



FSAN/ELEG815: Statistical Learning

Gonzalo R. Arce

Department of Electrical and Computer Engineering
University of Delaware

Introduction

Acknowledgement

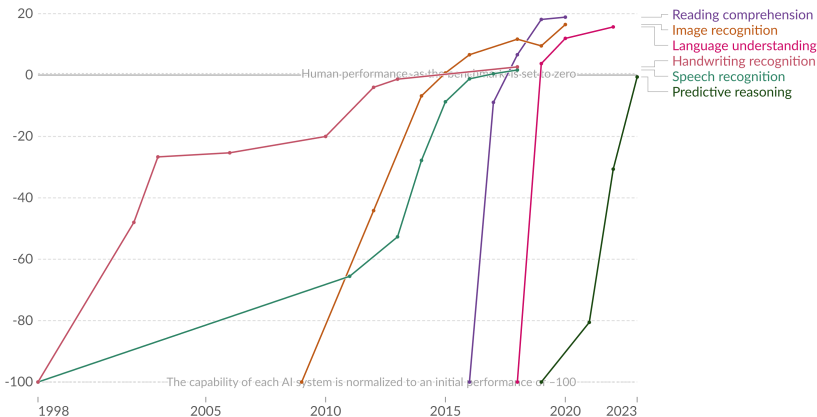
- ▶ Adapted from
 - ▶ *Artificial Intelligence: The Good, the Bad, and the Ugly* - Yaser Abu-Mostafa
 - ▶ *MIT Introduction to Deep Learning* - Alexander Amini

Deep Learning: Rapidly Moving field

Test scores of AI systems on various capabilities relative to human performance

Our World
in Data

Within each domain, the initial performance of the AI is set to -100. Human performance is used as a baseline, set to zero. When the AI's performance crosses the zero line, it scored more points than humans.



Data source: Kiela et al. (2023)

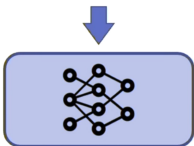
OurWorldInData.org/artificial-intelligence | CC BY

Note: For each capability, the first year always shows a baseline of -100, even if better performance was recorded later that year.

Where are we?

Generating Images from Natural Language

“A photo of an astronaut riding a horse.”



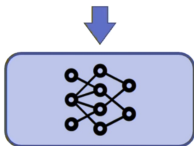
Ramesht+ arXiv 2022



Where are we?

Generating Language from Natural Language

“Write code in
TensorFlow to train a
neural network.”



OpenAI.



Certainly! Here is an example of how you can use TensorFlow to train a neural network for binary classification:

Copy code

```
import tensorflow as tf

# Load the data
(x_train, y_train), (x_test, y_test) = load_data()

# Define the model
model = tf.keras.Sequential([
    tf.keras.layers.Input(shape=x_train.shape[1:]),
    tf.keras.layers.Dense(16, activation='relu'),
    tf.keras.layers.Dense(1, activation='sigmoid')
])

# Compile the model
model.compile(optimizer='adam', loss='binary_crossentropy', metrics=
['accuracy'])

# Train the model
model.fit(x_train, y_train, epochs=10, batch_size=32, validation_data=
(x_test, y_test))
```

This code assumes that you have a function `load_data()` that returns the training and test datasets as tuples of NumPy arrays. The input shape of the model is determined by the shape of the training data, and the model has two dense layers with 16 and 1 units, respectively, with ReLU and sigmoid activations. The model is then compiled with an Adam optimizer and a binary cross-entropy loss function, and is trained using the `fit()` method.

I hope this helps! Let me know if you have any questions.

Where are Deep Learning and AI headed?

Deep Learning is revolutionizing so many fields.

You'll learn the fundamentals of this field and how you can build some of these incredible algorithms.

What is Deep Learning?

ARTIFICIAL INTELLIGENCE

Any technique that enables computers to mimic human behavior



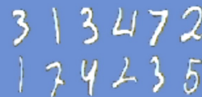
MACHINE LEARNING

Ability to learn without explicitly being programmed



DEEP LEARNING

Extract patterns from data using neural networks



Teaching computers how to **learn a task** directly from **raw data**

Outline

- ▶ The Science
- ▶ The Upside
- ▶ The Downside

Outline

- ▶ **The Science**
- ▶ The Upside
- ▶ The Downside

Brute-force approach to Artificial Intelligence

The '**Lookup Table**':

1. Store gigantic amount of information in a computer.
2. Look up the relevant information when someone asks.



[Credit: Jeopardy!]



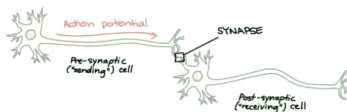
[Credit: Reuters]

Learning

Brute-force approach to Artificial Intelligence



The learning algorithm is built in:

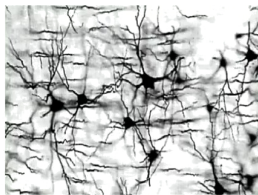


[Credit: Becoming Human]

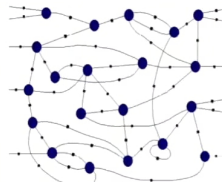
Learning from Data



How to imitate the brain?



[Credit: Alan Turing.net]



[Credit: Alan Turing.net]

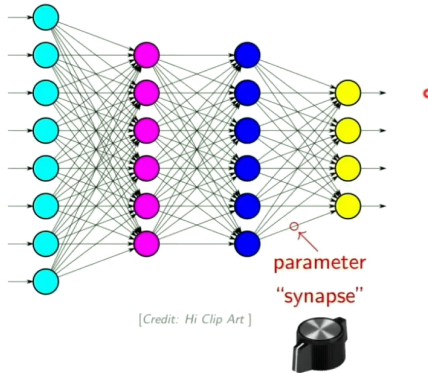


[Credit: wikimedia]



[Credit: dw.com]

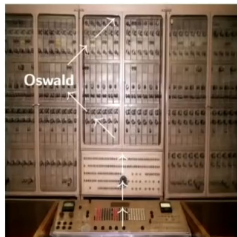
The Neural Network



How information is stored

Expert System

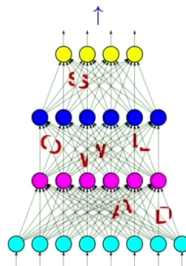
Oswald



Who killed Kennedy?

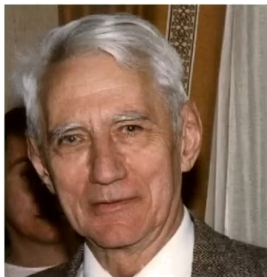
Neural Network

Oswald



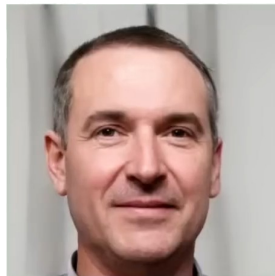
Who killed Kennedy?

Discriminative vs Generative

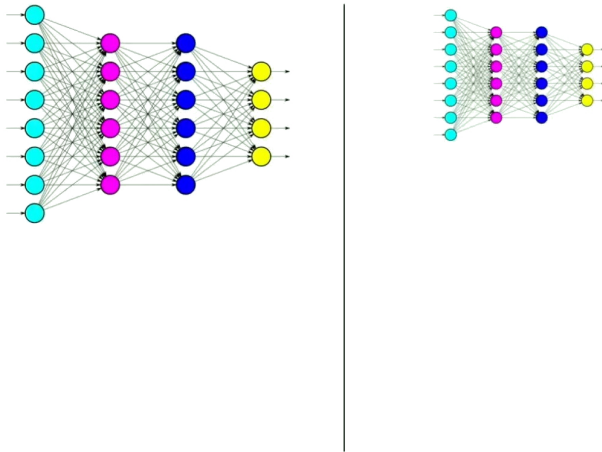


Claude Shannon

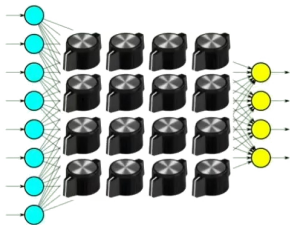
Create a face



Creating the Network vs Using the Network



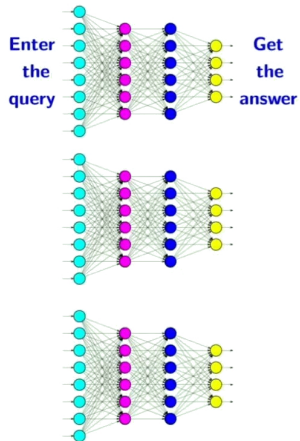
Creating the Network vs Using the Network



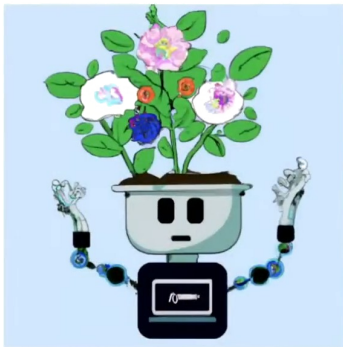
Training

1. Very intensive computation
2. Shapes the function

↑ Bias, Malice, ... ↑

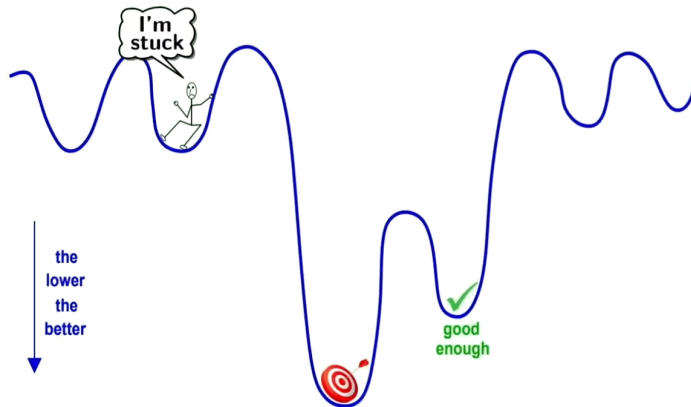


Where is AI going?

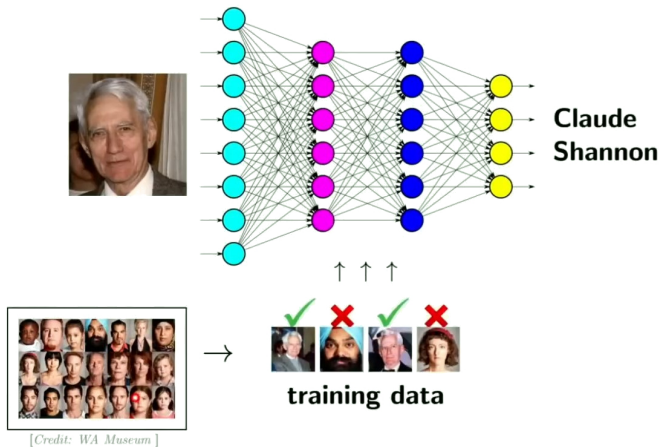


[Credit: Creativity Explorer]

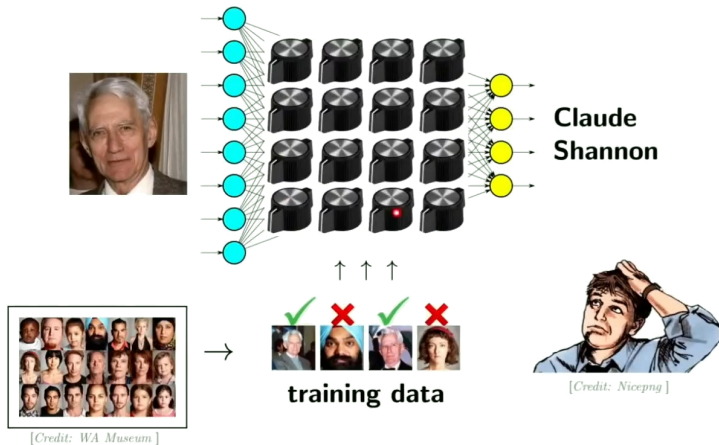
Lucky Break #1: Local Minima



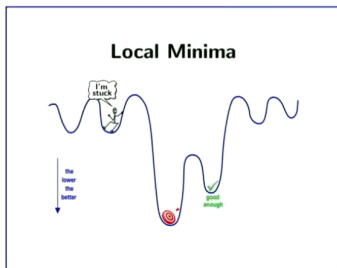
Lucky Break #2: Over-Parameterization



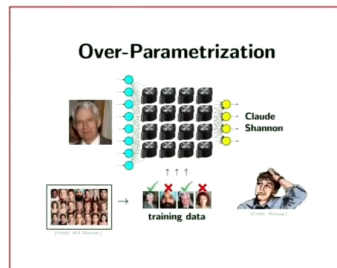
Lucky Break #2: Over-Parameterization



The two lucky breaks

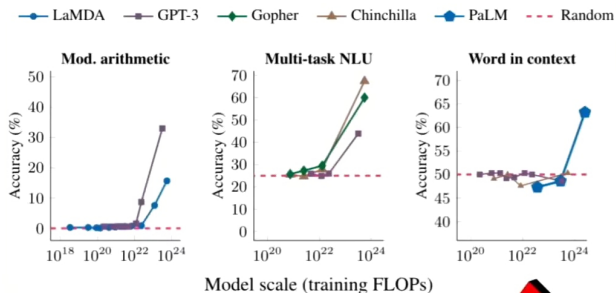


**Need Far Less
Computation**



**Need Far Less
Information**

More Luck: emergent Abilities



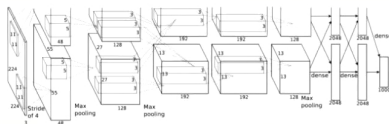
[J. Wei & Y Tay, Google (2022)]



Can You Cross The **Bridge**?
Can You Play **Bridge**?

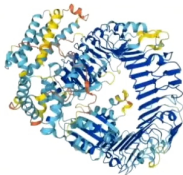
Academia → Industry

AlexNet (2012)



[A. Krizhevsky et al.]

AlphaFold (2021)



[DeepMind/Google]

↑ Academia ↑

← Industry →

GAN (2014)



[J. Goodfellow et al.]

ChatGPT (2022)



[OpenAI/Microsoft]

Outline

- ▶ The Science
- ▶ **The Upside**
- ▶ The Downside

Two things AI offers

1. Doing what we already do.
2. Doing what we cannot do.

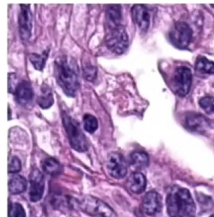
Two things AI offers

1. Doing what we already do.



[Credit: hearstaps]

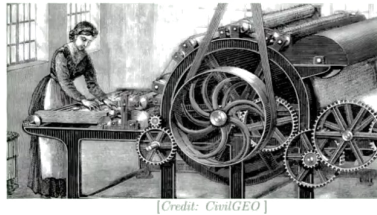
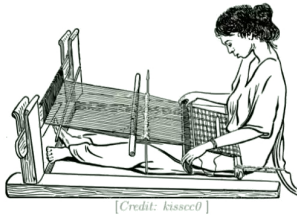
2. Doing what we cannot do.



[H. Zhou et al (2023)]

The Industrial Revolution

Relieving manual labor

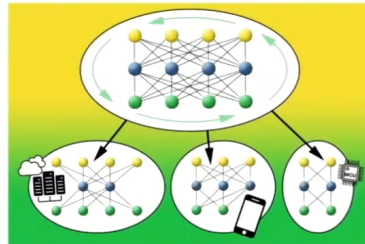


The AI Revolution

Relieving ^{mental} manual labor



[Credit: M&G]



[Credit: MIT]

The key difference

The pace

- The Industrial Revolution took hold in approximately 80 years

1760 - \approx 1840

- A.I. Revolution will have taken taken hold in less than 20 years

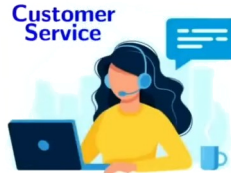
2012 - 2030?

By the year 2030

AI can automate any task that requires routine intelligence



[Credit: Vox]



[Credit: signpost]



[Credit: ailoitte]



[Credit: visme.co]

Outline

- ▶ The Science
- ▶ The Upside
- ▶ **The Downside**

The risks of AI

1. Malicious use of AI
2. Rogue AI systems
3. Social Impact

1. Malicious use of AI



[Credit: iStock]



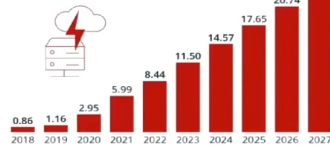
[Credit: Simplilearn]



[Credit: Adobe]

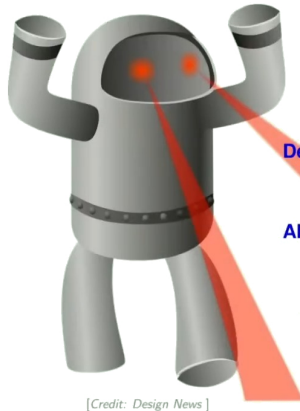
Cybercrime Expected To Skyrocket in the Coming Years

Estimated cost of cybercrime worldwide
(in trillion U.S. dollars)




[Credit: Statista]

2. Rogue AI systems



Three Issues:

- 
- Desire
 - Bundling human traits
 - Ability
 - Ignoring the logistics
 - Link
 - Intelligence vs. Control

General vs Machiavellian Intelligence

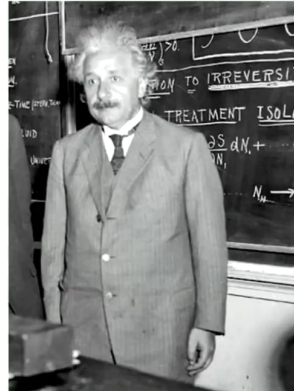
Studies That Relate Machiavellianism to Measures of General Intelligence

Correlation	Measurement	Reference
None	IQ	Christie & Geis (1970b)
None	IQ	Christie & Geis (1970b)
None	Verbal ability test	Christie & Geis (1970b)
None	General classification test	Christie & Geis (1970b)
None	MCAT	Christie & Geis (1970b)
None	Moore-Castore test	Singer (1964)
None	Guilford-Zimmerman test	Wrightman & Cook (1965)
None	Digit span	Steininger & Colsher (1979)
None	GPA	Ames & Kidd (1979)

Note. The first seven studies are reviewed by Christie and Geis (1970a, pp. 36–37). The two IQ tests were performed on separate samples of 115 college students and 218 preparatory school students, respectively. MCAT = Medical College Admissions Test; GPA = grade point average.

[D. Wilson et al (1996)]

*General Intelligence
has no correlation with
Machiavellianism*



3. Social Impact

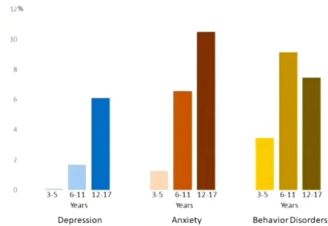


[Credit: PC Magazine]



[Credit: OrangeCrush]

Depression, Anxiety, Behavior Disorders, by Age



[Credit: CDC]