

Basics of Machine Learning & Artificial Intelligence

Decision Intelligence Confronts A.I.

Friday November 10, 2023

John Mount

Win Vector LLC

All refs/links:

https://github.com/WinVector/Examples/tree/main/decision_intel_ai

Disclaimer

- Produced to orient us on some current terms and claims
- A *greatly* simplified and abbreviated view and opinion
- Going to a bit of a “potted history”, opinion, and we gotta go fast to hit 15 minutes! (sorry!!!!)
- Trying to stay towards points that may be relevant to decision intelligence, quality, and general consequences
 - Talk goal: set up terms so the summary slide makes sense
 - “How to talk about it”



Outline

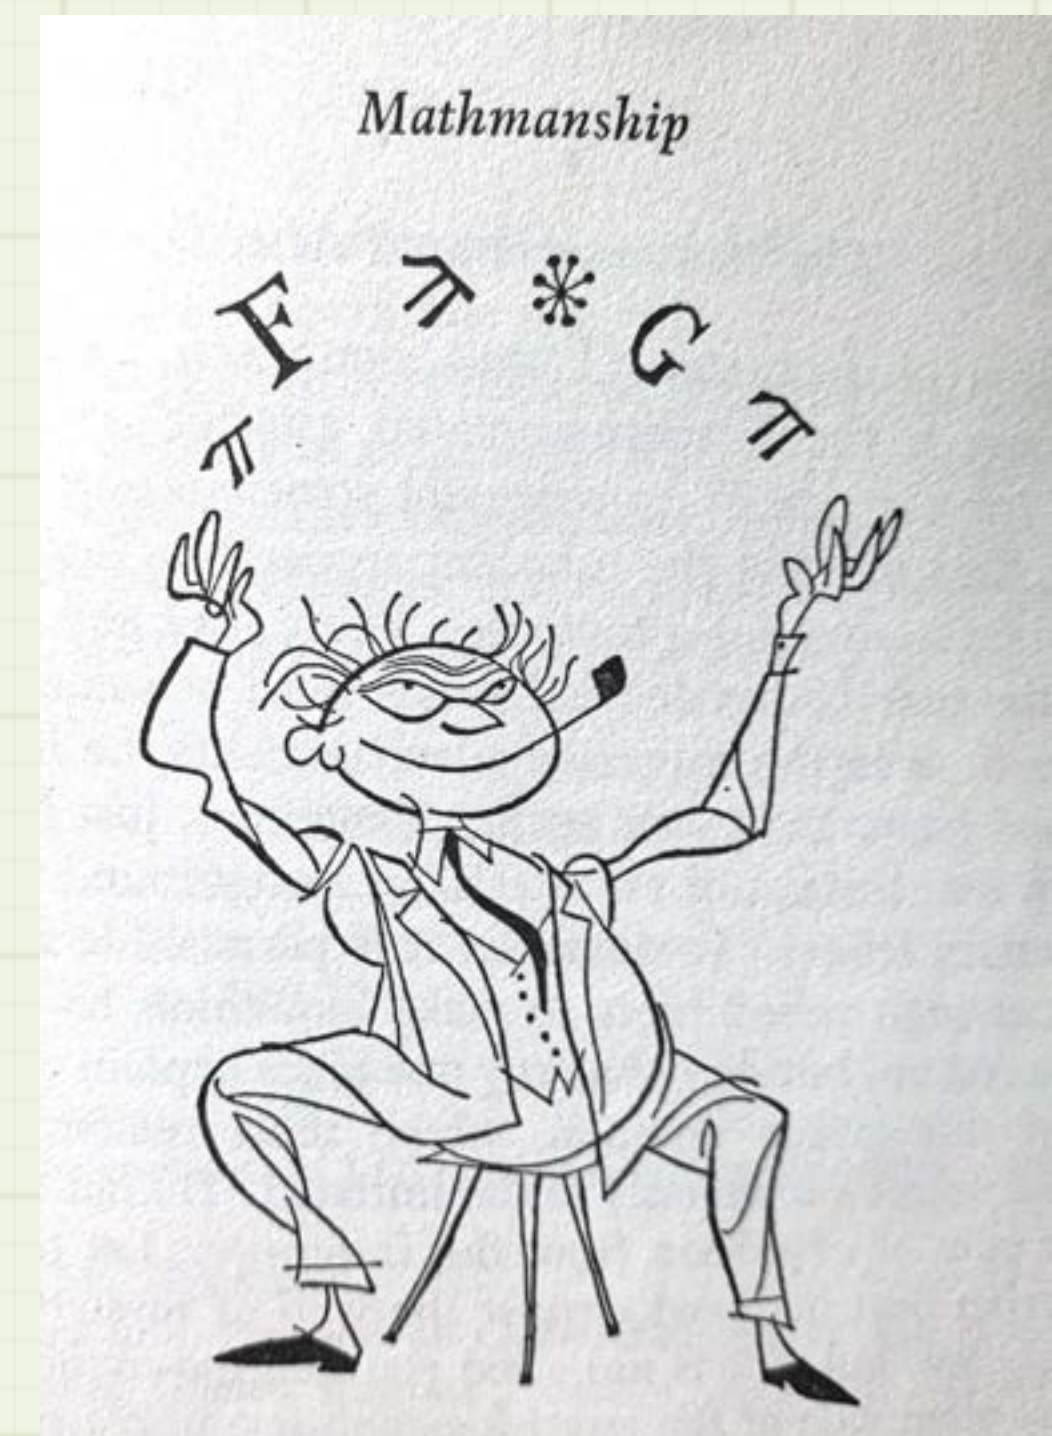
- **Statistical Foundations of Machine Learning** (the start)
- **Deep Learning** (evolutionary explosion)
- **Current AI** (amazing, but has some caveats)

Statistical Foundations of Machine Learning

Statistical Foundations of Machine Learning

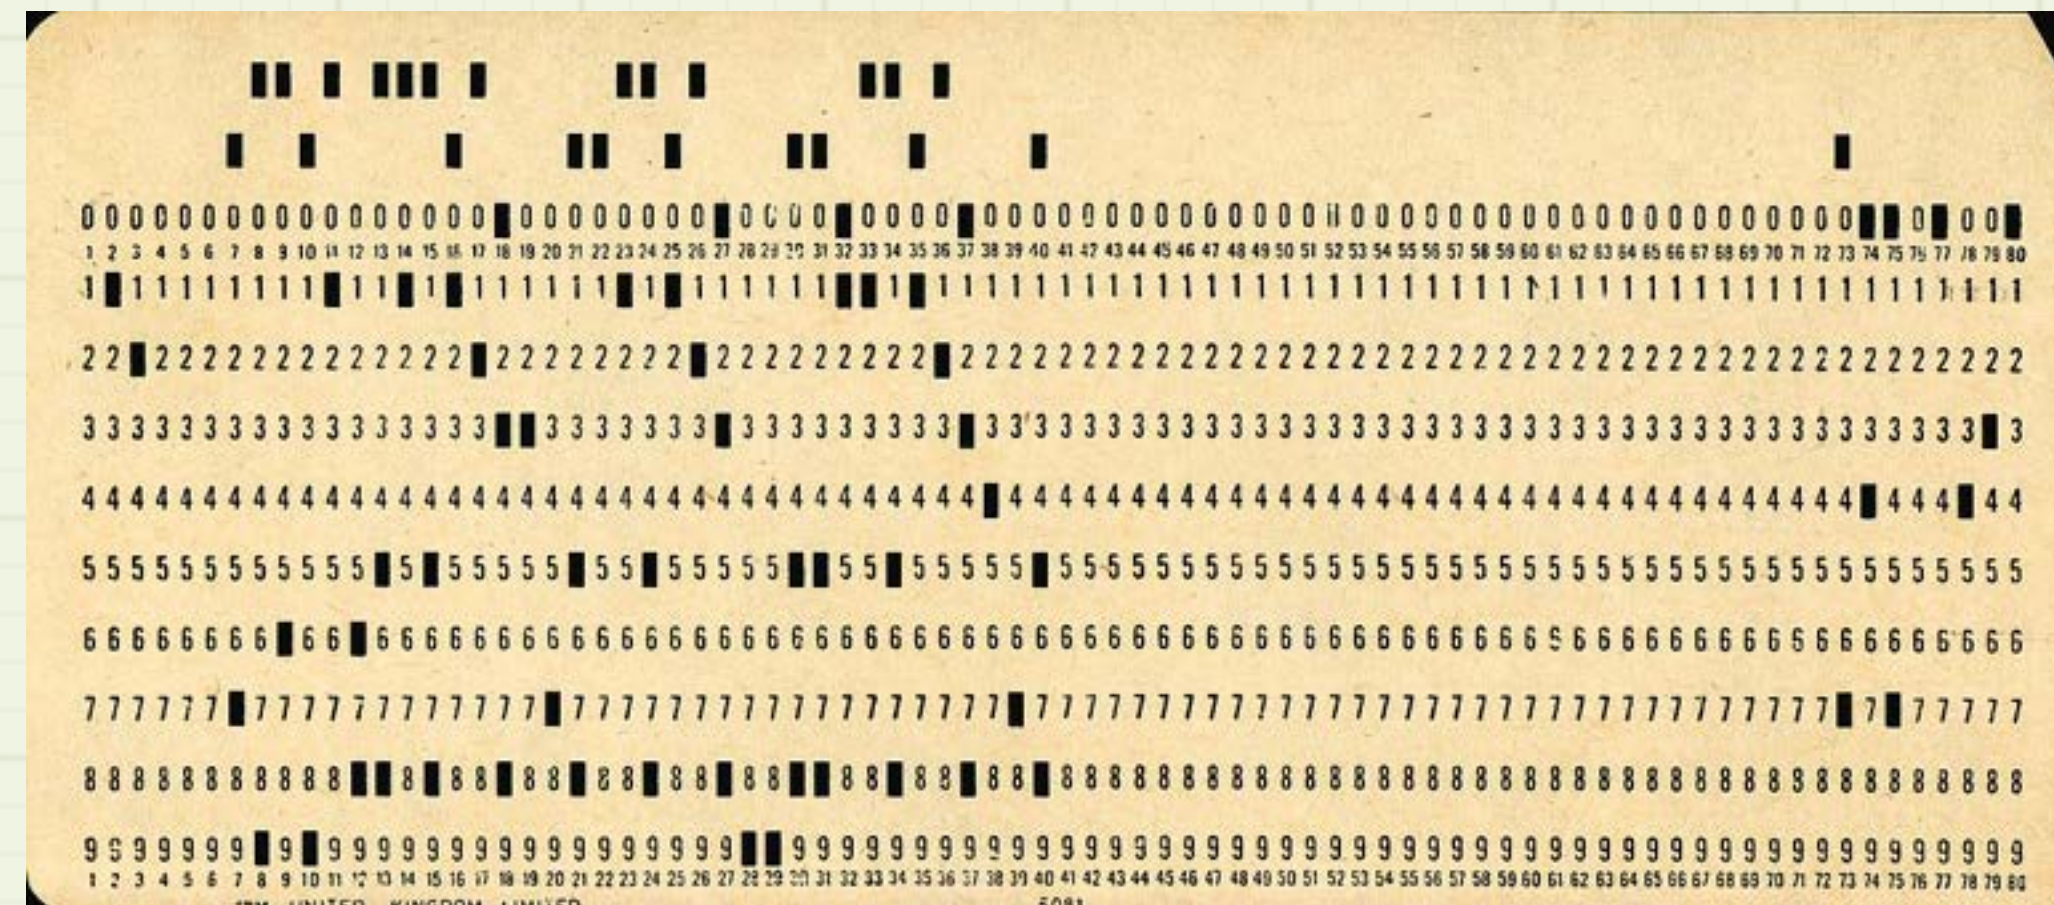
- Machine Learning techniques became dominant in the 1990s through 2010s
 - Linear and Logistic Regression
 - Support Vector Machines and Kernels
 - Classification and Regression Trees
 - Random Forests
 - Gradient Boosting
 - Neural Networks (pre deep learning)
 - Ensemble Methods
 - Sparse Machine Learning / Compressed Sensing
- Capstone texts include:
 - Hastie, Tibshirani, Friedman, *The Elements of Statistical Learning*, Springer 2001.
 - Bishop, *Pattern Recognition and Machine Learning*, 2006.

This is what is often meant by “machine learning”



Tabular Data Example

```
> library(palmerpenguins)
> data(package = 'palmerpenguins')
> head(penguins)
# A tibble: 6 × 8
  species island    bill_length_mm bill_depth_mm flipper_length_mm body_mass_g sex    year
  <fct>   <fct>          <dbl>          <dbl>          <int>          <int> <fct> <int>
1 Adelie Torgersen        39.1           18.7           181            3750 male   2007
2 Adelie Torgersen        39.5           17.4           186            3800 female 2007
3 Adelie Torgersen        40.3           18             195            3250 female 2007
4 Adelie Torgersen        NA             NA             NA             NA NA     2007
5 Adelie Torgersen        36.7           19.3           193            3450 female 2007
6 Adelie Torgersen        39.3           20.6           190            3650 male   2007
```

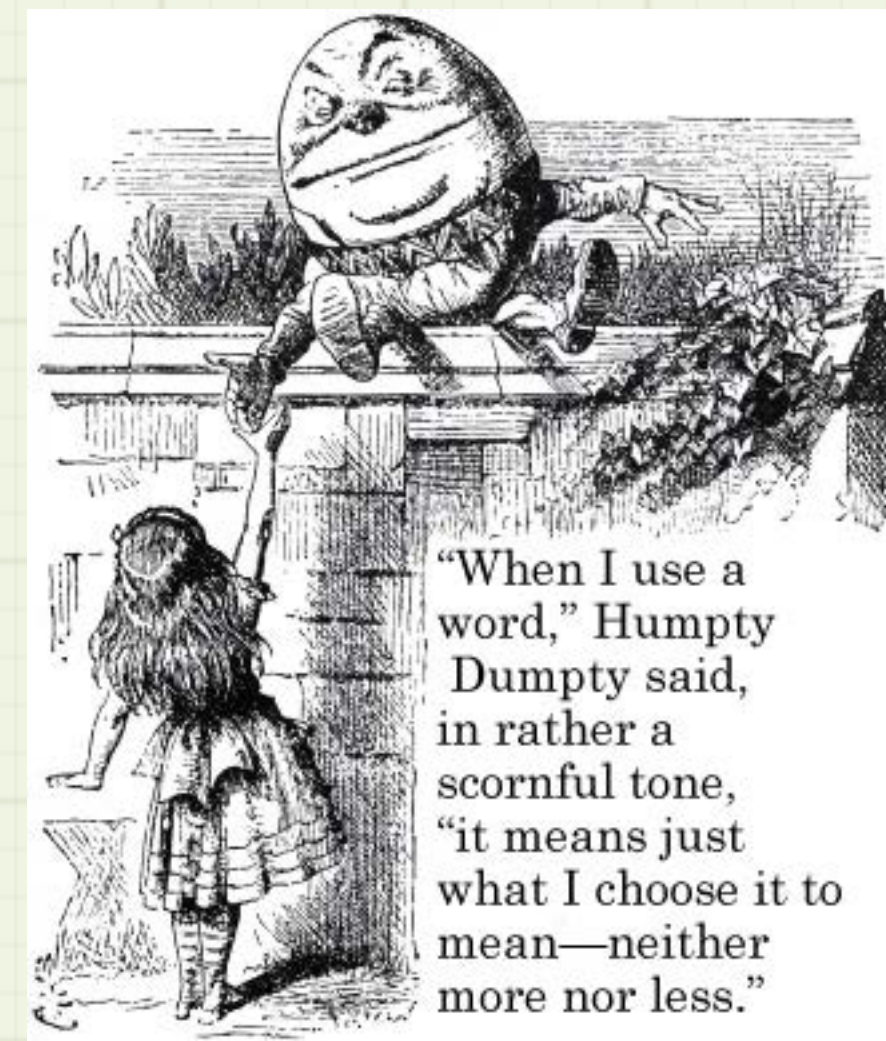


Supervised Machine Learning Task

- Assume one column is both valuable to know, and expensive to know in a timely fashion. Call this column the “outcome” or “dependent variable”: y .
- Label all columns that are cheap to know as the “explanatory variables”: x (x a row vector per instance).
- Use historic data where all columns are known to infer a function $f()$ such that $f(x) \sim y$ on many rows of training data.
- Assume or ensure there are reasons so that $f(x) \sim y$ should continue on new or application data.
- Now when only x is available, we have an estimate of y !

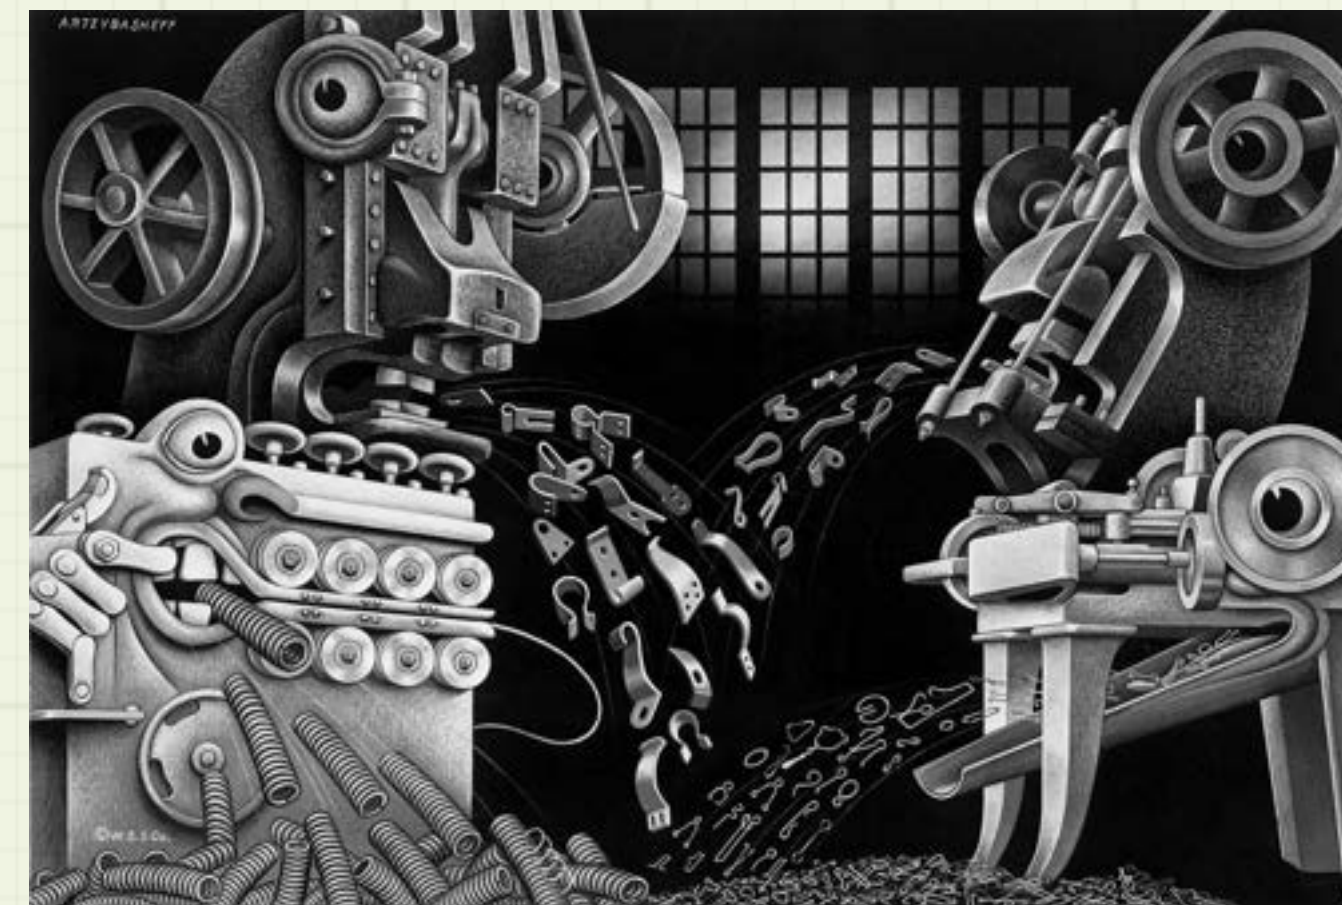
Some Terminology Traps

- Inference
 - Statisticians use term to mean estimating $f()$. Also called “training.”
 - ML Engineers use term to mean computing $f(x)$, given $f()$ and x , others call this prediction.
- Prediction
 - At best statistical machine learning models predict $f(x) \sim E[y | x]$, not actual precise values for given instances.
- Independent variables
 - Traditional (horrible) name for explanatory variables.
- Regression
 - Traditional name for predicting a number. Comes from a early criticism of the process: regression to mediocrity, or inferring the regression line.
- Classification
 - Traditional name for predicting a choice from a fixed list of alternatives.
 - *Always* inferior to predicting probabilities and leaving choice to a controlled policy or decision procedure.
 - Insistence on “classification returns a class label” leads to harmful procedures: such as naively re-balancing training data.



A Difference Between Machine Learning and Classic Modeling

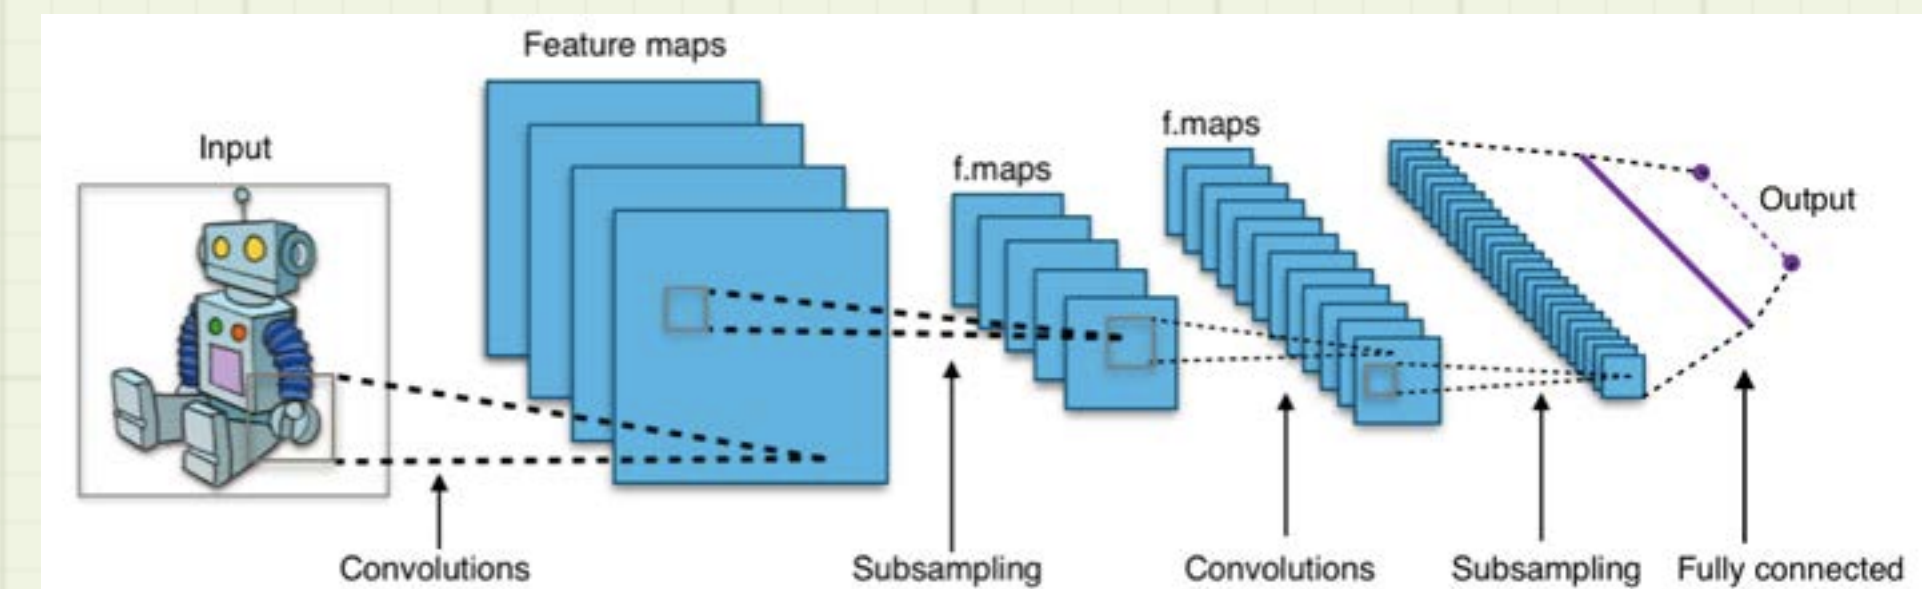
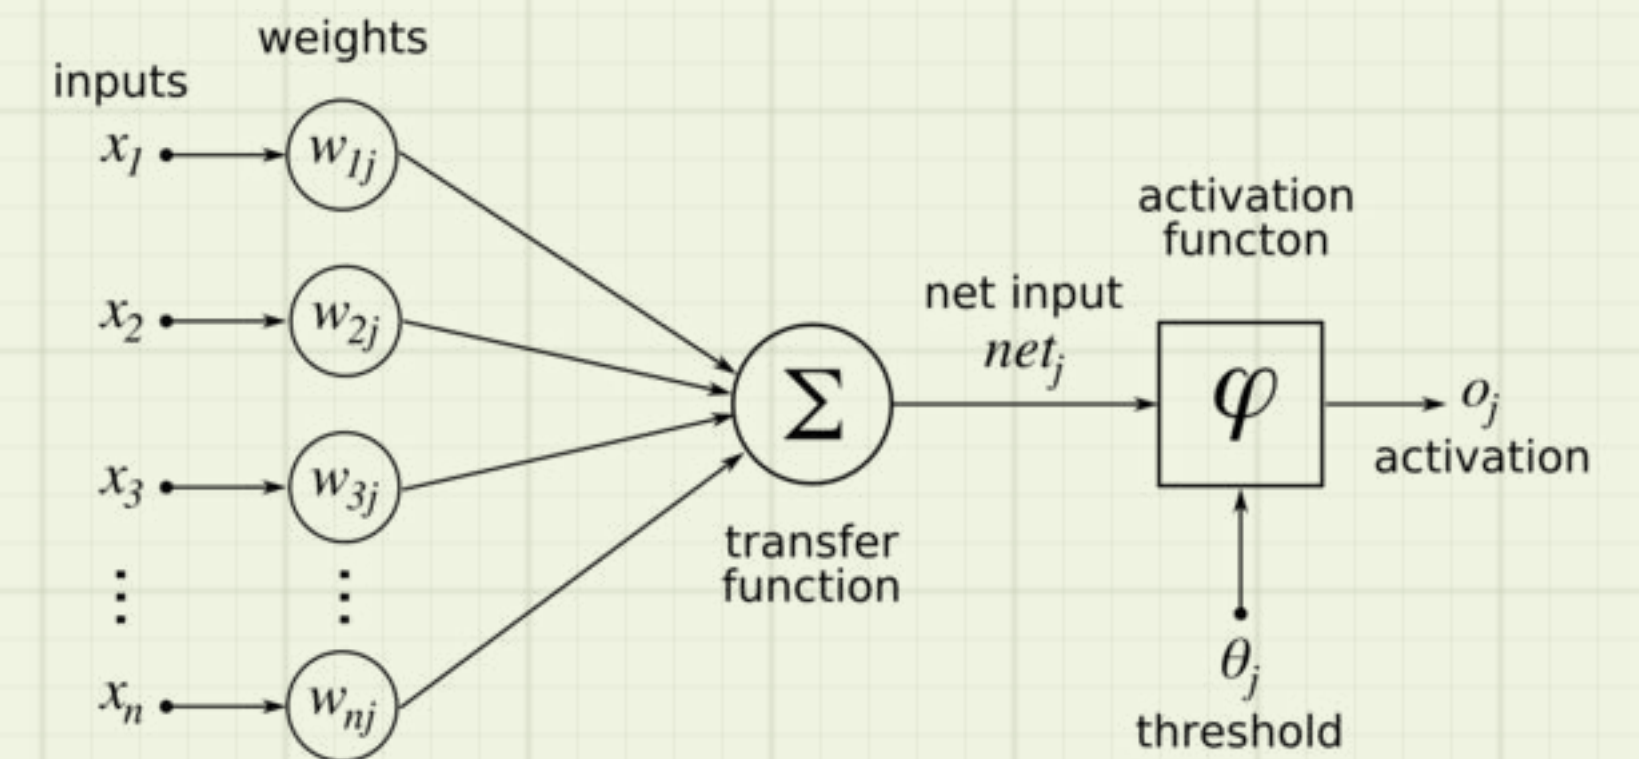
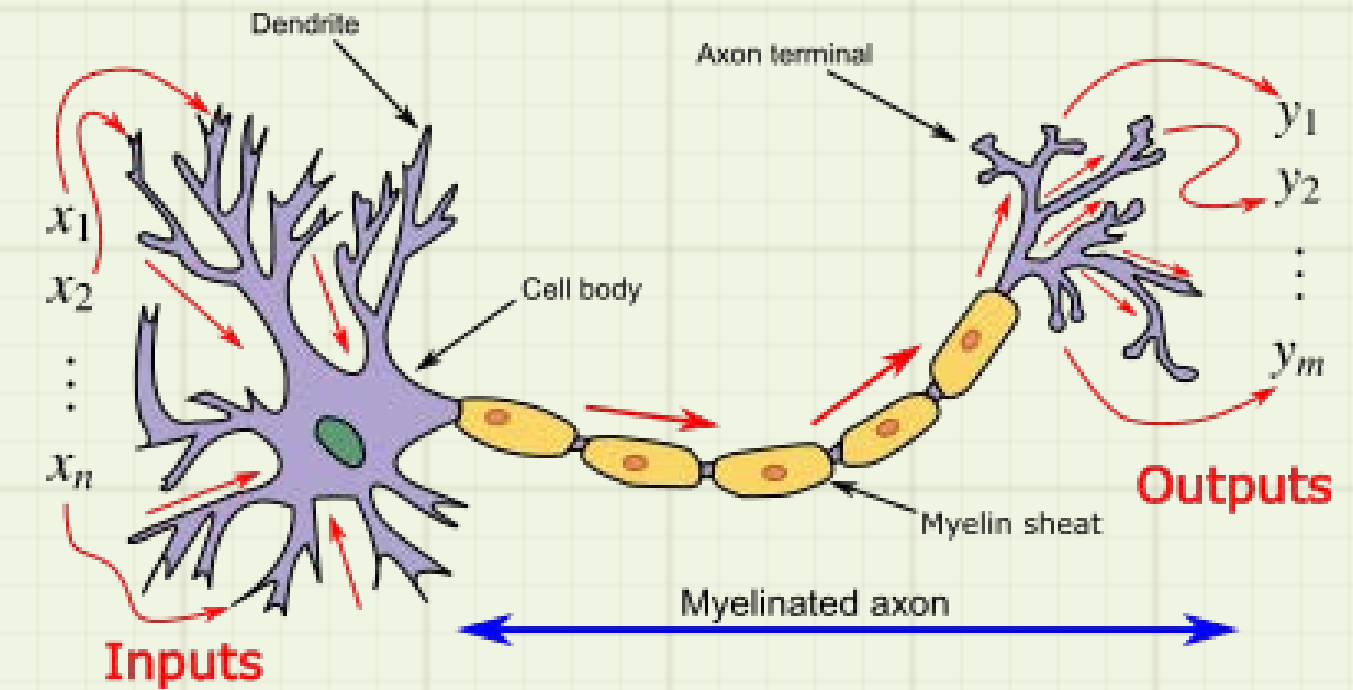
- In classic modeling one often models aggregates directly
 - Model the defect *rate* of a factory as a function of physically modeled *controlled* factors such as temperature and humidity.
 - Modeling considers functions over aggregates.
- In machine learning one usually models individuals and then aggregates later
 - Model the probability of an individual part having a defect as function of *observed* factors such as historic temperature and humidity.
 - Defect rate then estimated by aggregation of predictions over proposed individuals.



Deep Learning

What is Deep Learning?

- Large directed acyclic graphs of mathematical expressions
 - Suggestively called “neural networks.”
- Plus a specific set of training from data strategies
 - Stochastic gradient descent
 - Drop out (a protection against over fitting)
 - Simulating un-censorship to generate training data
- Huge growth since the 2010s.

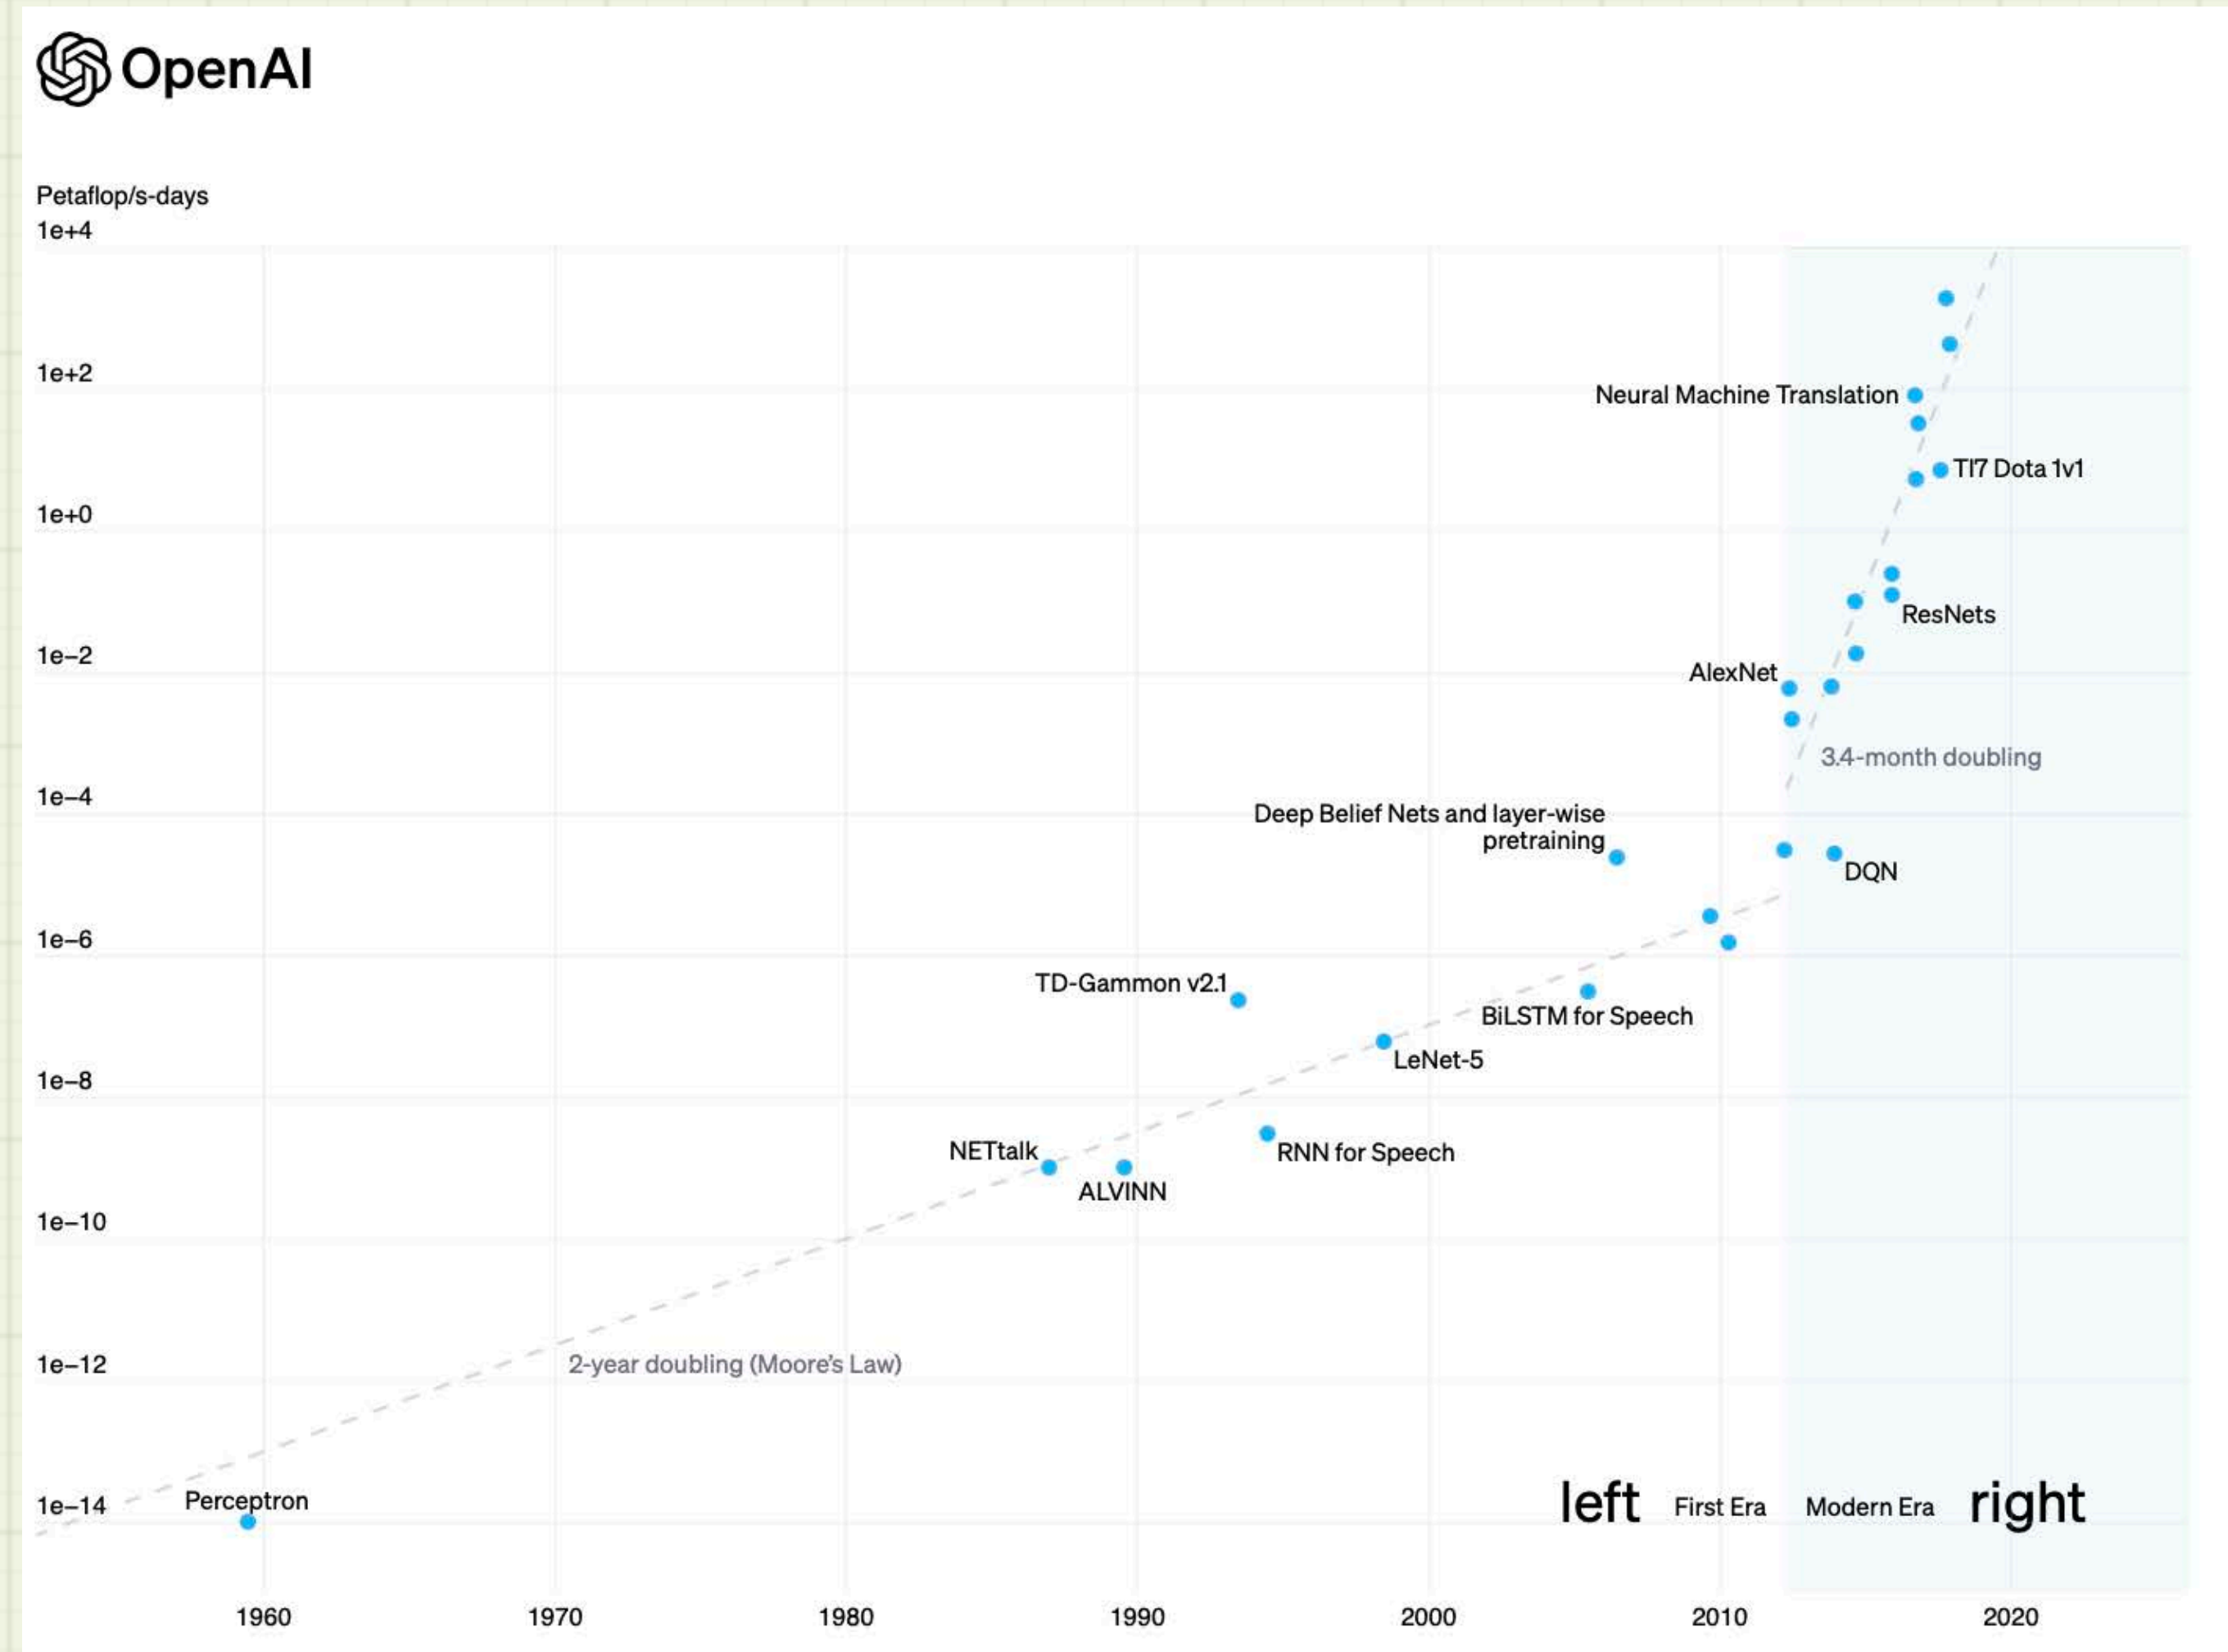


Un-Censorship

- Learn to replace the missing word
 - Original “As right as rain”
 - As training data
 - explanatory content (input): “As right as ?????”
 - training outcome: “???? = rain”
 - Same ideas as building a Doob Martingale by controlling exposure order.
- First *huge* and famous success of this method: `word2vec` (Mikolov, Chen, Corrado, Dean, 2013)
 - Observing internal activation pattern of neural net leads to auto encoders and semantic embeddings
 - Moderate dimension: 200 to 300 columns
 - Related concepts near each other in embedding!



Deep Learning is measured in \$

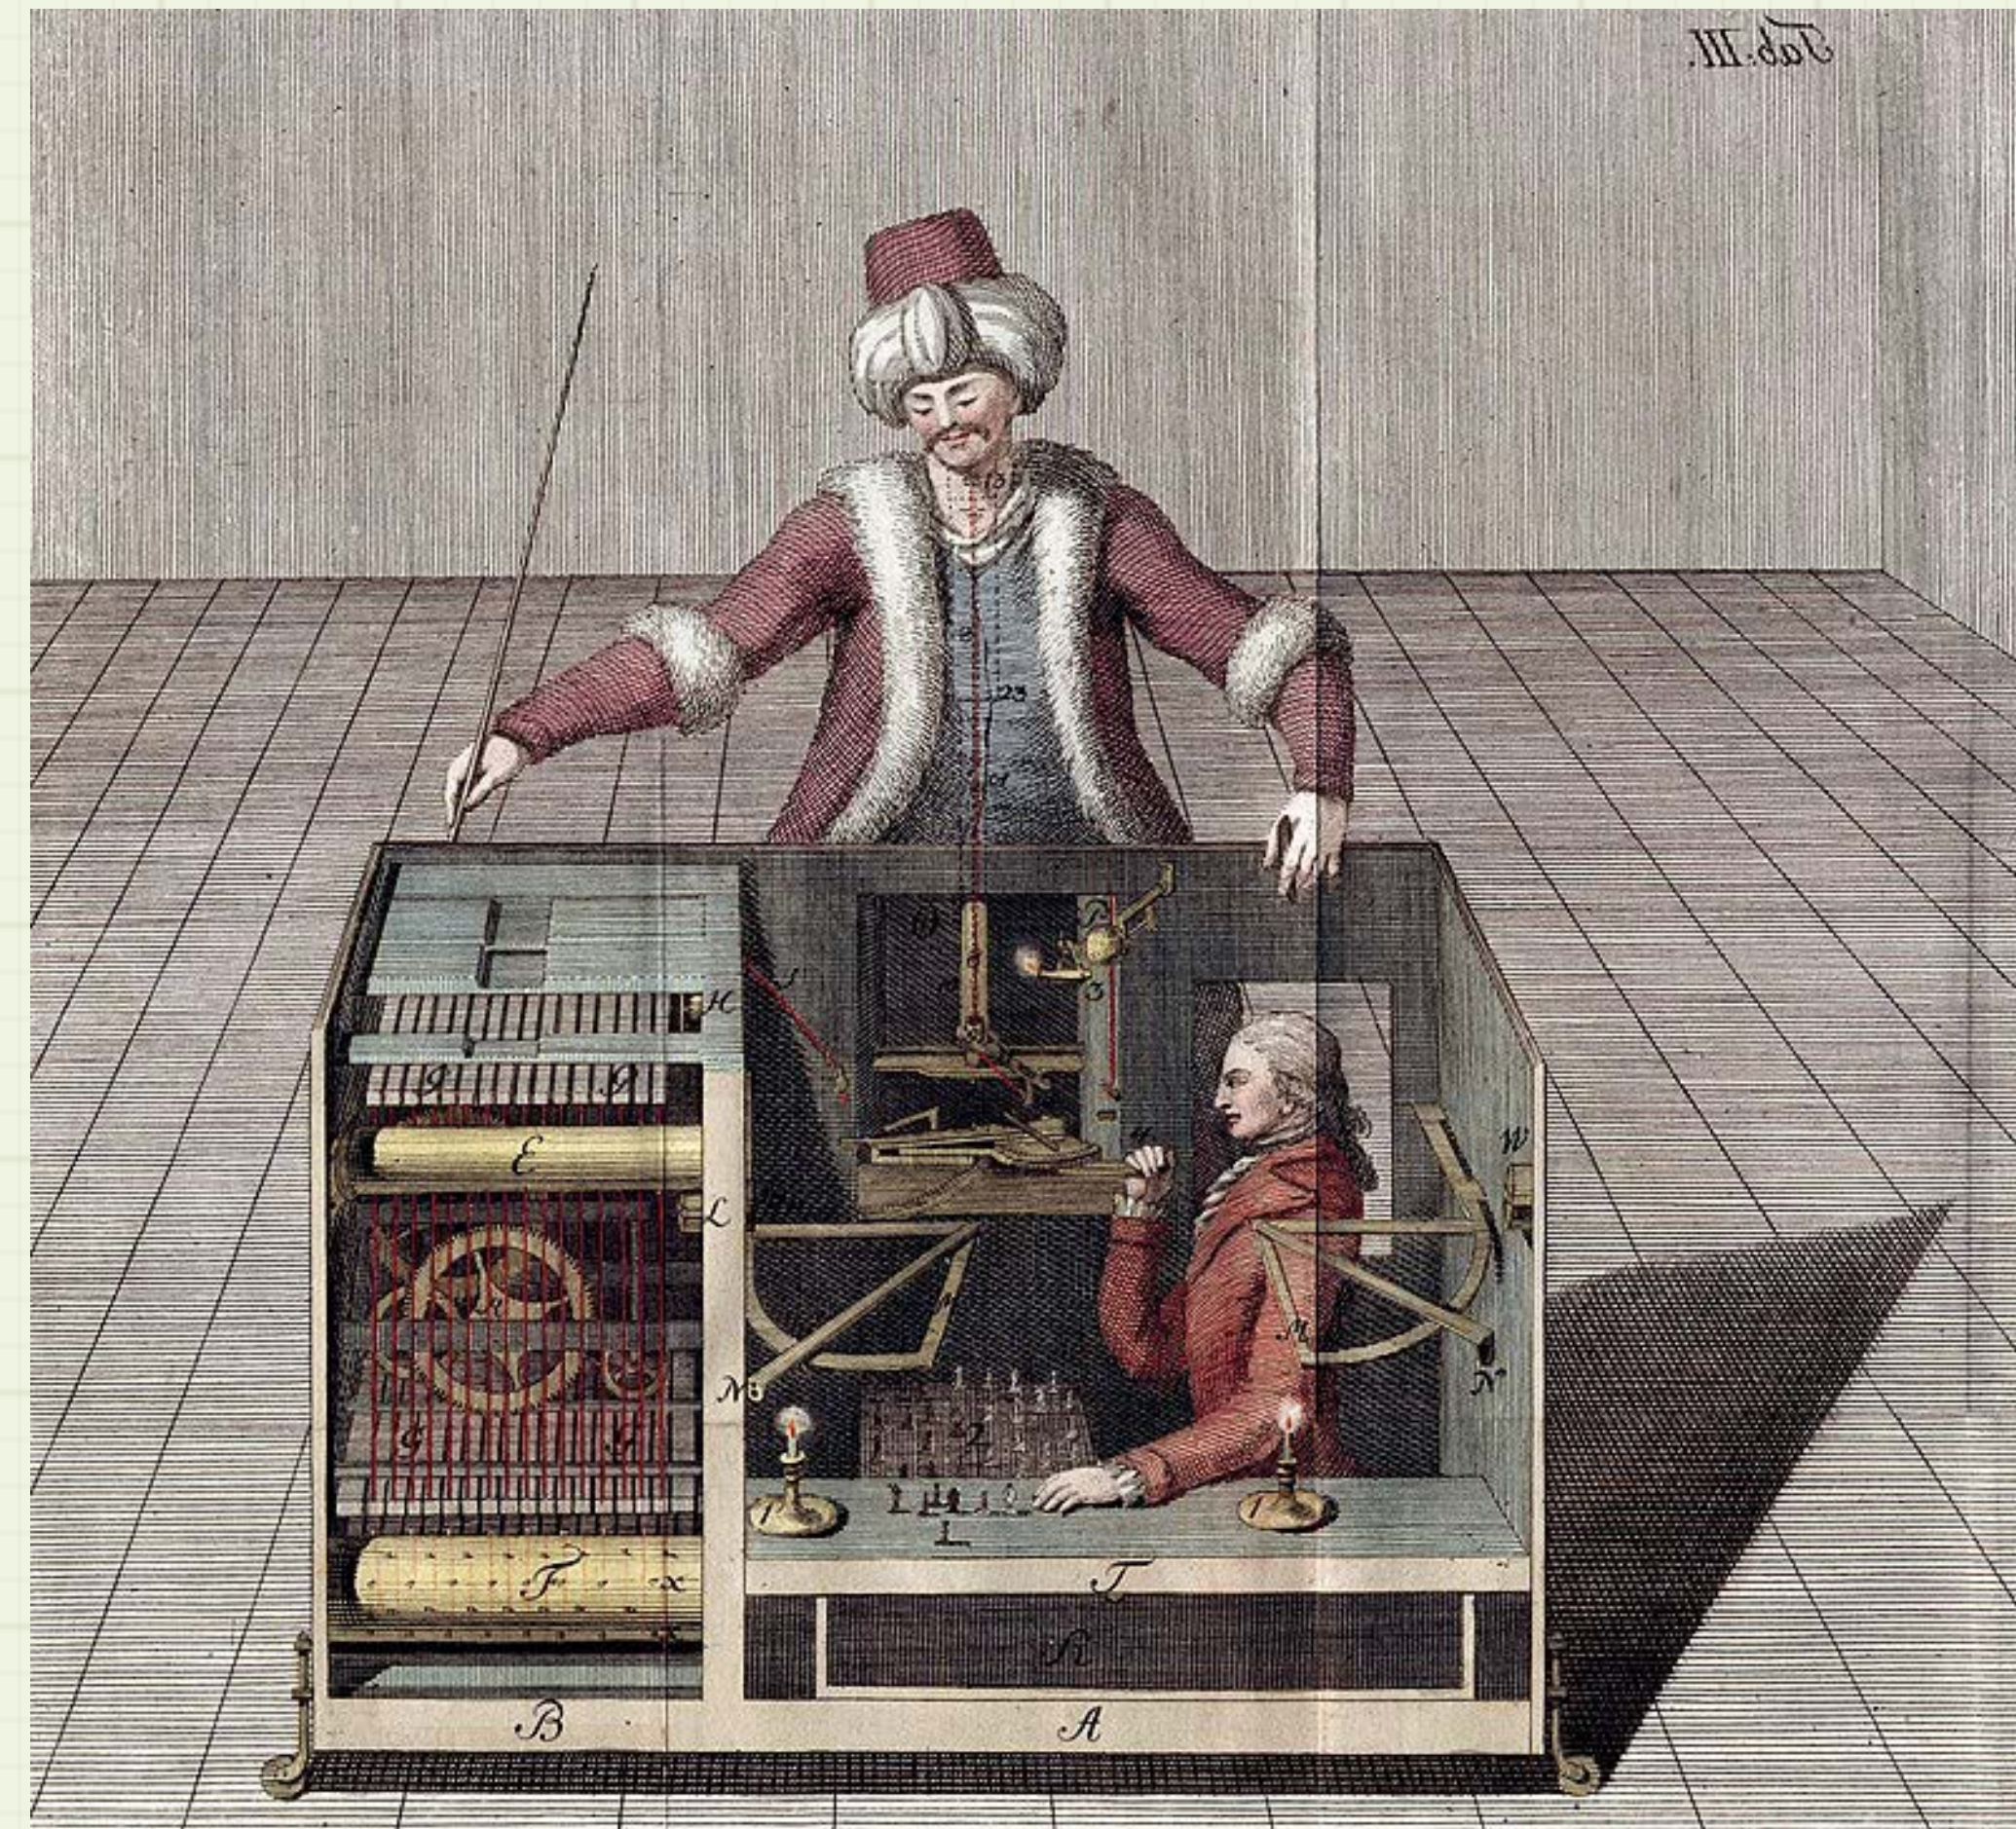


<https://openai.com/research/ai-and-compute>

AI

AI appears to be here

- Components
 - Large Language Models (LLMs)
 - Dominant since late 2010s early 2020s
 - Deep Neural Net trained on *huge* corpuses
 - Transformers
 - Self-supervised and semi-supervised training techniques
 - Prompt engineering (replacing user queries)
 - Post-processing
 - Specialized sub-systems
 - *Human* taggers
- Generally ***behind an API*** (application programming interface or service)
 - useful and impressive
 - yet, unfalsifiable



Actual use of a LLM based AI



What does the following code do?

```
import numpy as np
import pandas as pd

def get_digit_sum(anumber):
    sum=0
    while (anumber != 0):
        sum = sum + int(anumber % 10)
        anumber = int(anumber/10)

    return sum

def get_number(astring):
    cleanstr = astring.replace(' ', '').lower()
    tokens = [ord(c) - 96 for c in cleanstr]
    mynumber = sum(tokens)

    # reduce to a single digit
    while mynumber >= 10:
        mynumber = get_digit_sum(mynumber)

    return mynumber
```

ChatGPT3.5 Answer



The provided code defines two functions, `get_digit_sum` and `get_number`, which are used to calculate a numerical value associated with a given input string.

1. `get_digit_sum(anumber)`: This function calculates the sum of the digits of a number. It takes an integer `anumber` as input and iteratively adds up the digits of that number until the number becomes 0. It then returns the sum of the digits.
2. `get_number(astring)`: This function takes a string `astring` as input. It first cleans the string by removing spaces and converting it to lowercase. Then, it assigns a numeric value to each character in the cleaned string based on its position in the English alphabet (e.g., 'a' is 1, 'b' is 2, and so on). After calculating the sum of these values, it reduces the result to a single digit using the `get_digit_sum` function if necessary. The final single-digit result is returned.

Here's how you might use the `get_number` function:

```
python
```

Copy code

```
result = get_number("Hello, World")  
print(result) # This would print a single-digit result based on the input s
```

The purpose of this code is to assign a numeric value to a string by summing the values associated with its characters in the English alphabet and reducing the result to a single digit. This can be used for various applications, such as numerology or assigning values to words or phrases.

What are the AIs doing?

- Gave up on modeling mind or process
 - Modeling the mere structure of text or images
 - RAG or Retrieval Augmented Generation (replace portions of LLM answer with looked up results!)
- Produce results that are ***statistically indistinguishable from correct answers.***
- Takes Searle's "Chinese Room" *criticism as an actual implementation blueprint*
 - As a criticism
 - Does a book with every possible Mandarin phrase next to its English equivalent "understand Mandarin and English?"
 - As a blueprint
 - Levers the amazing amount of existing digital content
- Training and testing issues
 - GIGO (garbage in, garbage out) remains an issue!
 - Racist and sexist biases often retained from training data
 - May not be an easy ask an AI a ***novel*** question at this point
 - We are now ***"testing" on training data (not a reliable procedure!)***



Some issues working with AIs

- Models commonly produce *confident* plausible hallucinations instead of correct answers.
- **Currently hidden behind a “pennies per task” network API**, so we don't know much about the actual (versus claimed) implementation or unit cost.
- Better at producing spam than filtering spam.
- Impact on human taggers and editors in the process.
- **Current leading companies using funding and first-mover advantage to complete regulatory capture of the benefits** (see Yann LeCun's criticism).
- Partnering with AI companies may be risky
 - AI's are probably *nowhere near* as dangerous as the corporations that own them.



Additional Issue

- May not be safe to allow the general public to directly access your AI. (i.e. **don't do that!!!!**)
- **No** history of having a publicly facing AI stay on designed policy.



<https://spectrum.ieee.org/in-2016-microsofts-racist-chatbot-revealed-the-dangers-of-online-conversation>

<https://www.jumpstartmag.com/top-5-epic-fails-of-ai-chatbots/>

<https://www.bloomberg.com/news/newsletters/2022-12-08/chatgpt-open-ai-s-chatbot-is-spitting-out-biased-sexist-results>

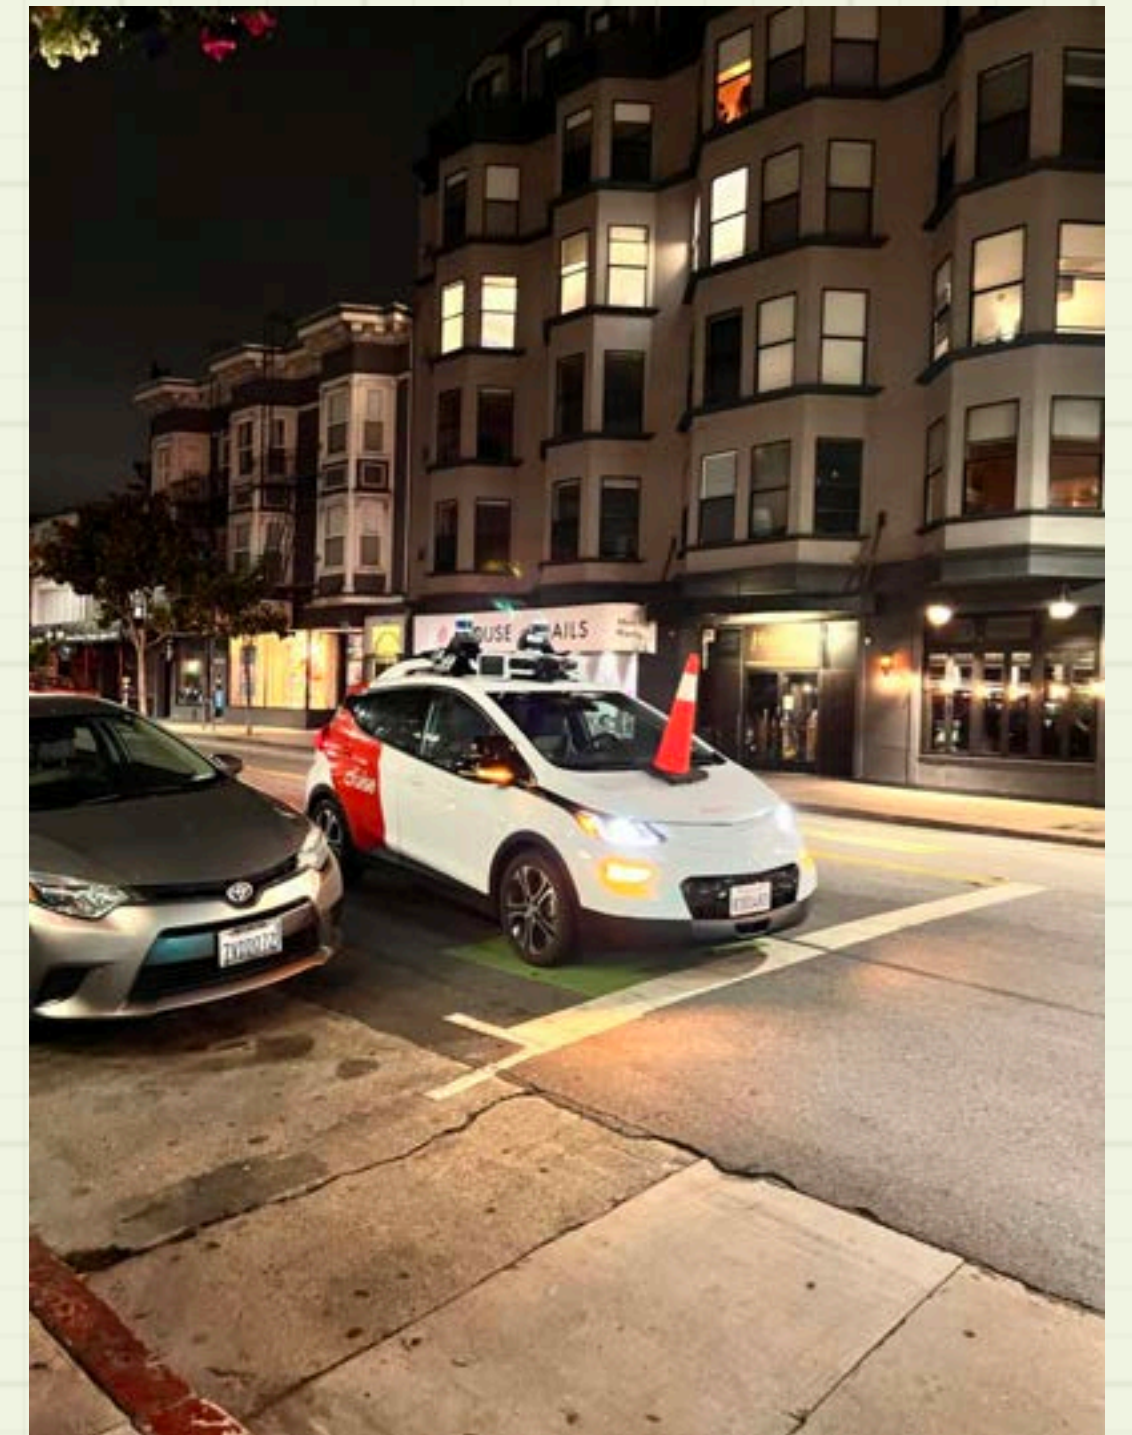


Behind the API curtain

<https://www.jwz.org/blog/2023/11/once-again-ai-is-revealed-to-be-an-army-of-mechanical-turks-in-a-call-center/>

Once again, "AI" is revealed to be an army of mechanical turks in a call center

“[Cruise “autonomous”] vehicles were supported by a vast operations staff, with 1.5 workers per vehicle. The workers intervened to assist the company's vehicles every 2.5 to five miles”



Conclusion



Summary Opinion

- Machine learning and AI are not the same thing
 - The days of selling a linear regression or some if-else statements as “AI” may be over
 - Have to go back to selling analytics as analytics (not a bad thing)
- Both Approaches ***treat observed quality of outcomes as inherent quality of decisions or process*** (a somewhat bad thing)
- AI ***is*** here
 - State of the art ***is currently a captive service***
 - AI augmented systems ***will*** outperform non-augmented systems
 - No longer just deep learning (prompt engineering, and sub-modules are critical)
 - Dominates in unstructured processes such as text and images
 - Right now probably best as a staff force multiplying tool
 - Not something to be ignored, a significant increase in capabilities in information retrieval



Thank You