

# How AI Works & Why You Are Important

(in 45 minutes)

---

Lisa Yan

(with slides from Chris Piech)




# Announcements

---

- Final Project due at 6pm!
- Ceren's user study:
  - Consent forms will be handed out Lab 3
  - Email to come soon (with instructions)
  - Do this AFTER you've submitted your project 😊
- CS Bridge exit survey: fill out during Lab 3

Where is my robot?

# Sci-Fi Has Promised Me Robots




# House Cleaning Robot



# House Cleaning Robot



# Robots?



What is AI?



[suspense]

Computer programs



AI: The study and design of intelligent **agents**.

Better than chance



As well as humans



# Narrow Intelligence

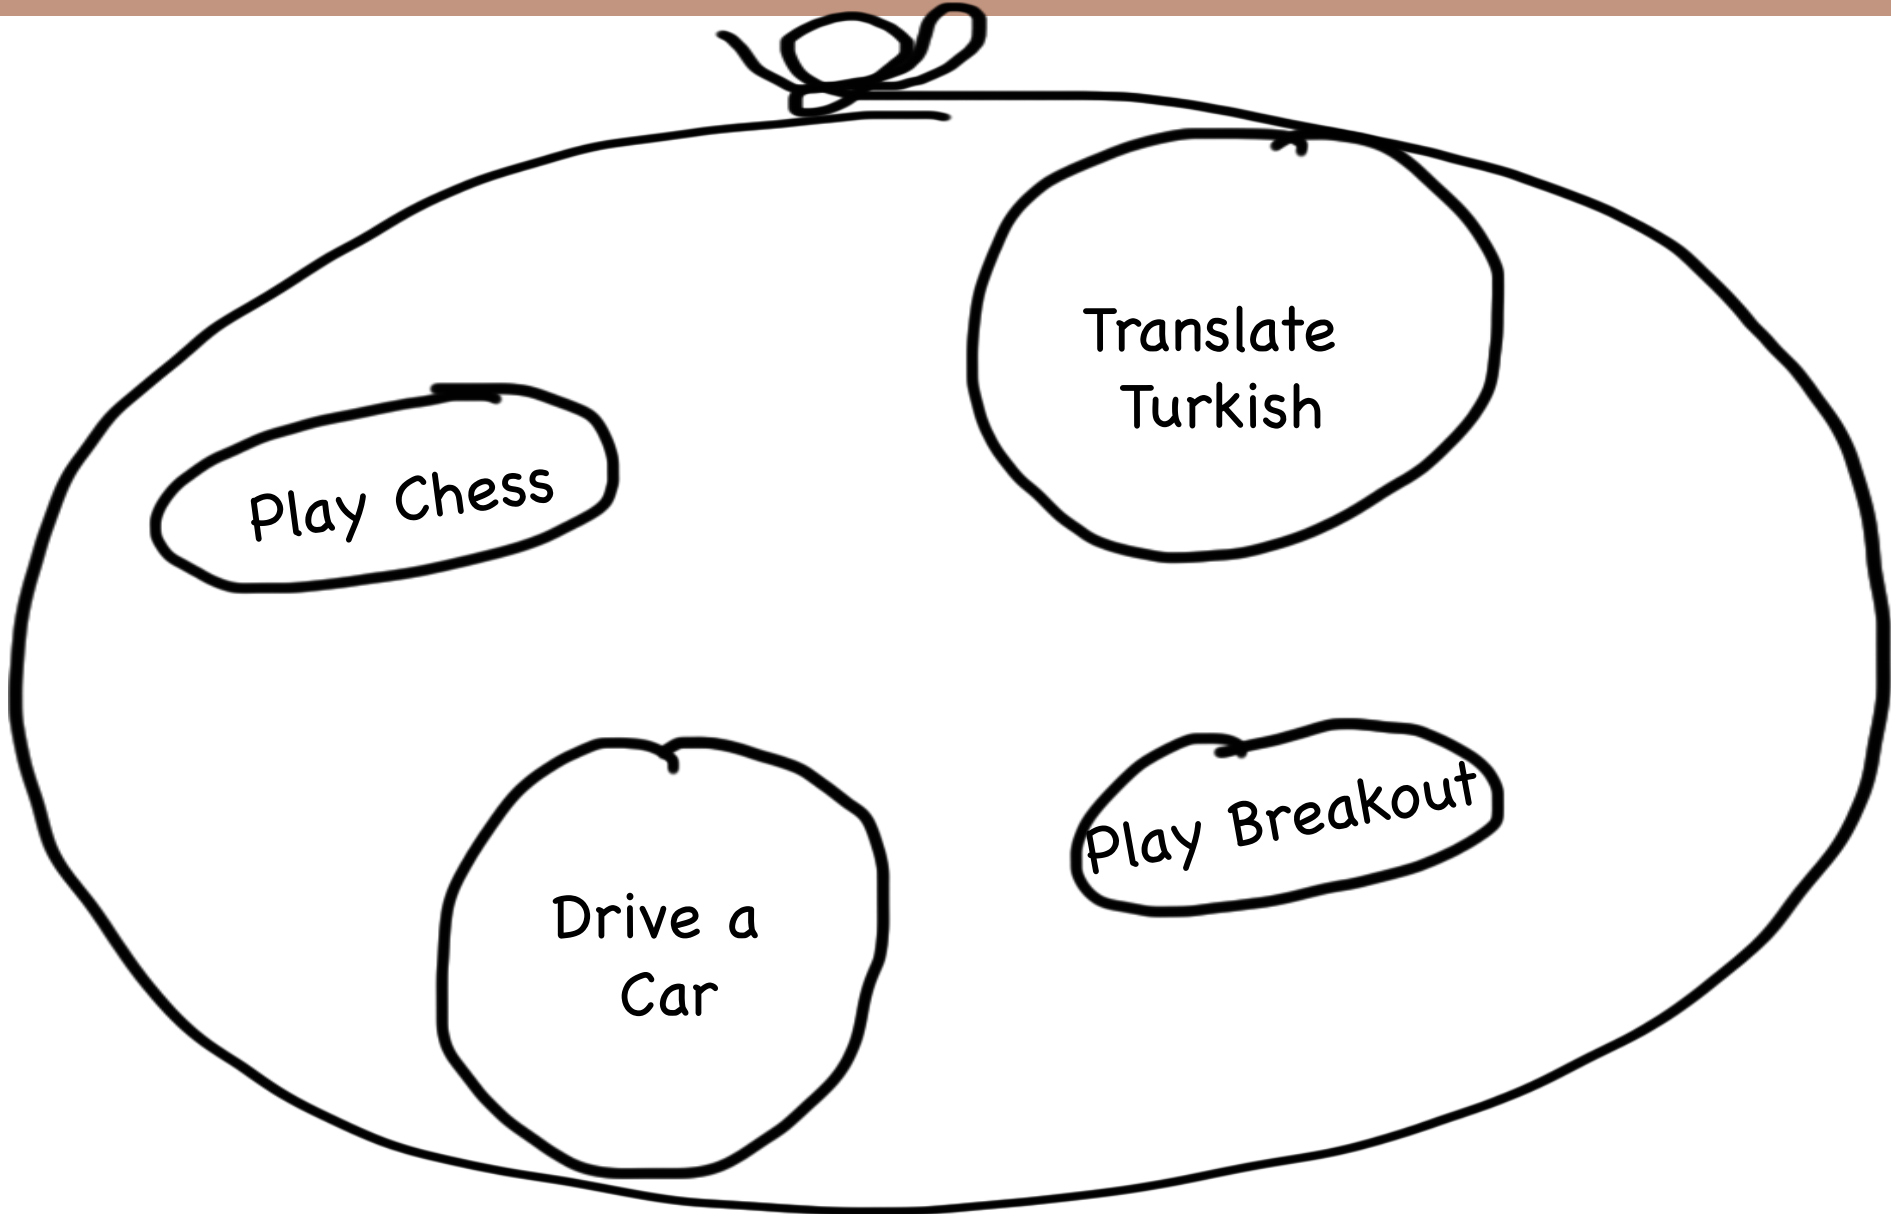
Play Chess

Translate  
Turkish

Drive a  
Car

Play Breakout

# General Intelligence



# Brief History



# Early Optimism 1950s



1952

# Early Optimism (1950s)

“Machines will be capable, within twenty years, of doing any work a man can do.”

–Herbert Simon, 1952

# Underwhelming Results (1950s–1980s)

*The spirit is willing but the flesh is weak.*



(Russian)




*The vodka is good but the meat is rotten.*

The world is too complex!



# Modern Game of AI



Istanbul Airport  YOLCU TİGARI KURUMSAL

## Evinize Hoş Geldiniz


Uçuş numarası, havayolu, şehir, havalimanı ile uçuşunuzu arayabilirsiniz.


**Giden** **Gelen**

Uçuş Numarası / Şehir / Havayolu Kodu



**IstanbulAirport Mobil Uygulaması Yayında!**  
Seyahatlerinizi eğlenceli ve kolay hale getirmek için ihtiyacınız olan her şey uygulamada.


**Giden** **Gelen**

Planlanan	Tahmini	Havayolu	Uçuş	Çıkış	Varış
14:10	17:00	 <b>TURKISH AIRLINES</b>	TK 0005	İSTANBUL	ŞİKA



2014


3:48   

**Turkish**  **English**


ISTANBUL AIRPORT  PASSENGER COMMERCIAL CORPORATE

## YOUR HOME WELCOME TO


FLIGHT NUMBER-AIRLINE-CITY-AIRPORT NAVIGATE WITH ARAYABİLİRSİNİZ


**V-THE**  FROM THE

FLIGHT NUMBER / CITY / AIRLINE CODE



**OUR** 

**IstanbulAirport MOBILE APP IS ONLINE!**  
YOUR TRAVELS FUN AND EASY TO MAKE YOU NEED HE' IN PRACTICE

**V-THE**  FROM THE

PLANNED	THE ESTIMATED	AIRLINE	FLIGHT	DEPART	THE RE
14:10	17:00	 <b>TURKISH AIRLINES</b>	TK 0005	İSTANBUL	ŞİKA

# Told Speech Was 30 Years Out




Almost perfect...

# The Last Remaining Board Game



# Computers Making Art



# Self-driving Cars



What is going on?

[more suspense]

Story of Modern AI:

Focus on one problem



# Make a Harry Potter Sorting Hat



# Classification



That is a picture  
of a **one**



# Classification



That is a picture  
of a **zero**



# Classification




That is a picture  
of an **zero**



\* It doesn't have to be  
correct all of the time

Can you do it?

# What number is this?



# What number is this?



# How about now?

What a computer sees

0	0	1	0	1	0	1	0	0	0	1	1	1	0	1
1	0	0	1	0	1	1	1	0	1	0	0	0	0	0
1	1	1	0	1	0	0	1	1	0	0	1	0	1	0
1	1	1	1	1	0	0	0	0	0	1	1	0	1	1
0	0	0	1	1	0	0	1	0	0					
1	0	0	1	1	0	0	0	1	0					
1	1	0	1	1	0	0	1	1	0					
1	0	1	0	0	1	0	0	1	0					
0	0	0	0	1	0	1	0	1	1					
0	1	1	0	0	0	0	0	1	1					
0	0	1	0	1	1	1	0	0	0					
0	1	1	1	0	1	0	0	1	0					
1	1	0	0	0	0	0	0	0	0					
0	0	0	0	0	0	0	0	1	1					
0	0	1	1	1	0	1	0	1	1					



What a human sees



# Why is it easy for Humans?



About 30% of your cortex is used from vision  
3% is used to process hearing

# Very hard to Program



```
public class HarryHat extends ConsoleProgram {  
  
    public void run() {  
        println("Todo: Write program");  
    }  
  
}
```


Perhaps there is an insight?

# One Algorithm Hypothesis

Much of perception in the brain can be explained with a single learning algorithm.



# One Algorithm Hypothesis




Auditory cortex learns to see

[Roe et al., 1992]

[Andrew Ng]

# One Algorithm Hypothesis



Somatosensory cortex learns to see

[Metin & Frost, 1989]

[Andrew Ng]


# Sensor Representations



Seeing with your tongue



Human echolocation (sonar)



Haptic belt: Direction sense



Implanting a 3<sup>rd</sup> eye

# Two Great Ideas

Story of Modern AI:

**1. Artificial Neurons**

**2. Learn by Example**


Machine Learning




# 1. Artificial Neurons




# Neuron




# Neuron




# Neuron




# Neuron





# Neuron



# Some Inputs are More Important




# Artificial Neuron

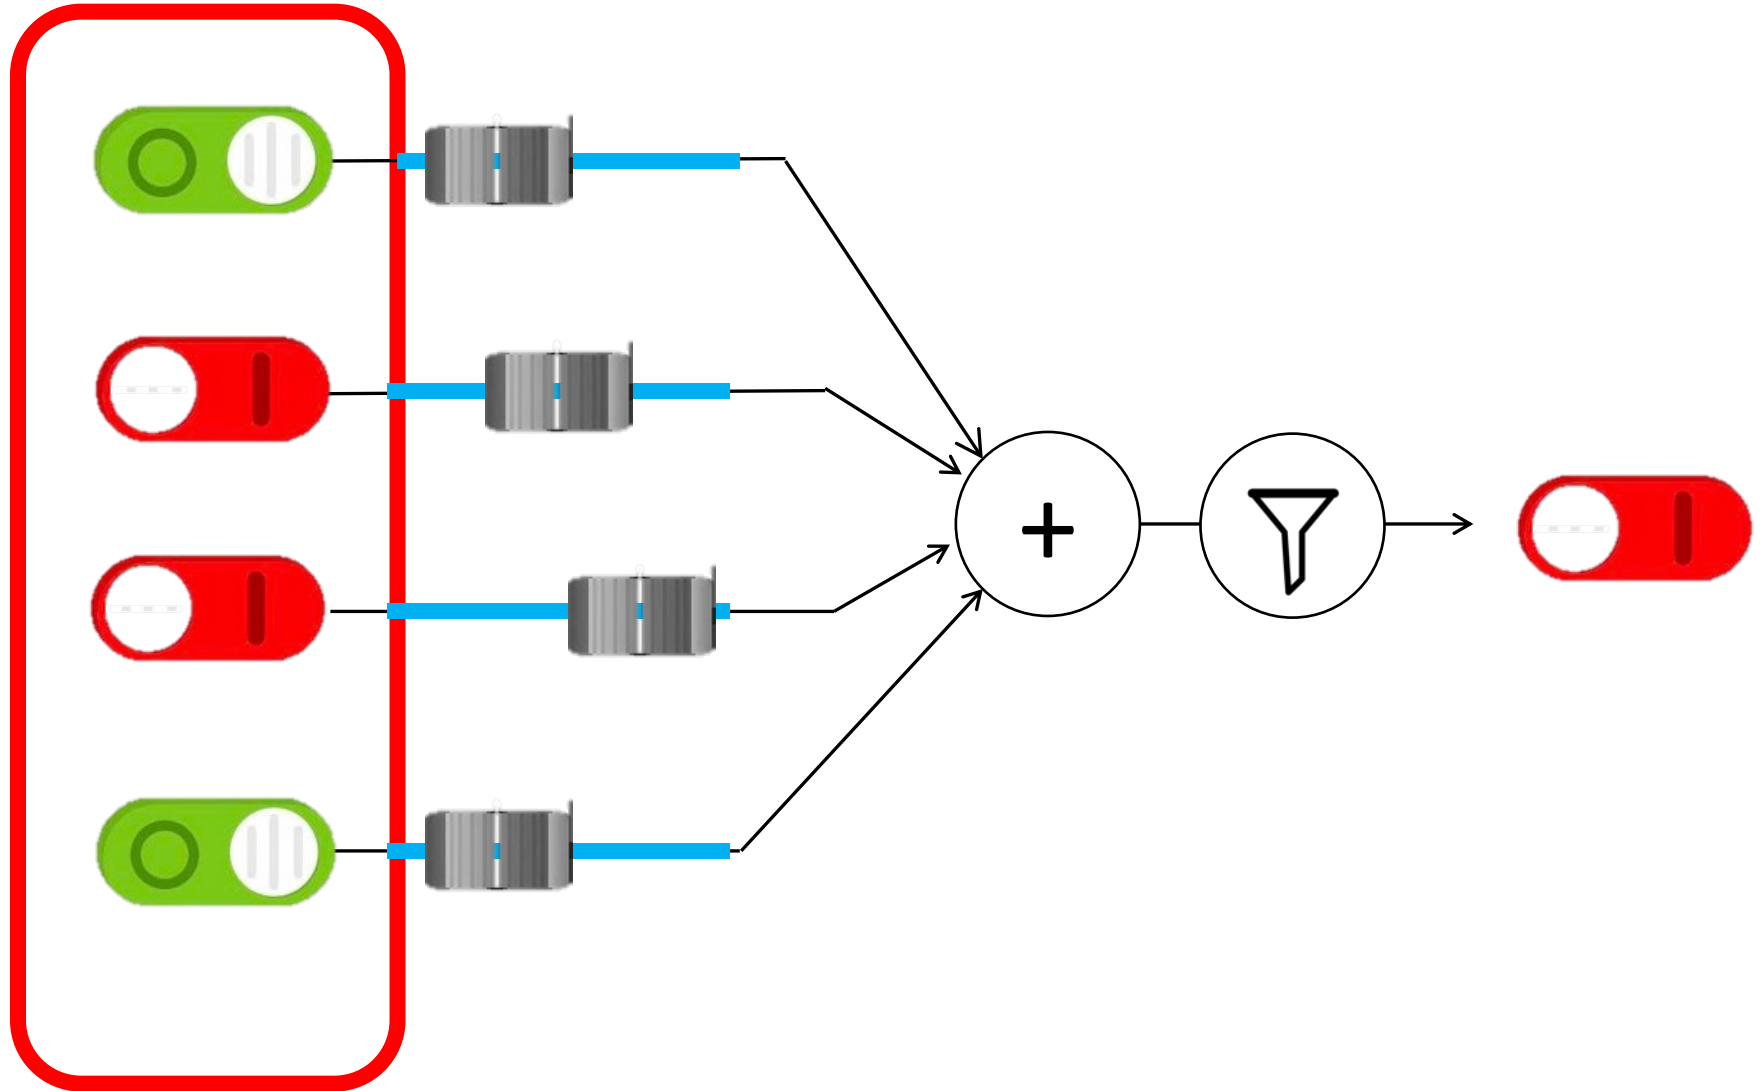





# Inputs




# Inputs




# Inputs




# Weights




# Weights



# Weights




# Weighted Sum



```
double weightedSum = 0;  
weightedSum += input0 * weight0;  
weightedSum += input1 * weight1;  
weightedSum += input2 * weight2;  
weightedSum += input3 * weight3;
```



# Filter and Output






# Biological Basis

A neuron




Your brain

(actually, probably someone else's brain)




???


# Put Many Together




# Put Many Together



Input Neurons

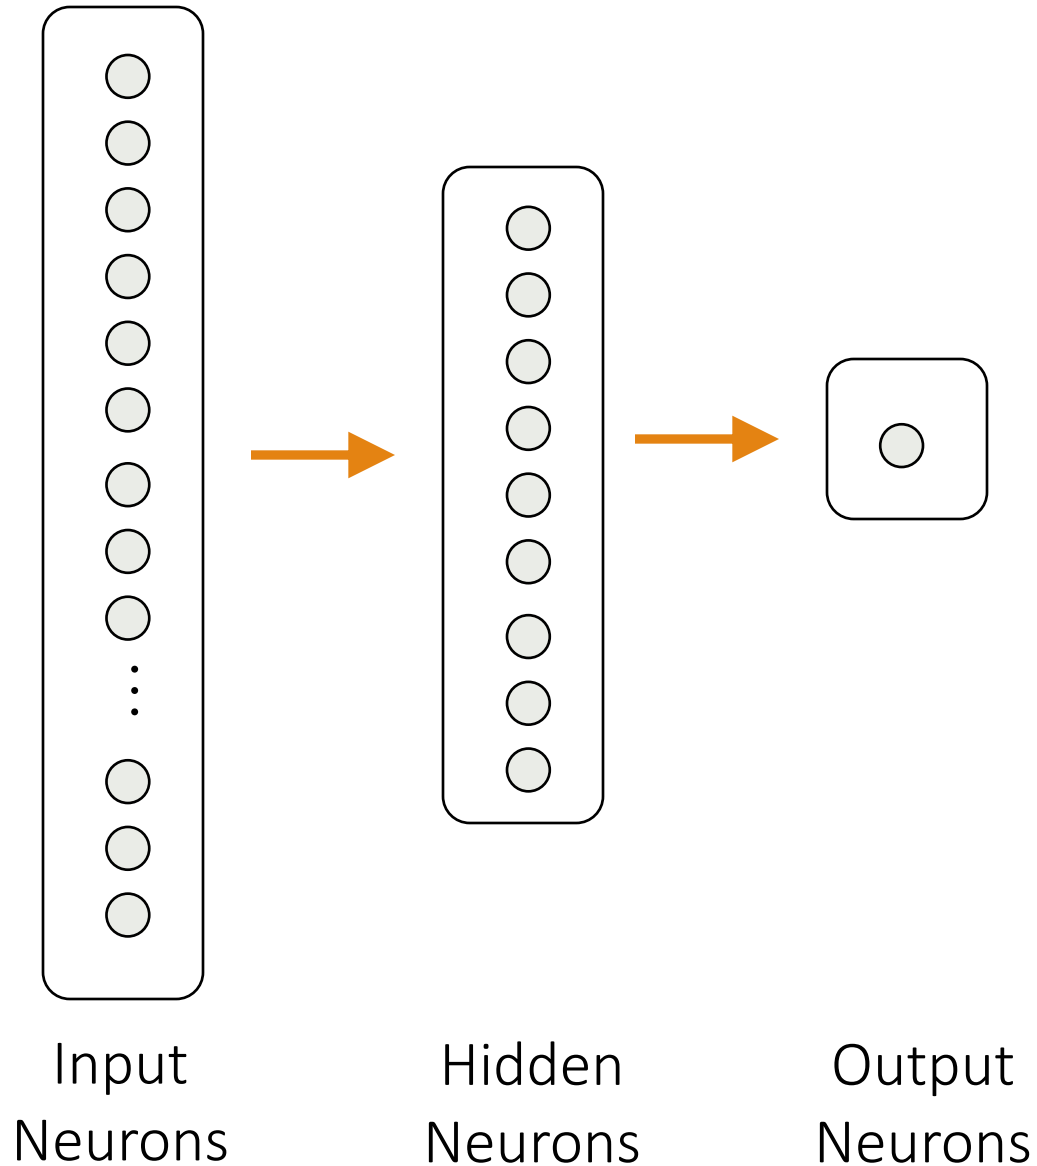


Hidden Neurons




Output Neurons

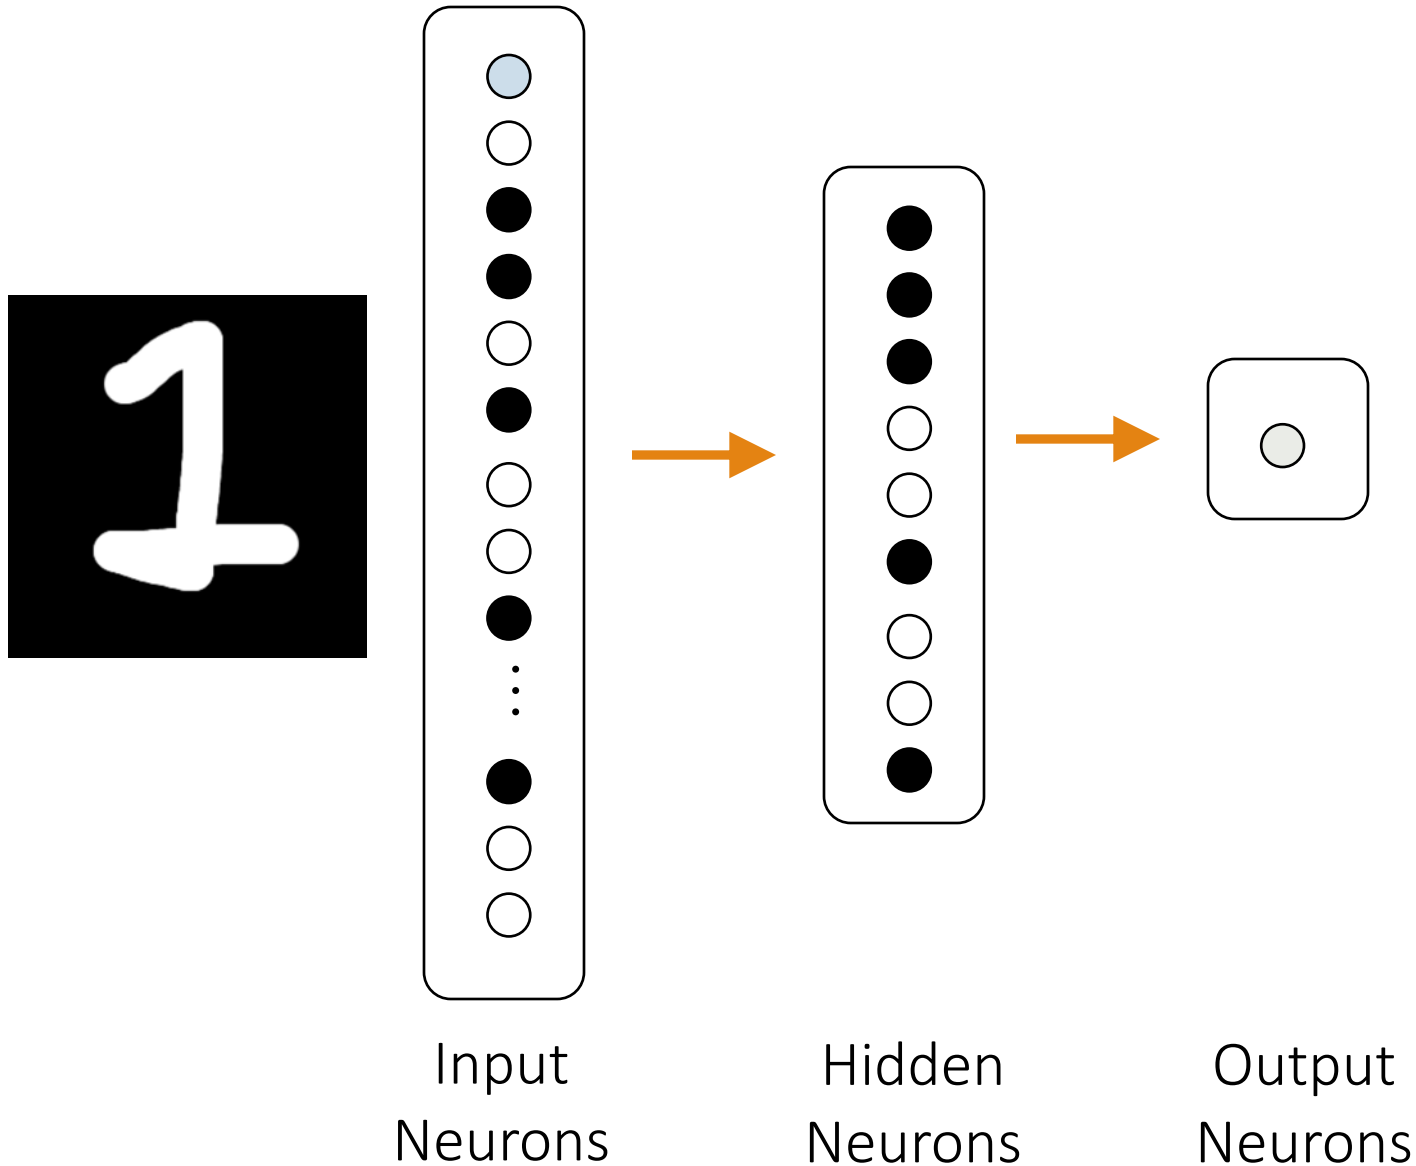
# Making a Prediction



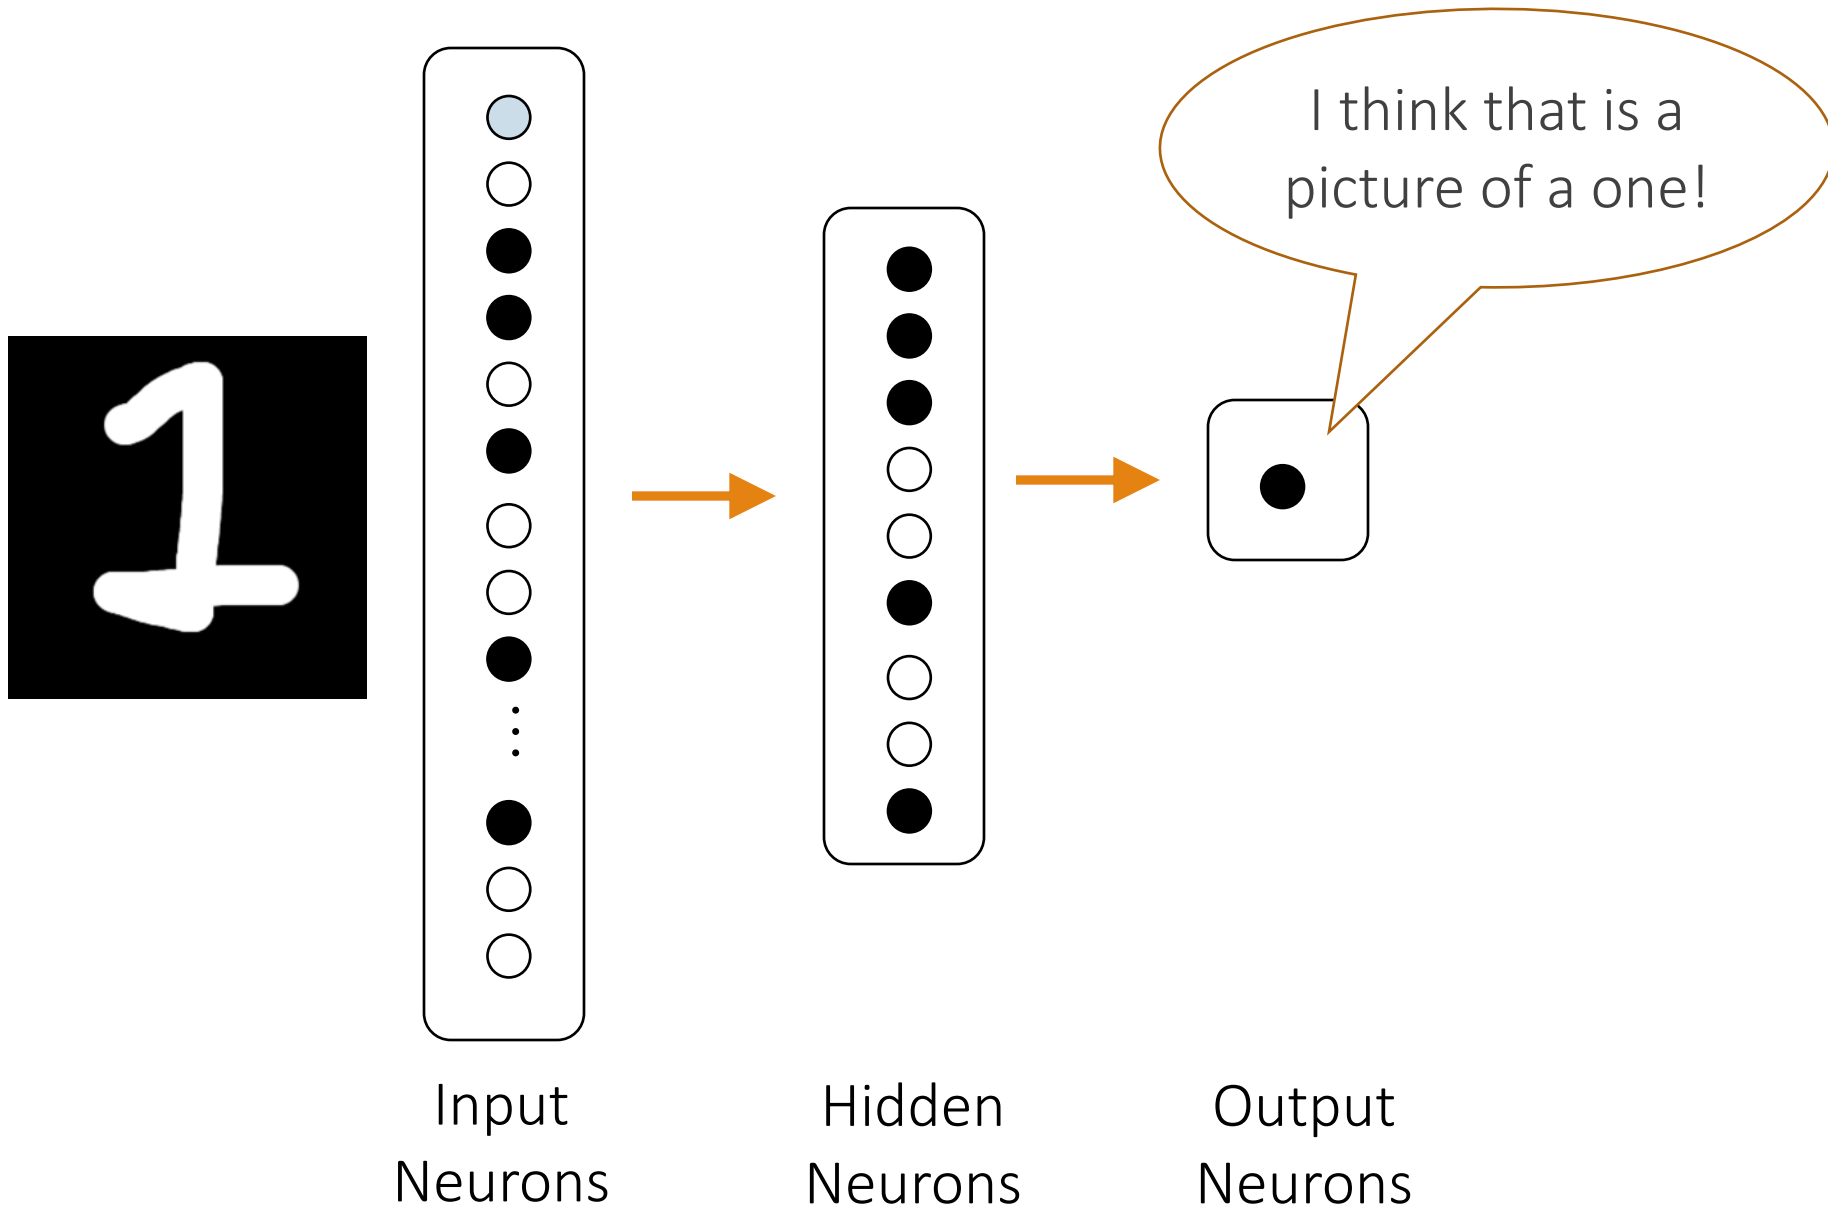
# Making a Prediction



# Making a Prediction




# Making a Prediction



# You Can Try It Yourself

Draw your number here



Downsampled drawing:

First guess:

Second guess:

## Layer visibility


Input layer

Convolution layer 1

Downsampling layer 1

Convolution layer 2



Downsampling layer 2



<http://scs.ryerson.ca/~aharley/vis/conv/>



# Great Idea: Artificial Neurons




# Two Great Ideas

**1. Artificial Neurons**

**2. Learn by Example**


## 2. Learn From Experience






Neural Networks gets their *intelligence* from its sliders (aka its weights)


# Neural Network



Input Neurons




Hidden Neurons



Output Neurons


# Neural Network



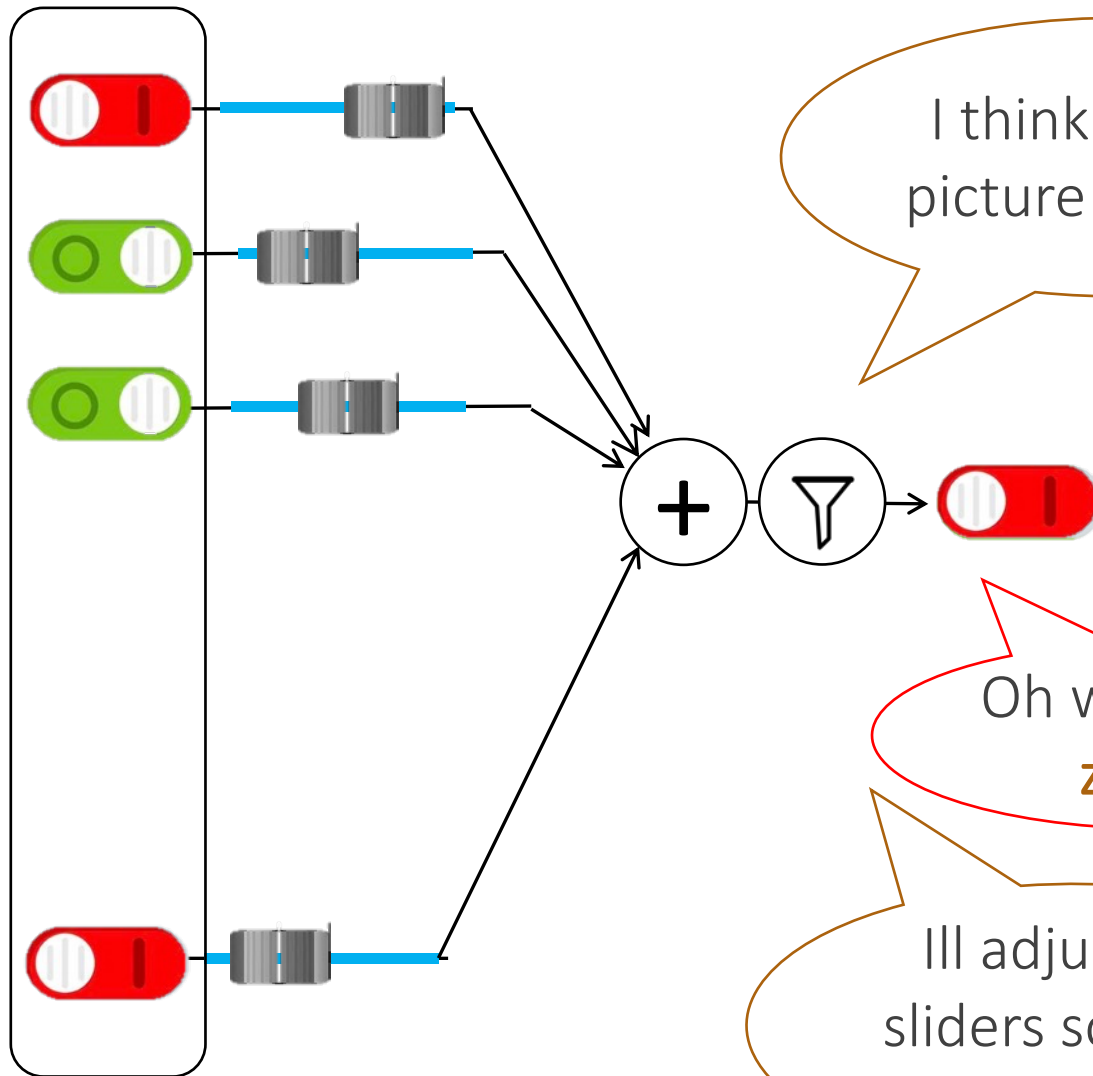
# Learn by Example

0 0 0 0 0 0 0 0 0 0 0 0 0 0 0  
1 1 1 1 1 1 1 1 1 1 1 1 1 1 1  
2 2 2 2 2 2 2 2 2 2 2 2 2 2 2  
3 3 3 3 3 3 3 3 3 3 3 3 3 3 3  
4 4 4 4 4 4 4 4 4 4 4 4 4 4 4  
5 5 5 5 5 5 5 5 5 5 5 5 5 5 5  
6 6 6 6 6 6 6 6 6 6 6 6 6 6 6  
7 7 7 7 7 7 7 7 7 7 7 7 7 7 7  
8 8 8 8 8 8 8 8 8 8 8 8 8 8 8  
9 9 9 9 9 9 9 9 9 9 9 9 9 9 9





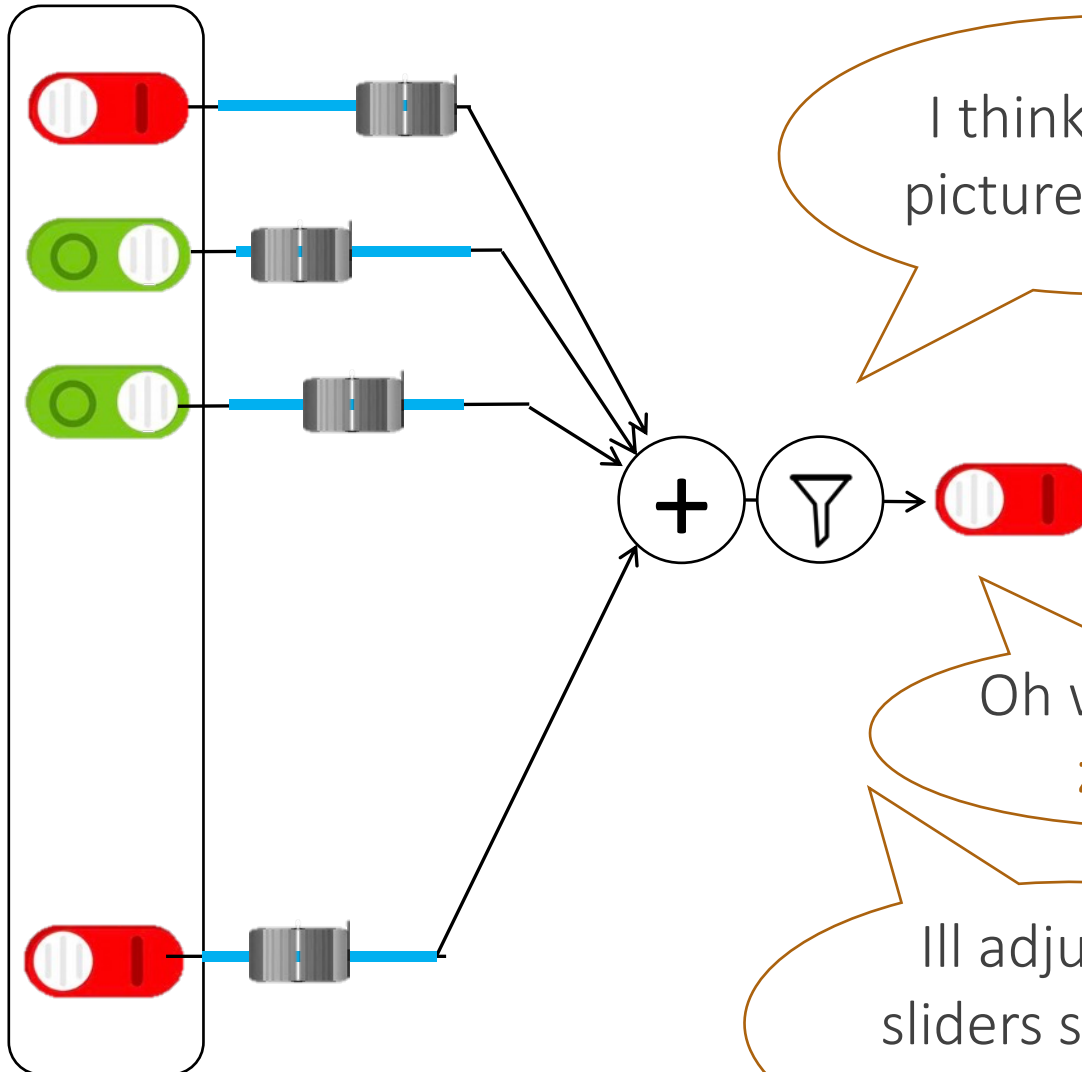




I think that is a picture of a **one**!

Oh what, it's a **zero**??


I'll adjust my sliders so that I do better.




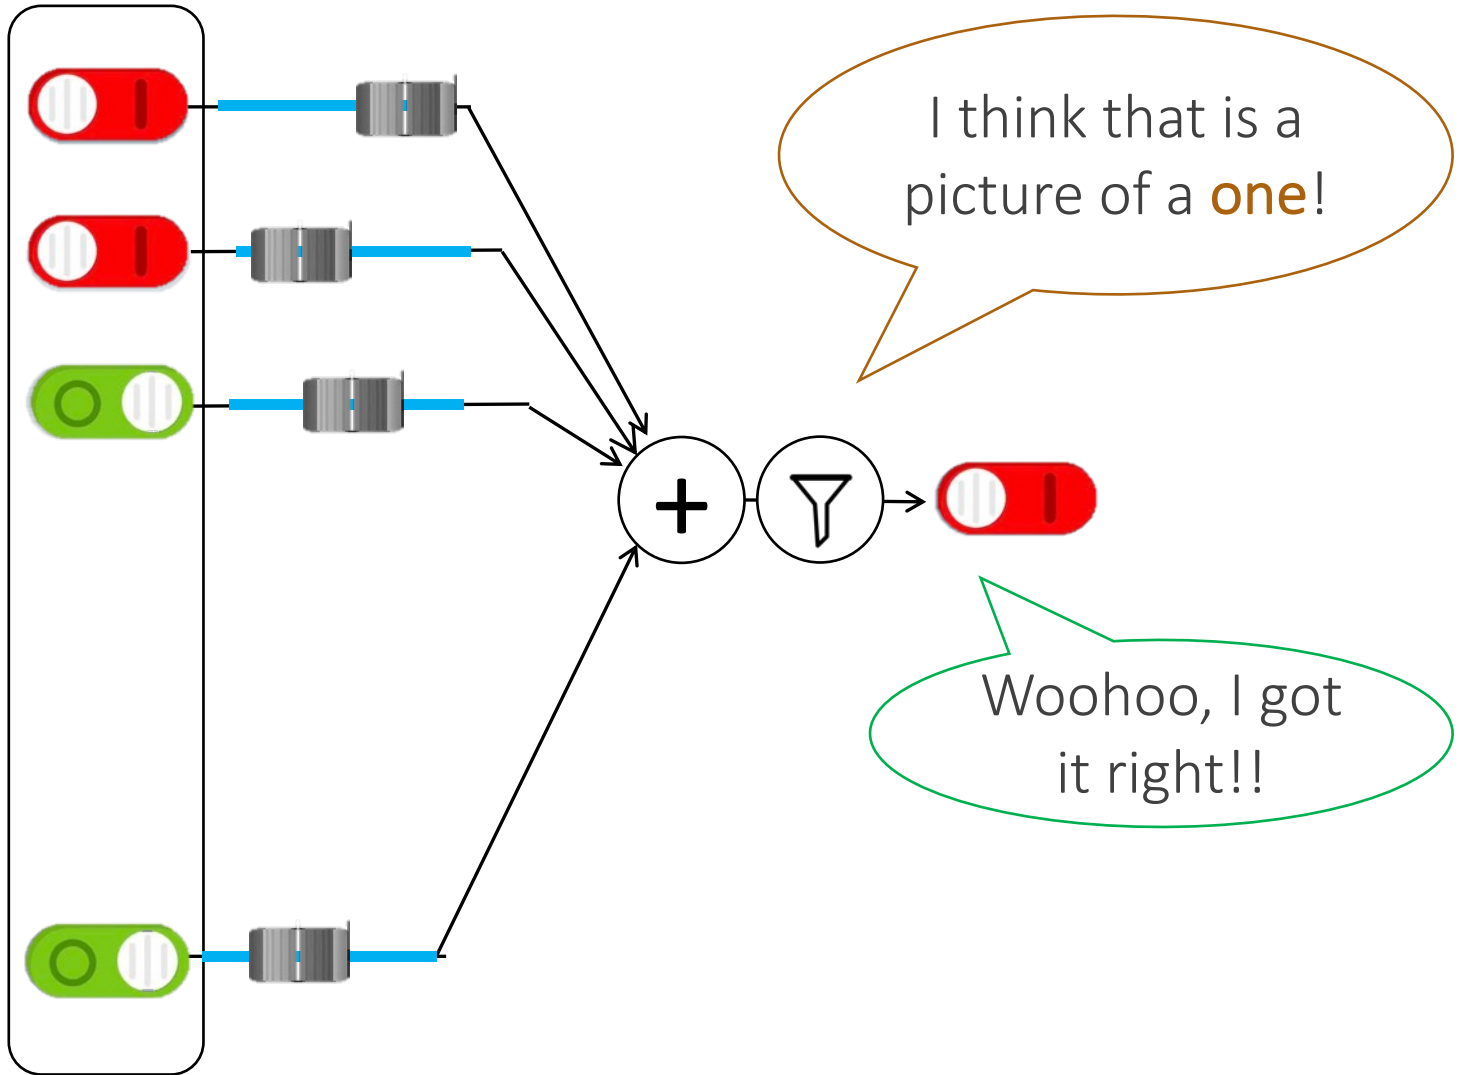
I think that is a picture of a **one**!


Oh what, it's a **zero**??

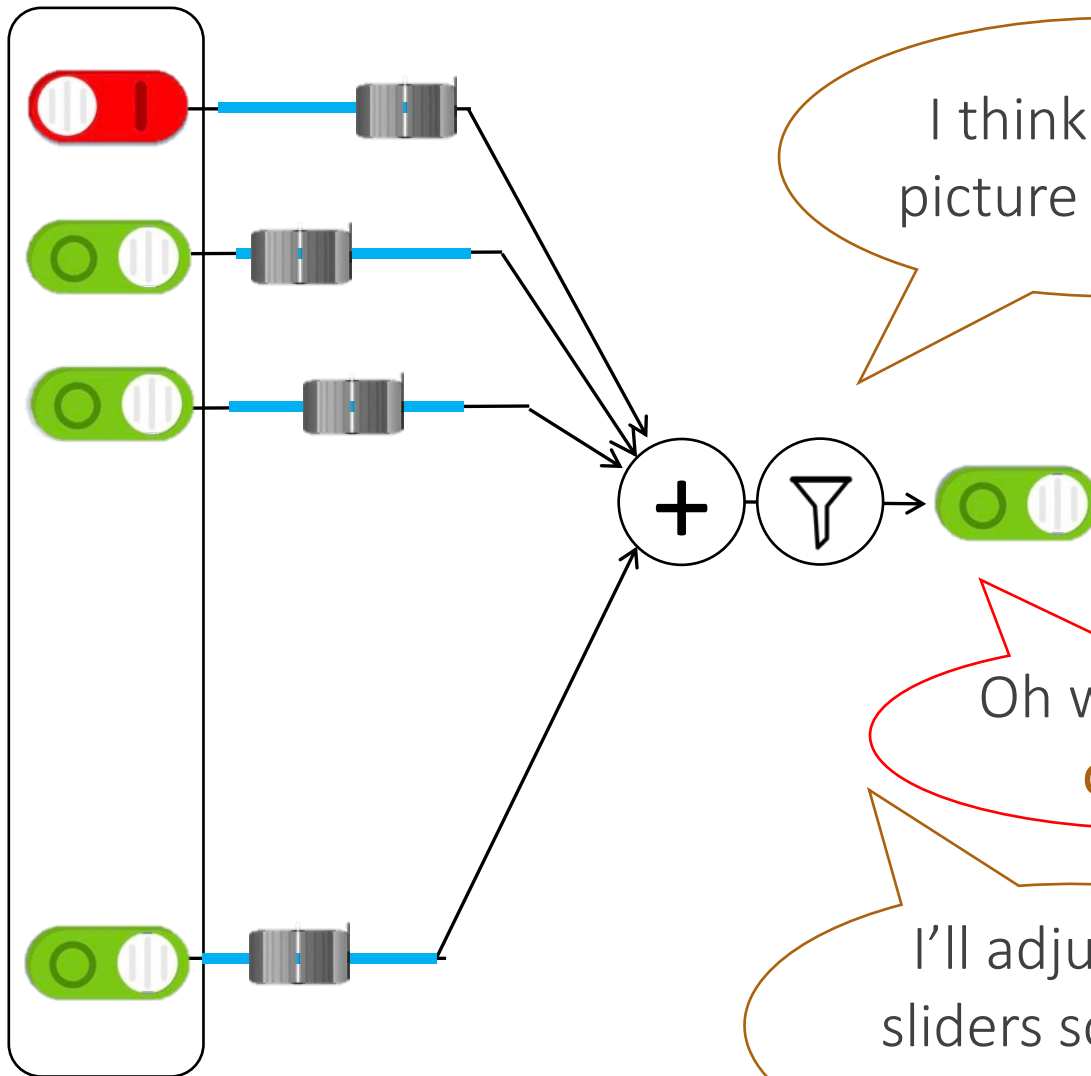
I'll adjust my sliders so that I do better.







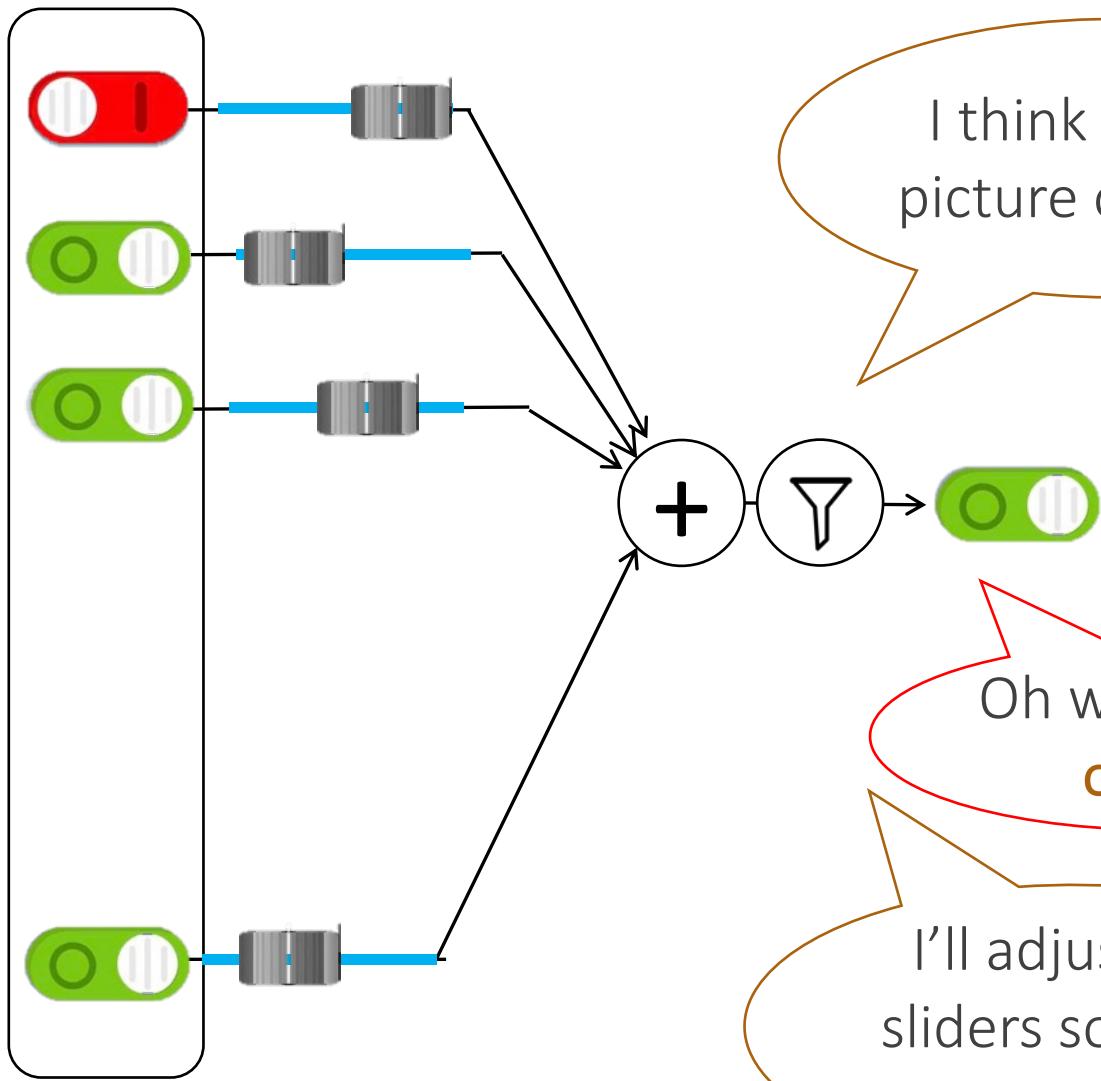




I think that is a picture of a **zero**!

Oh what, it's a **one**??

I'll adjust my sliders so that I do better.



I think that is a picture of a **zero**!


Oh what, it's a **one**??

I'll adjust my sliders so that I do better.





# Train on Faces




# Visualize the Sliders




Training set: Aligned images of faces.



object models



object parts  
(combination  
of edges)



edges



pixels

Woah... that's like a brain...

True.

# ImageNet Decomposition

...

smoothhound, smoothhound shark, *Mustelus mustelus*

American smooth dogfish, *Mustelus canis*

Florida smoothhound, *Mustelus norrisi*

whitetip shark, reef whitetip shark, *Triaenodon obesus*

Atlantic spiny dogfish, *Squalus acanthias*

Pacific spiny dogfish, *Squalus suckleyi*

hammerhead, hammerhead shark

smooth hammerhead, *Sphyrna zygaena*

smalleye hammerhead, *Sphyrna tudes*

shovelhead, bonnethead, bonnet shark, *Sphyrna tiburo*

angel shark, angelfish, *Squatina squatina*, monkfish

electric ray, crampfish, numbfish, torpedo

smalltooth sawfish, *Pristis pectinatus*

guitarfish

rougthead stingray, *Dasyatis centroura*

butterfly ray

eagle ray

spotted eagle ray, spotted ray, *Aetobatus narinari*

cownose ray, cow-nosed ray, *Rhinoptera bonasus*

manta, manta ray, devilfish

Atlantic manta, *Manta birostris*


devil ray, *Mobula hypostoma*

grey skate, gray skate, *Raja batis*


little skate, *Raja erinacea*

...

## Stingray



## Mantaray



0.005%

Random guess

1.5%

Pre Neural Networks (2012)

43.9%

GoogLeNet (2015)

73.1%


2017

Le, et al., *Building high-level features using large-scale unsupervised learning*. ICML 2012

Szegedy et al, *Going Deeper With Convolutions*, CVPR 2015

<http://image-net.org/challenges/LSVRC/2017/results>

# Google Brain



1 Trillion Artificial Neurons

(btw, human brains have 1 billion neurons)



# A Neuron That Fires When It Sees Cats




Top stimuli from the test set



Optimal stimulus  
by numerical optimization

**HIRE THE SMARTEST PEOPLE IN THE  
WORLD**



**INVENT CAT DETECTOR**

[meme-generator.net](http://meme-generator.net)

# Other Neurons

Neuron 1



Neuron 2



Neuron 3



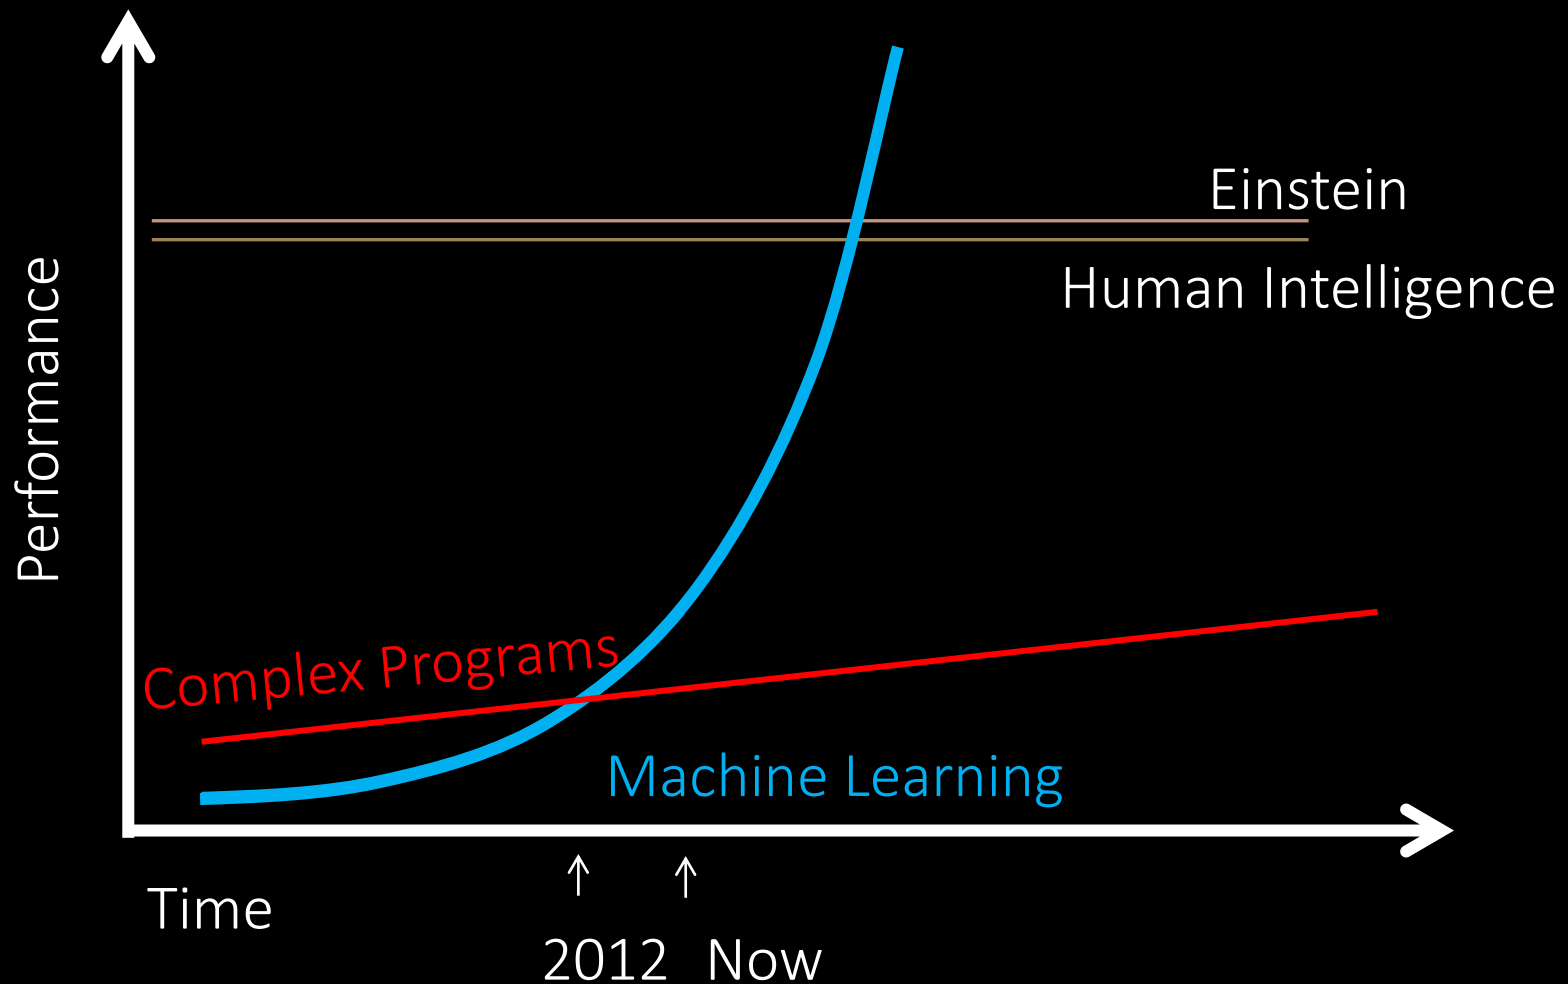
Neuron 4



Neuron 5




# The Future of AI



What's the catch?

# (1) Machine Learning Needs Data




# (1) Machine Learning Needs Data



(1) Get Data

Compiled by humans




(2) Train computer


Math and logic

# (1) Machine Learning Needs Data

Machine  
Translation



(1) Get Data



(2) Train computer





Language	# challenges
German	7
Czech	6.5
Russian	6
Finnish	5
French,Turkish,Chinese	3
German-Czech	1

One-time appearances:  
Hindi, Spanish, Lithuanian,  
Romanian, Latvian, Estonian, Gujarati

Because the engineers all speak English!



# (2) How can we explain decisions?




Visualize sliders

# (2) How can we explain decisions?



That is a picture  
of a one!

(probably fine)



# (2) How can we explain decisions?



This person is dangerous!

(not fine)



# (3) How can we make it fair?

$$\vec{\text{man}} - \vec{\text{woman}} \approx \vec{\text{king}} - \vec{\text{queen}}$$

$$\vec{\text{man}} - \vec{\text{woman}} \approx \vec{\text{computer programmer}} - \vec{\text{homemaker}}.$$


Should our data reflect society's systemic bias?

# The current challenge

*Understand* data.

Then train your model.

Then make your system usable for *real people*.



Tracy

- Quickly set up the bricks
- Spent 2 hrs implementing mouse interactor



Instructor

I can help you better!

Where is my robot?


... coming soon

# You can help



A little math

# You can help



Understand; don't assume



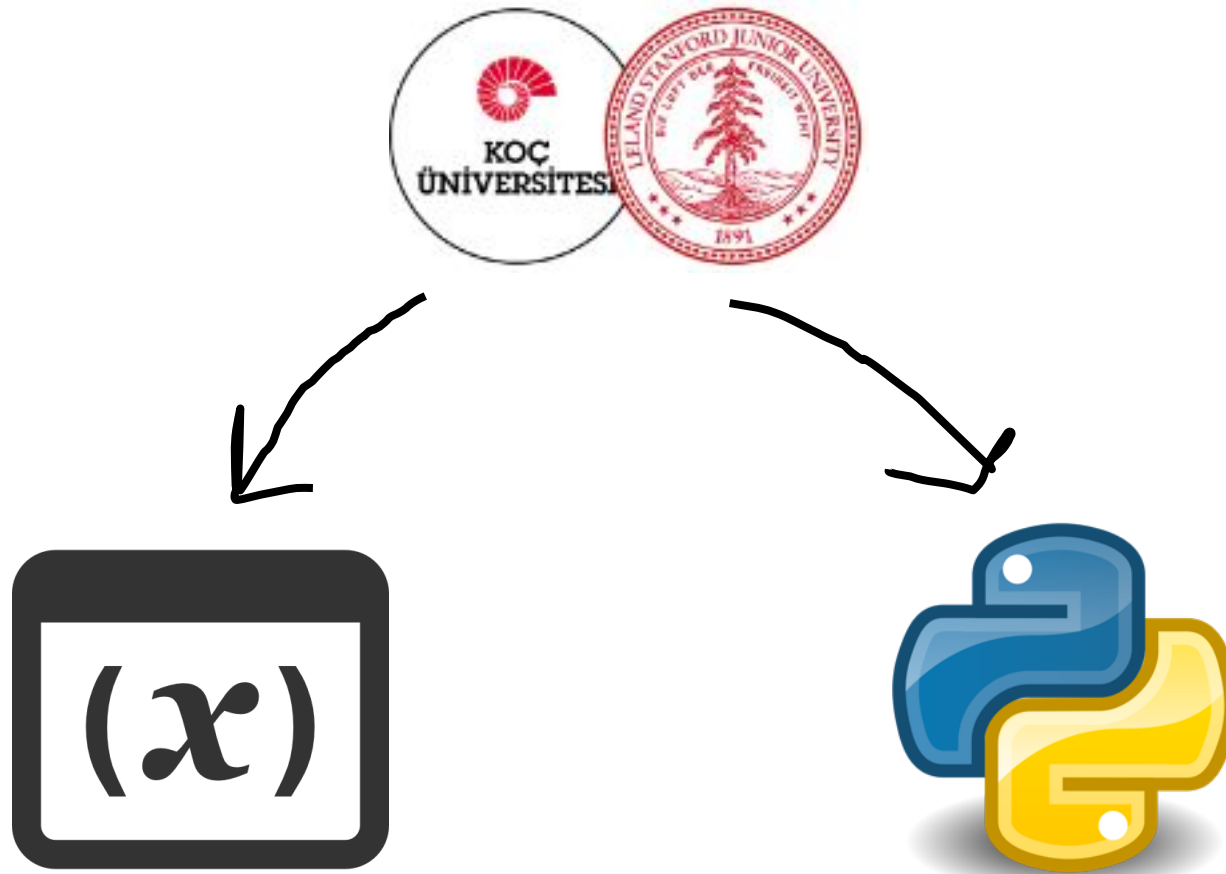
Can you do it?

(I explained to you the main components)

Not easy...

But yes. You can.

# Path to AI




Why?

# Closest Thing To Magic



# It's Useful

1,000,000 more jobs than students by 2020

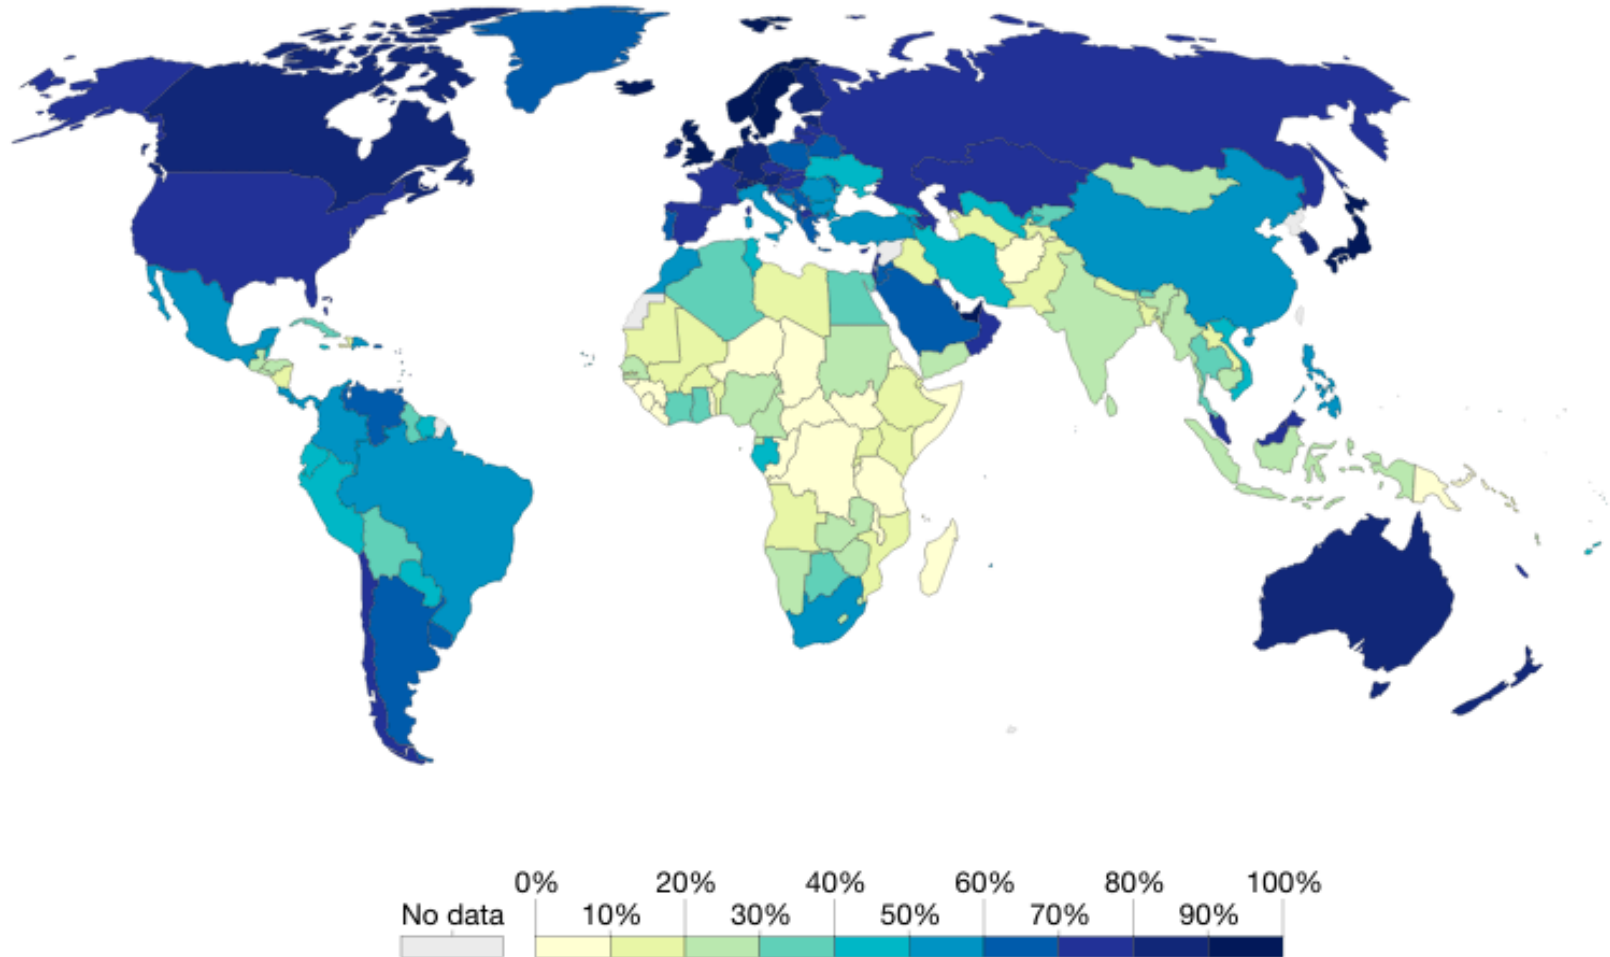


# It should fit your culture

## Share of the population using the Internet, 2015

All individuals who have used the Internet in the last 3 months are counted as Internet users. The Internet can be used via a computer, mobile phone, personal digital assistant, games machine, digital TV etc.

Our World  
in Data



**(make it so that) Everyone is Welcome**



A black gyroscope is shown spinning on a reflective surface. The gyroscope is a dark, metallic-looking object with a sharp point at the top and a sharp point at the bottom. It is tilted at an angle, and its reflection is visible on the surface below it. The background is a blurred, warm-toned surface.

The End?

Submit your projects by 6pm!