- Visual questioning answering!

Who is wearing glasses?

man



woman



Where is the child sitting?

fridge



arms



Is the umbrella upside down?

yes



no



How many children are in the bed?

2



1



- Generating images!



Figure 3. Example results by our proposed StackGAN, GAWWN [20], and GAN-INT-CLS [22] conditioned on text descriptions from CUB test set. GAWWN and GAN-INT-CLS generate 16 images for each text description, respectively. We select the best one for each of them to compare with our StackGAN.

- Self-driving cars!



- Image style transfer!

1

Upload photo

The first picture defines the scene you would like to have painted.



2

Choose style

Choose among predefined styles or upload your own style image.



3

Submit

Our servers paint the image for you. You get an email when it's done.



- Image style transfer!



- Image style transfer!



Original photo

Reference photo

Result

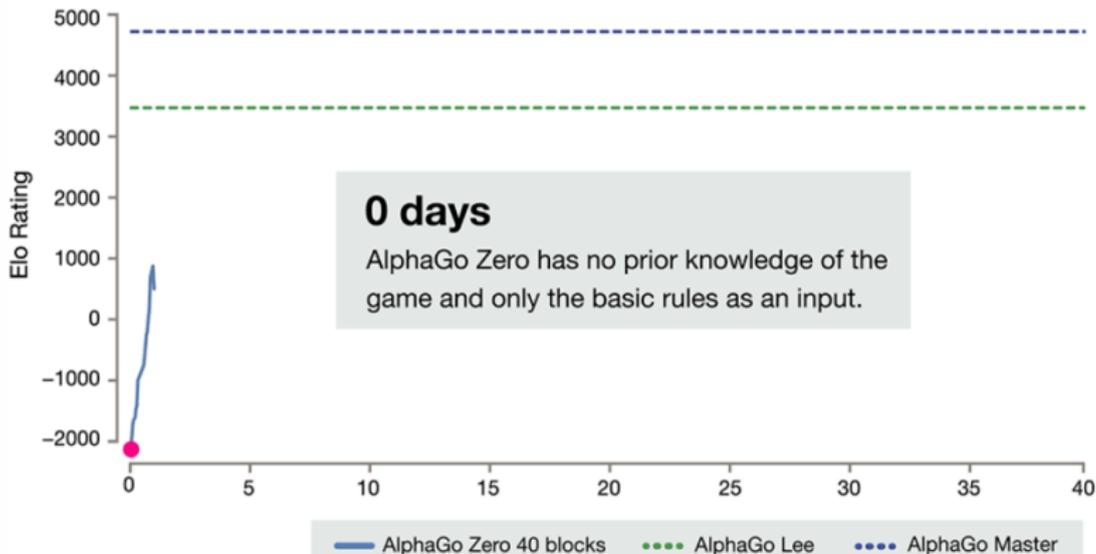
- Dreaming!

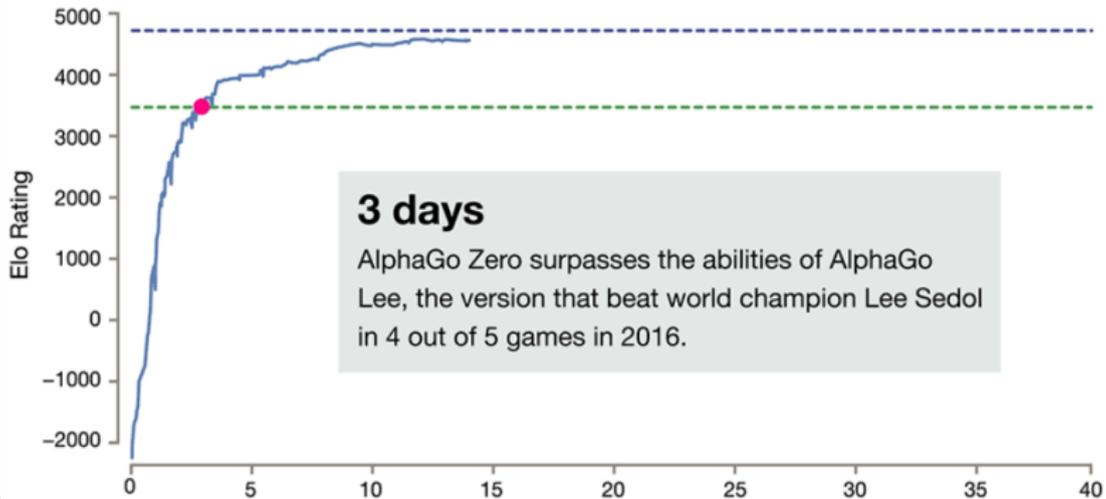


- Dreaming!



- Playing games better than humans!
- Alpha Go Zero learned to play Go from scratch better than any other human (or machine) in just 40 days!





3 days

AlphaGo Zero surpasses the abilities of AlphaGo Lee, the version that beat world champion Lee Sedol in 4 out of 5 games in 2016.





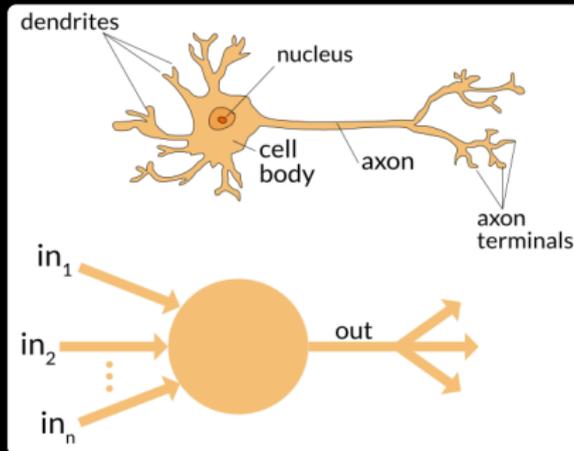
40 days

AlphaGo Zero surpasses all other versions of AlphaGo and, arguably, becomes the best Go player in the world. It does this entirely from self-play, with no human intervention and using no historical data.

So, how does a neural network work?

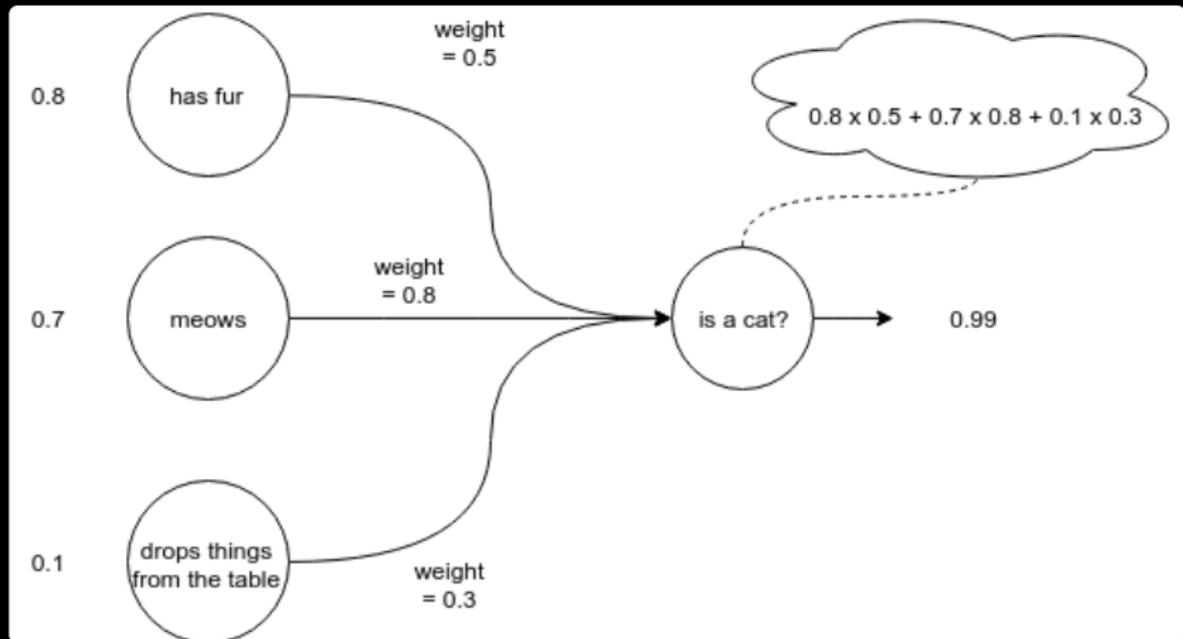
The neuron

- The idea is simple: every neuron receives the output of each other neuron and makes a decision!
- Each neuron does not *trust* equally each other neuron!
- Each input is associated with a *weight* that alters the importance of the corresponding input.



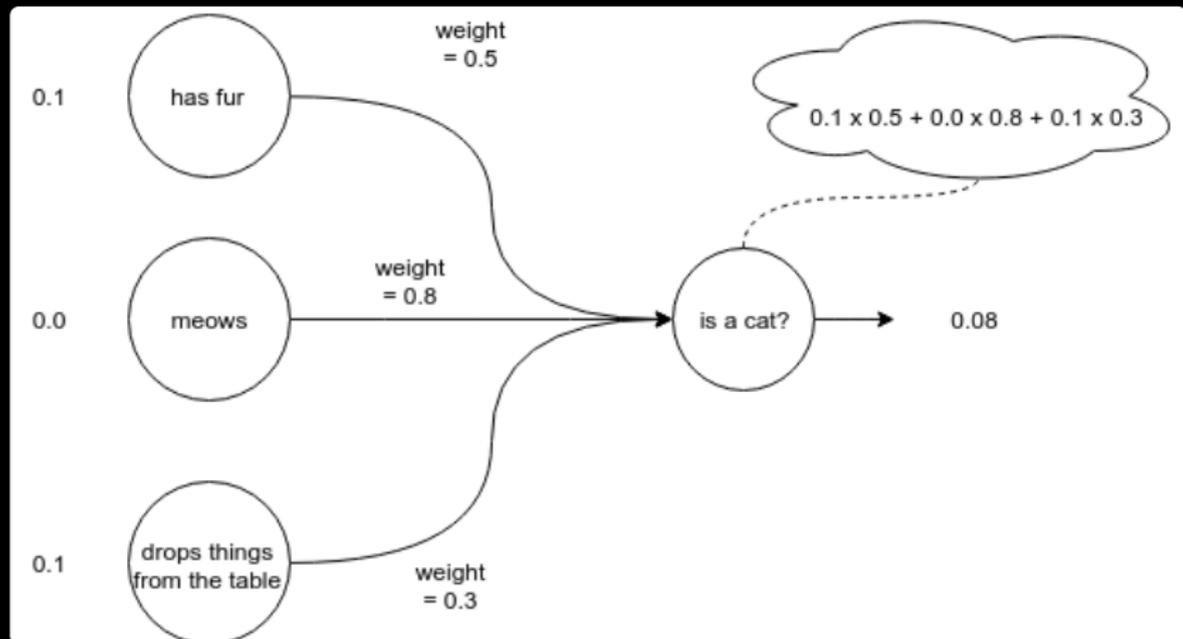
The neuron

- An example!



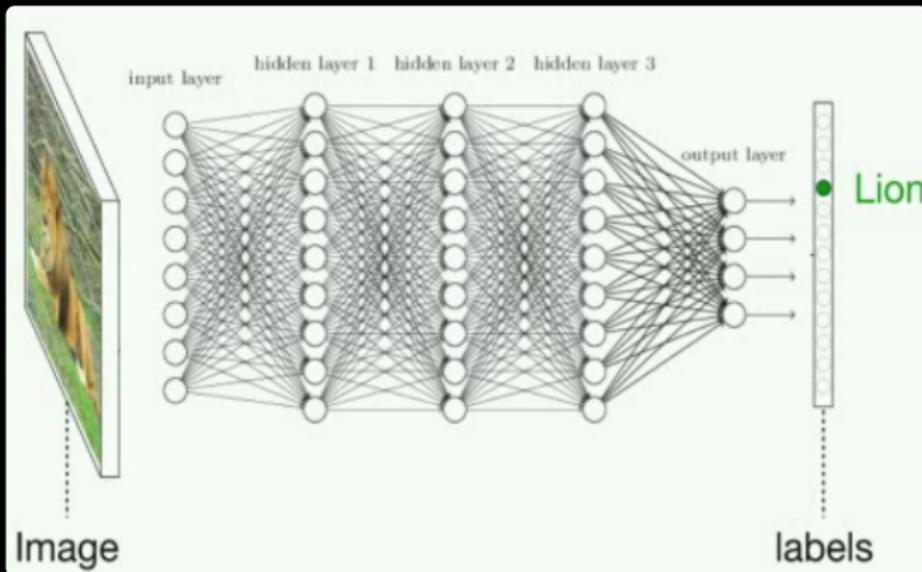
The neuron

- An example!



The neuron

- How do we discover neurons that know if “has fur”, “meows” or “drops things from the table”?
- We stack multiple neurons!



Training

- How do we calculate the appropriate weights for the neurons?
- A modern neural network might have more than 60,000,000 weights (parameters)!

Training

- How do we calculate the appropriate weights for the neurons?
- A modern neural network might have more than 60,000,000 weights (parameters)!
- Answer: Tell them what to do and then adjust the weights to be closer to the correct output.
- A whole field of mathematics is devoted to solving such problems (optimization)!
- Usually the “back-propagation” algorithm is used!

Back-propagation

- We leave it to mathematicians!

Back-propagation

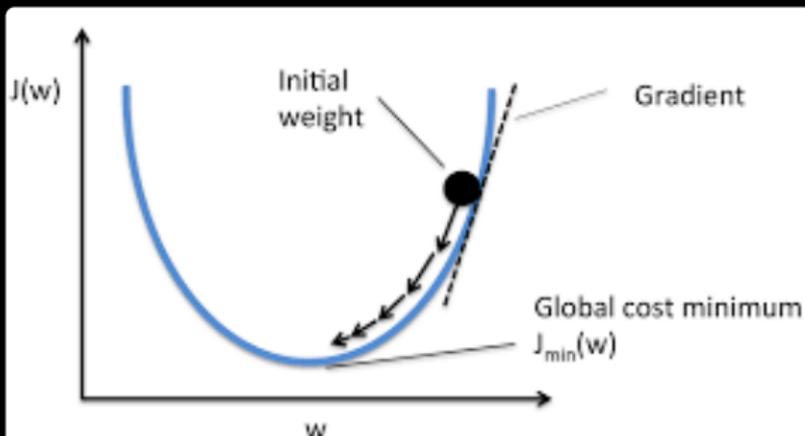
- We leave it to mathematicians!
- ... or try to understand it!

Back-propagation

- We leave it to mathematicians!
- ... or try to understand it!
- The wrong output induces an error!
- We want to minimize the error!

Back-propagation

- We leave it to mathematicians!
- ... or try to understand it!
- The wrong output induces an error!
- We want to minimize the error!
- High school mathematics: We can follow the *derivative* of the error function!
- We might converge to a local minimum!



The devil is in the details!

- Neurons also use activation functions!
 - It's just a transform of their output that increases their power!

The devil is in the details!

- Neurons also use activation functions!
 - It's just a transform of their output that increases their power!
- Neural networks are prone to “cheating”!
 - They tend to simply memorize the training data to cheat!
 - Several techniques proposed to avoid this phenomenon (called *overfitting*)!

The devil is in the details!

- Neurons also use activation functions!
 - It's just a transform of their output that increases their power!
- Neural networks are prone to “cheating”!
 - They tend to simply memorize the training data to cheat!
 - Several techniques proposed to avoid this phenomenon (called *overfitting*)!
- Vanishing/exploding gradients
 - The training algorithm can be unstable if the network is not correctly designed!
 - This was among the most important issues during the early years!

Dangers!

Artificial General Intelligence

- Currently, AI can only solve very specific problems (“weak AI”)
- In contrast, “strong AI” (or Artificial General Intelligence) aims to perform any task that a human can!
- Such machines will probably have “consciousness”.
- How far we are from that?

Artificial General Intelligence

'I fear that AI may replace humans altogether. If people design computer viruses, someone will design AI that improves and replicates itself. This will be a new form of life that outperforms humans.'

Stephen Hawking

'This is extremely important. The danger of AI is much greater than the danger of nuclear warheads. If humanity collectively decides that creating digital super-intelligence is the right move then we should do so very, very carefully.'

Elon Musk

Artificial General Intelligence

'I do not agree with Elon Musk about A.I. 'We shouldn't panic about it. The so-called control problem that Elon is worried about isn't something that people should feel is imminent.'

Bill Gates

'And I think people who are naysayers and try to drum up these doomsday scenarios I just, I don't understand it. It's really negative and in some ways I actually think it is pretty irresponsible.

In the next five to 10 years, AI is going to deliver so many improvements in the quality of our lives.'

Mark Zuckerberg

Artificial General Intelligence

We are just machines built from carbon instead of silicon.



So... stay tuned!