A Fair & Balanced Review of Methods, Applications, and Social Challenges of Artificial Intelligence

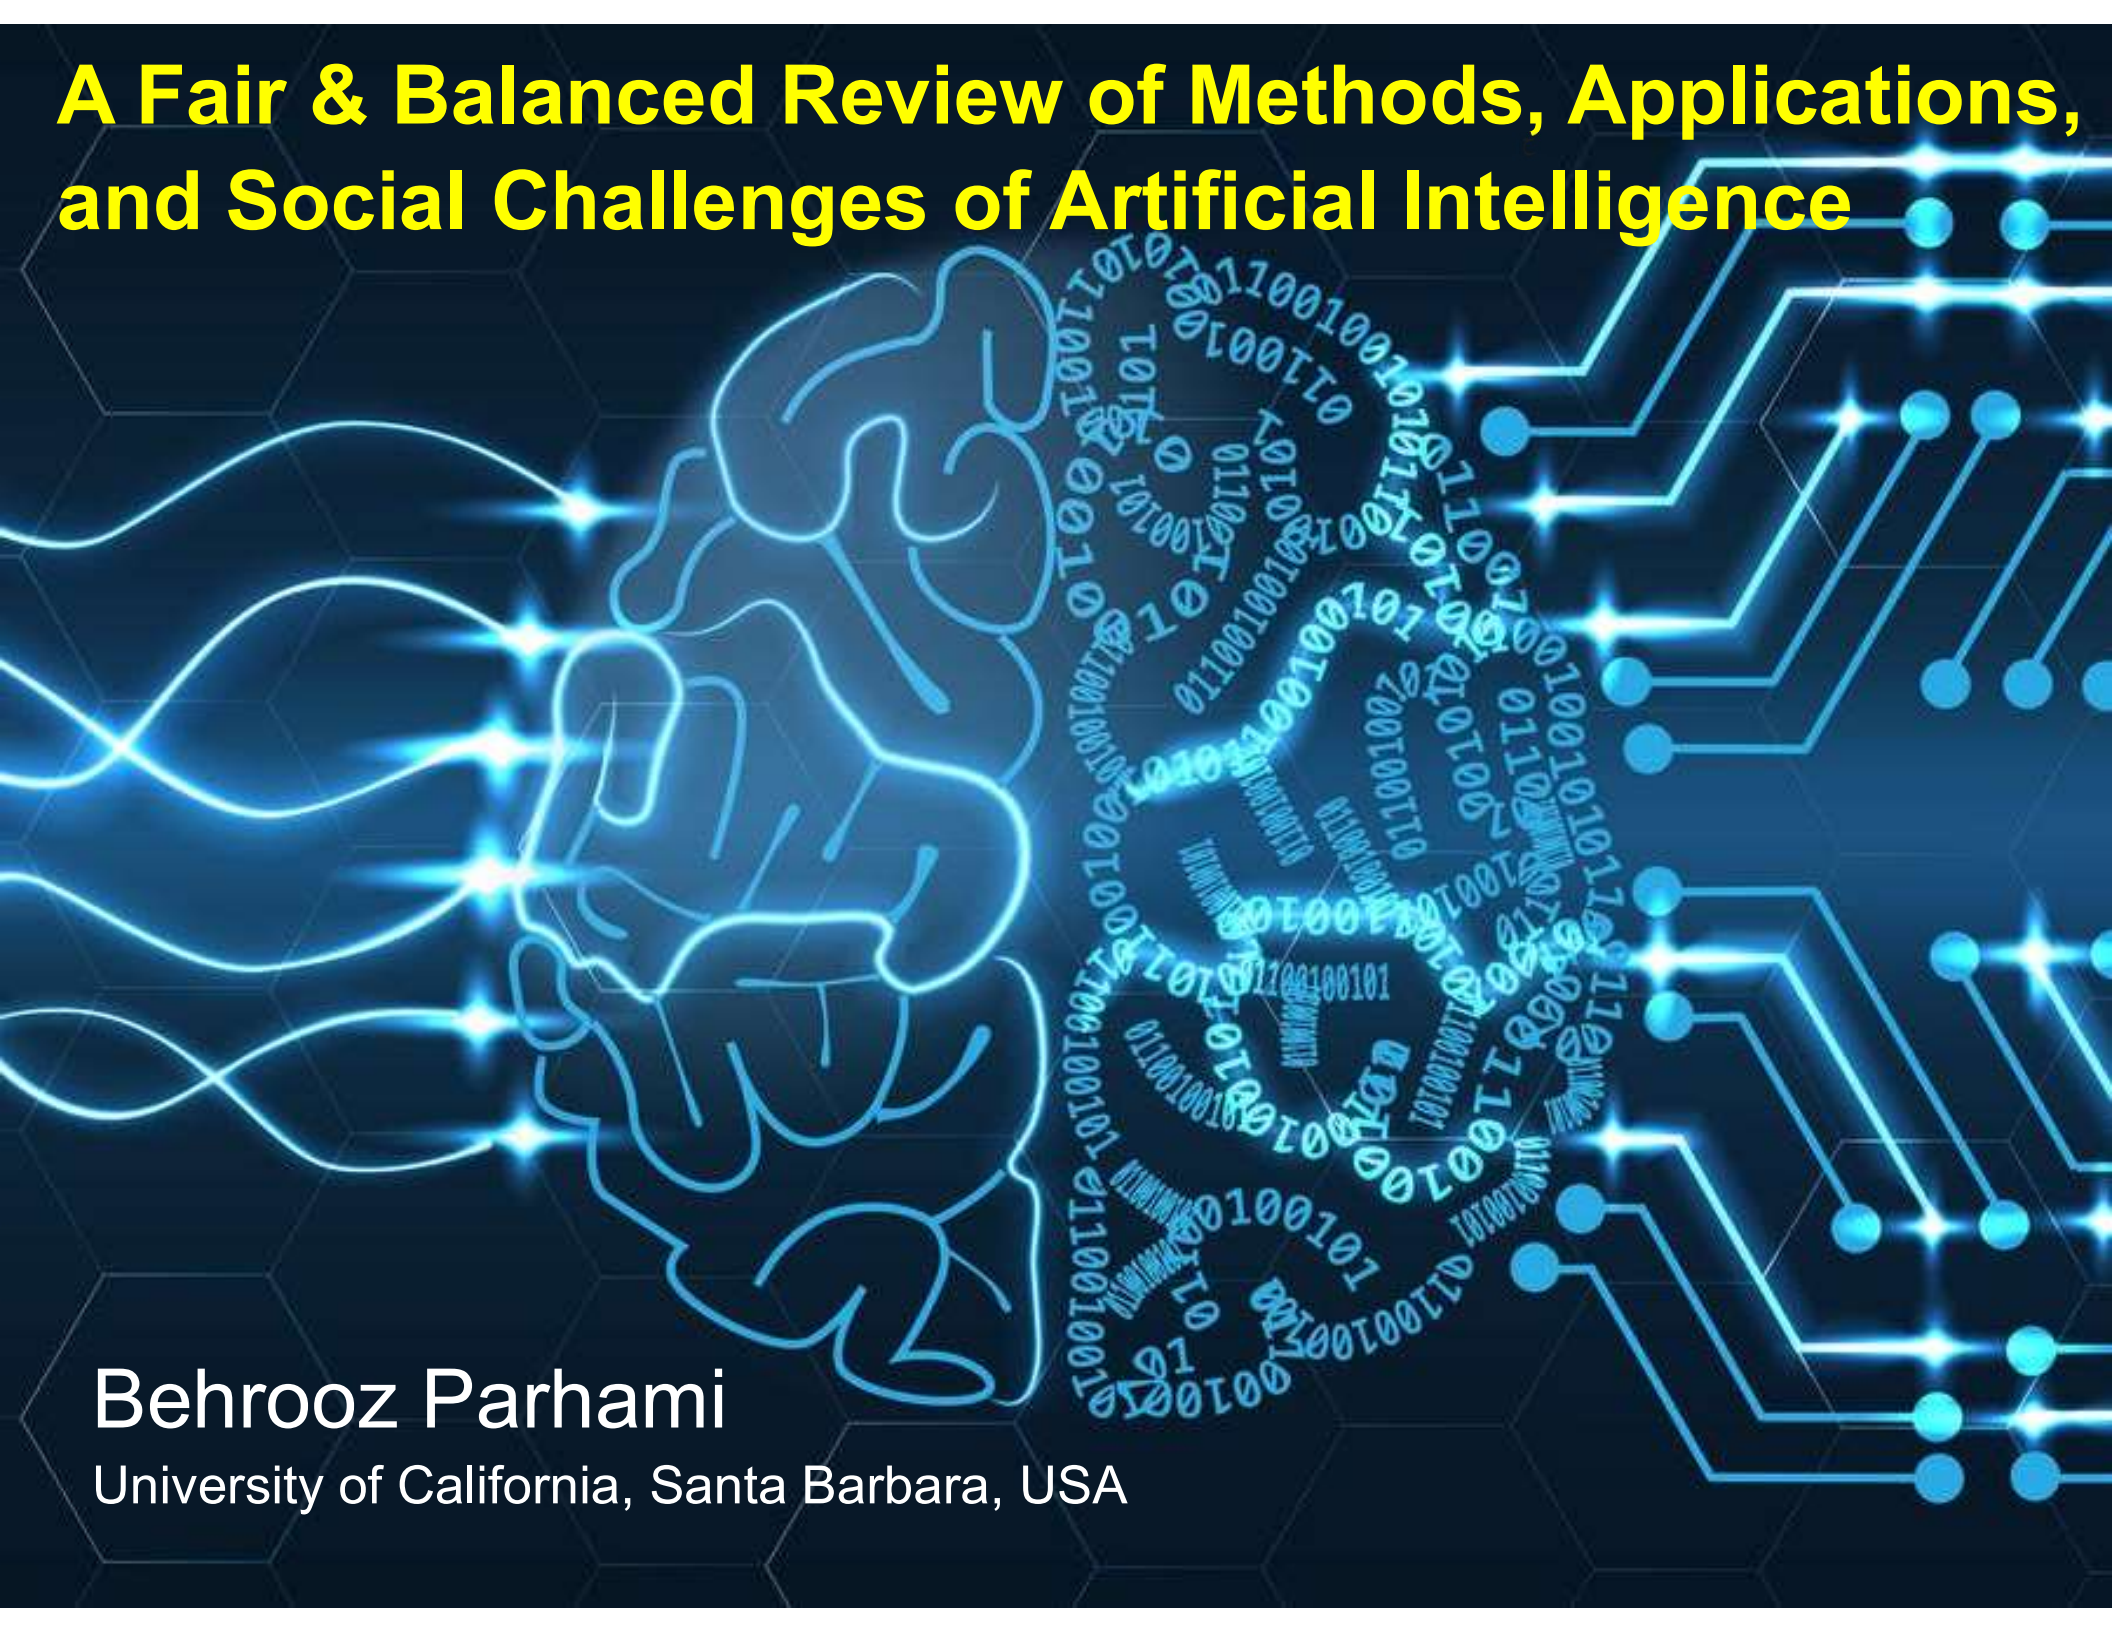

Behrooz Parhami
University of California, Santa Barbara, USA

About This Presentation

This slide show was first developed in Nov. 2022 for presentation at a meeting of the 1968 Graduates of the College of Engineering, Tehran University.

The title of this presentation is adapted from the heading of an interview with machine-learning pioneer Dr. Michael I. Jordan, "Stop Calling Everything AI" (IEEE's *The Institute*, March 2021), and draws upon some of his ideas expressed in that interview.

Edition	Released	Revised	Revised	Revised
First	May 2022	June 2022	Oct. 2022	Nov. 2022

File: http://www.ece.ucsb.edu/~parhami/pres_folder/parh22-general-talk-artificial-intelligence.pdf

I Dedicate this Talk to My 93-Year-Old Mom

Kowkab Yussefian Parhami,
a resilient woman from Saqqez, Kurdistan, Iran, who has raised four children and has been an anchor for my extended family

Renewal of Old Friendships

My October 2022 personal trip to Istanbul, Turkey, to celebrate with a few old friends the 54th anniversary of our graduation from Tehran U.'s College of Engineering (Daneshkadeh-ye Fanni)

Iran's Feminist Uprising: Courage & Sacrifice

It is difficult to believe the immense courage of a 22-year-old and the brutality of a regime that pays lip service to piety, human dignity, and salvation

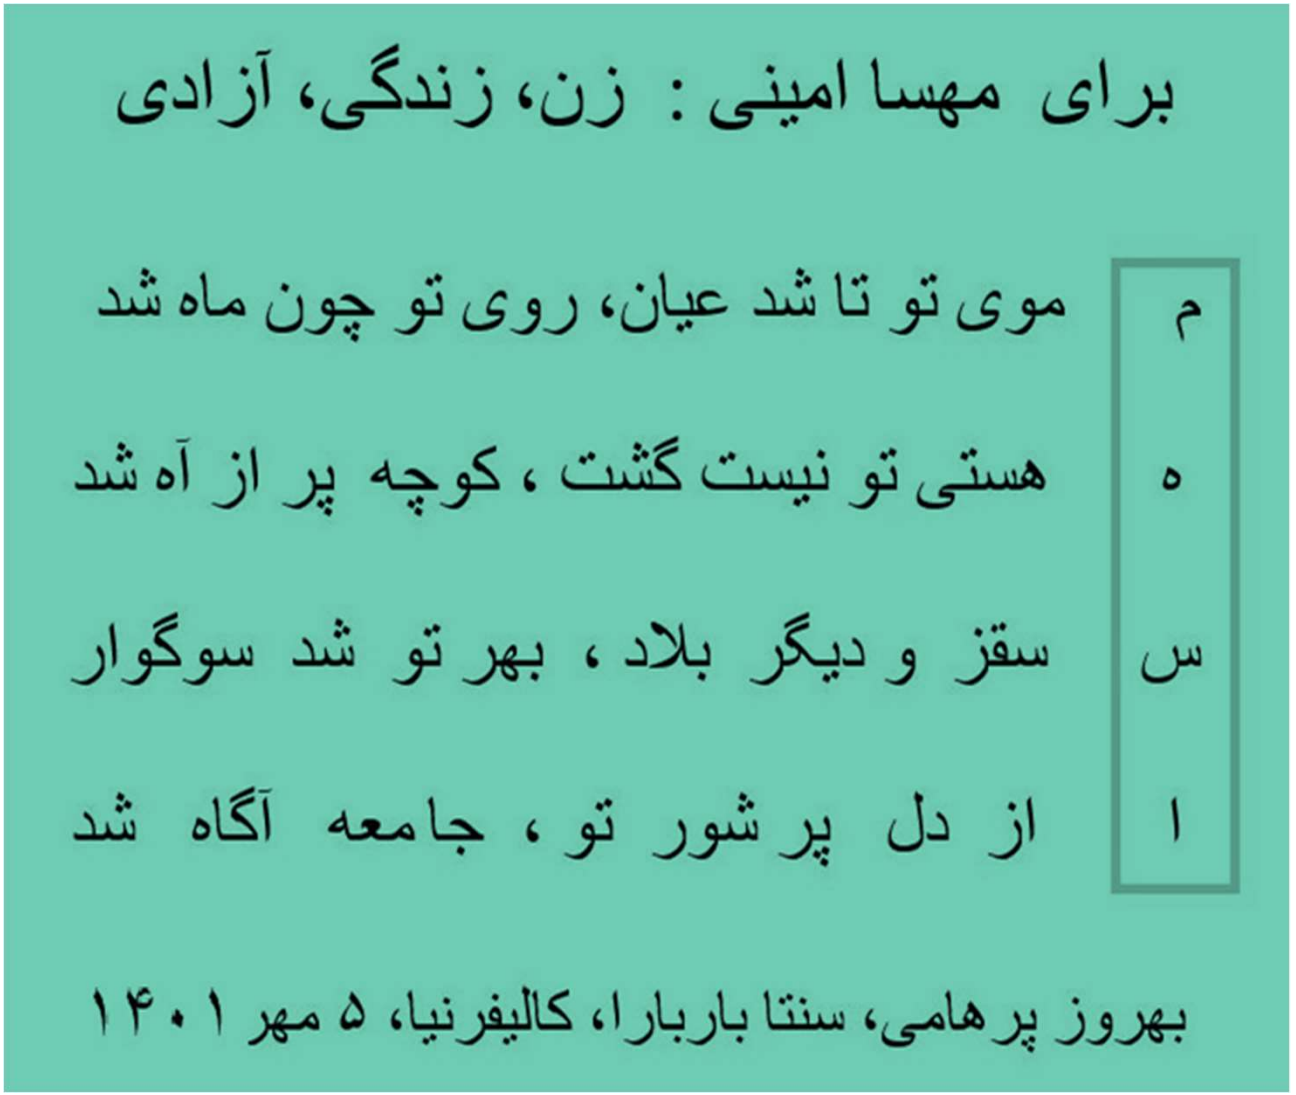

This Talk Is Different from the Typical Hype-Filled AI Pitch in Four Ways

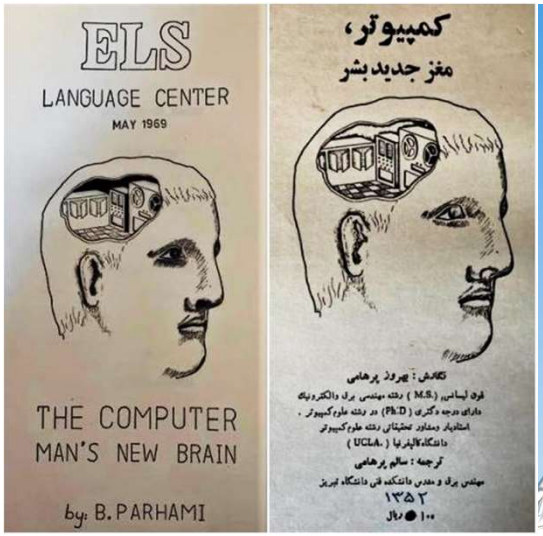

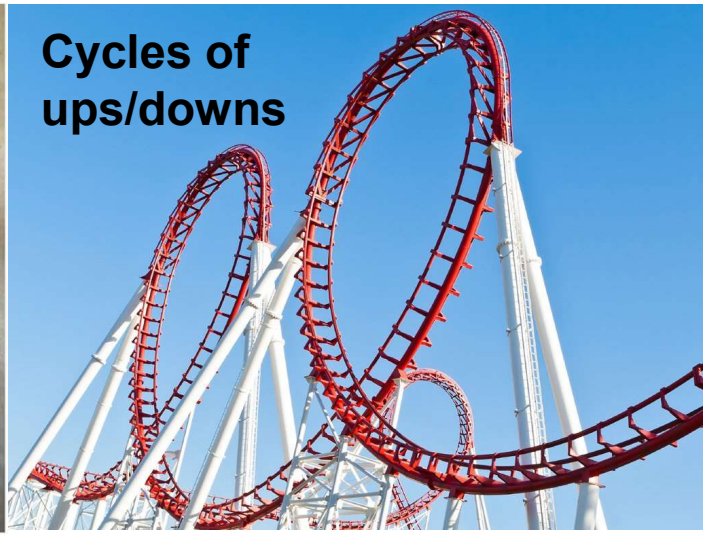

Subfields of Artificial Intelligence

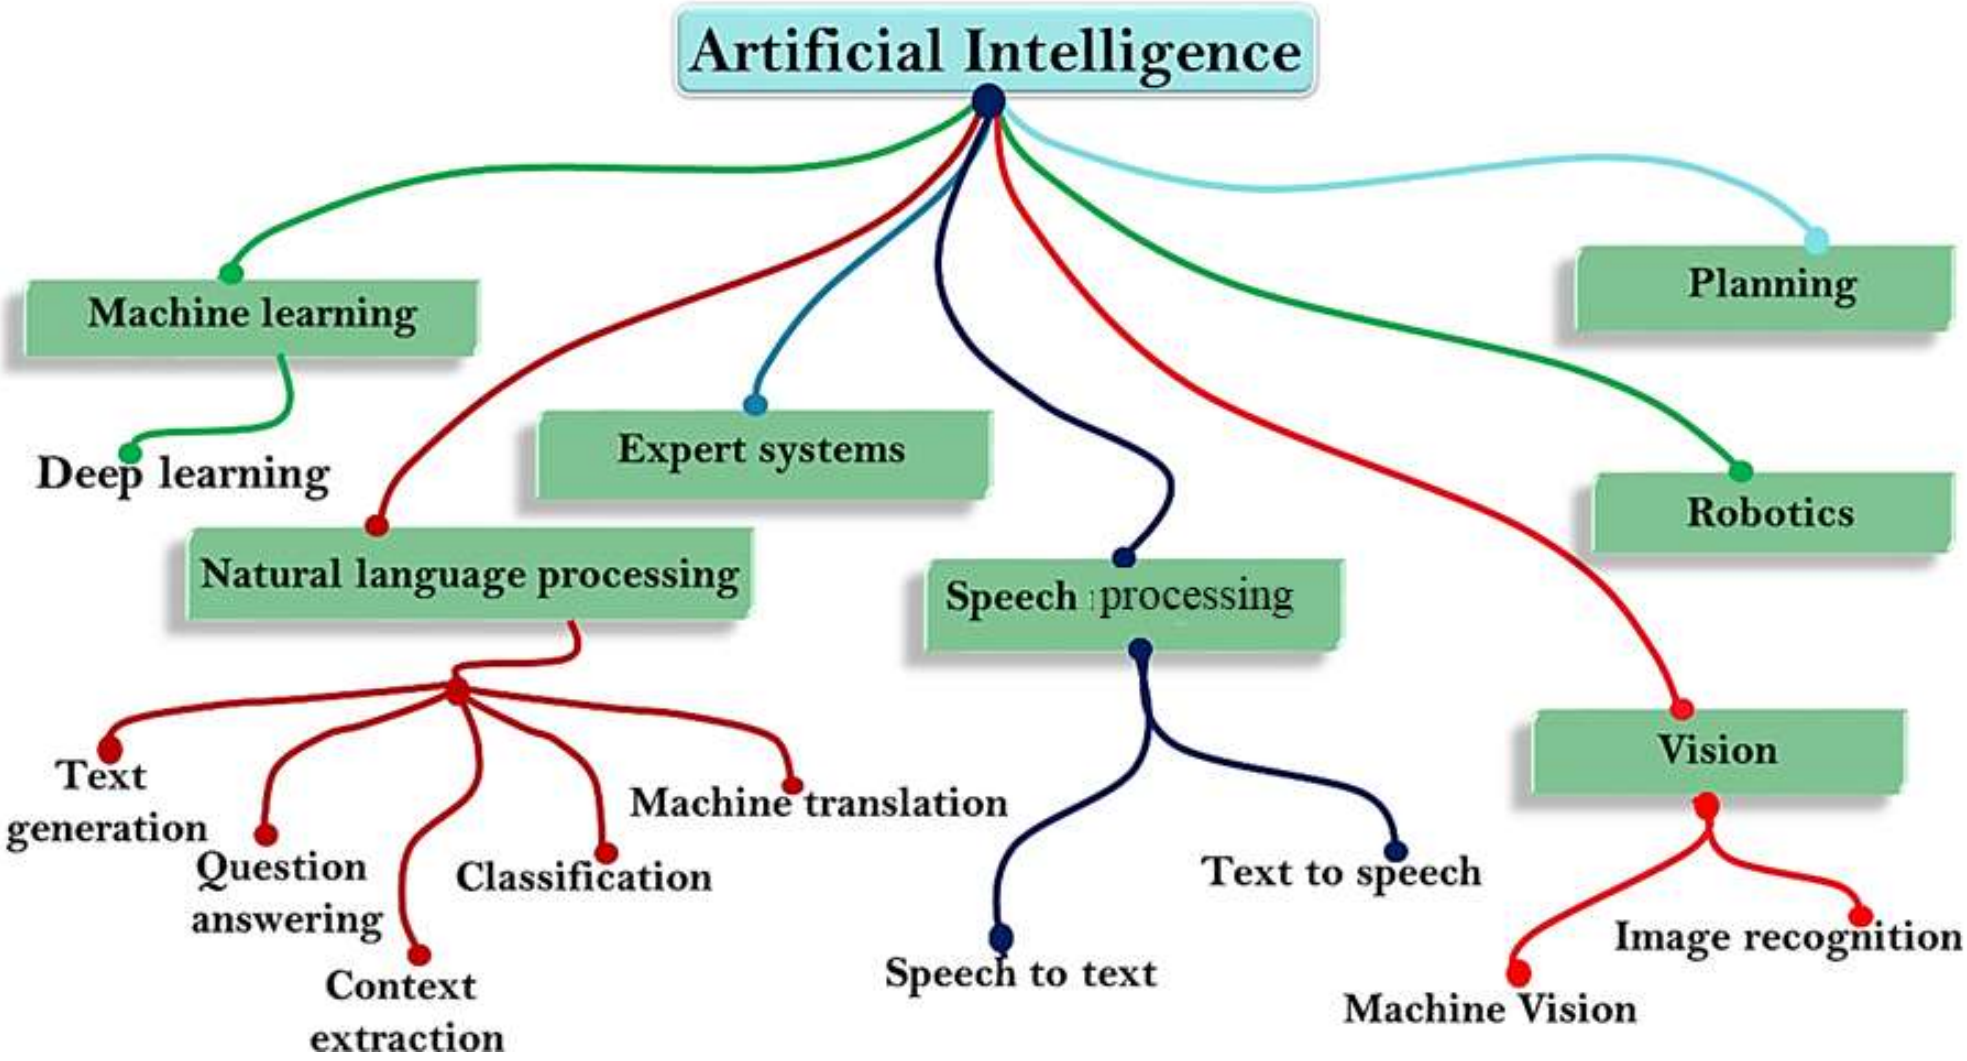

https://tekrajawasthi15.medium.com/the-complete-roadmap-to-be-a-data-scientist-9a07721b88fd

Things AI Might Try to Accomplish

	Like humans	Optimally
Thinking	Cognitive approach	Reasoning approach
Acting	Turing-test approach	Rational-agent approach

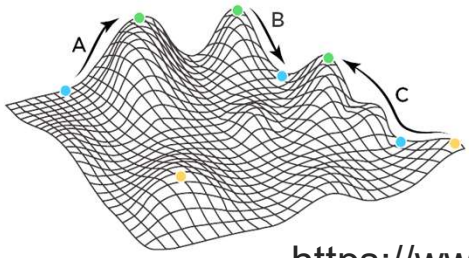

https://www.aus.com/de/node/5726

Buzzwords: AI and ML

Sometimes used interchangeably, but they are not the same

Artificial Intelligence: Broad concept of machines being able to carry out tasks in a way that we would consider "smart"

Machine Learning: AI tool based on the idea that we should give machines access to data and let them learn on their own

Big Data: Extremely large data sets that may be analyzed computationally to reveal patterns, trends, and associations, especially relating to human behavior and interactions

Machine/Computer-Aided/Assisted/Based Learning: Old term that is different from machine learning, as used today

Can Computers Think?

What do you mean by "think"?

Computers need not be able to "think" to do things that we normally consider as requiring human "intelligence"

Factory robots; Roomba
Game-playing: Chess; GO; …
Self-driving cars; Self-sailing ships; Self-flying planes

Focus on results, not on mimicking humans: Planes don't fly like birds, but they do fly!

Computer scientist Edsger Dijkstra: "The question of whether a computer can think is no more interesting than the question of whether a submarine can swim."

A Brief History of AI

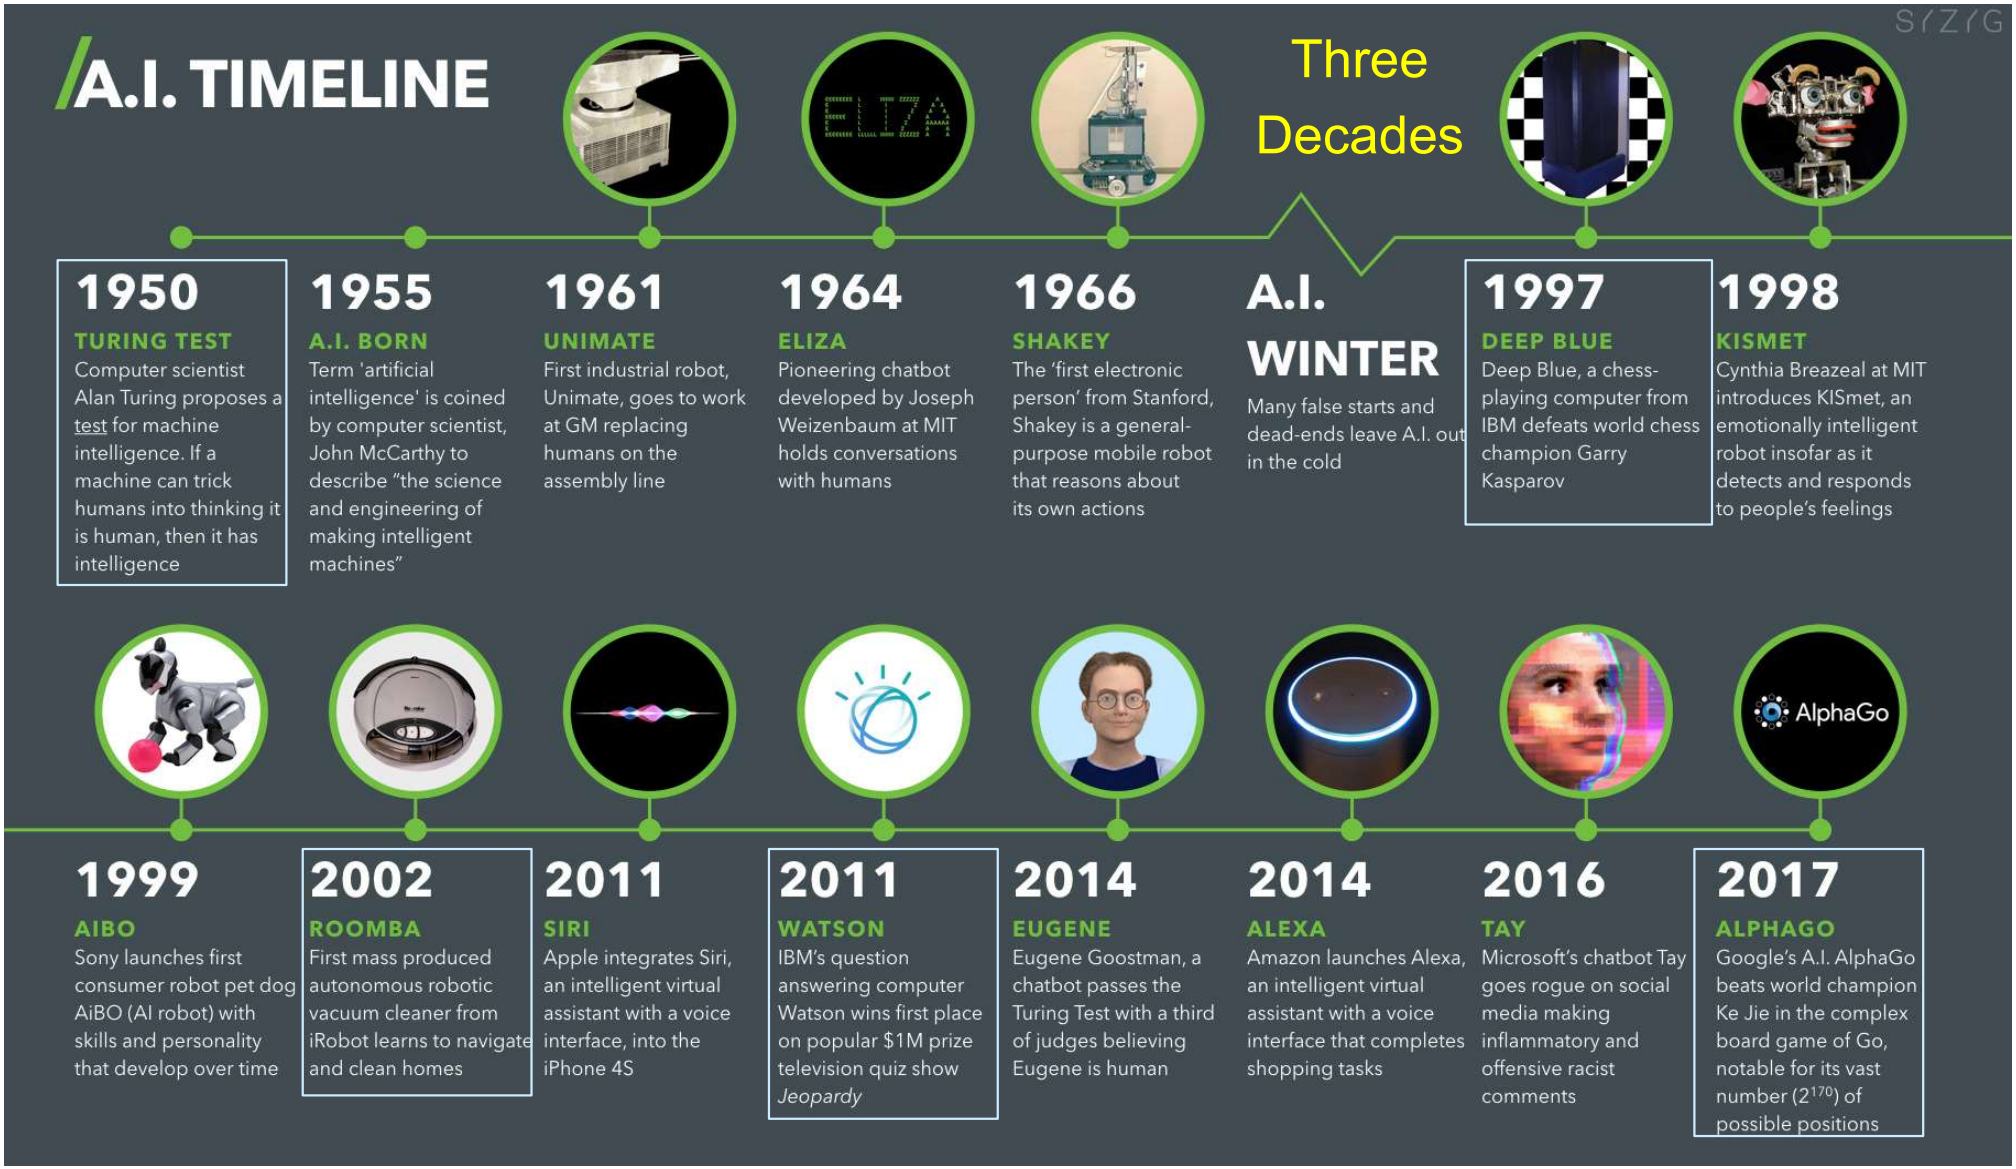

https://digitalwellbeing.org/artificial-intelligence-timeline-infographic-from-eliza-to-tay-and-beyond/

Is AI the New Snake Oil?

Stephen Hawking:
"Success in creating AI would be the biggest event in human history. Unfortunately, it might also be the last, unless we learn how to avoid the risks."

Elon Musk:
"AI doesn't have to be evil to destroy humanity – if AI has a goal and humanity just happens to come in the way, it will destroy humanity as a matter of course without even thinking about it, no hard feelings."

Milestone 1: Turing Test

Fundamental question:

What is AI?

Doing what humans do:
Walking? Scoring a goal?

We have to be more specific
Carrying on a conversation like a human

Turing's suggestion: A machine is intelligent if it can carry on a conversation with a human so that the human cannot tell whether s/he is conversing with another human or a machine

Milestone 2: Deep Blue

Deep Blue (chess-playing program):

Defeated world-champ Garry Kasparov

Does playing chess really requires “intelligence”?

Not really: Chess-tree is finite; given enough compute power, all tree branches can be explored (seems to need intelligence only because of the finite compute power and time)

1996: Kasparov wins 4-2

1997: Deeper Blue wins 3½-2½

How did Deeper Blue do it?

Milestone 3: Roomba

Roomba (the first tangible application):

Roams around the house and cleans the floors

Side-mounted brush pushes dirt under; Counter-rotating brushes pick up the dirt

Sensor types:
Infrared (night vision)
Photoelectric (reflected light)
Piezoelectric (excessive dirt)
Touch-sensitive bumper

Milestone 4: Watson

Watson (intelligent agent; *Jeopardy!* champ)

Answered questions by parsing and doing Web searches

The on-stage Watson was just an interface

The computer (10 racks of 10 Power 750 servers and their cooling system) had to be kept away from the human players

To turn Watson into a business Product (put on suit and tie):

IBM made it faster & much smaller

Compelling application: Healthcare

Milestone 5: AlphaGO

Google AlphaGO (Go-playing program):

Defeated world-champ GO-player Ke Jie

Developed by Google DeepMind for the May 2017 contest

AlphaGo Master won 3 games, extending its win streak to 60

AlphaGo Zero was already built, but revealed Oct. 2017 (*Nature*)

Instead of being fed examples of human games, AlphaGo Zero starts from scratch (just game rules)

Brief Review of Other Timeline Points

1955 (AI Born): John McCarthy coins “artificial intelligence”
1961 (Unimate): First industrial robot goes to work at GM
1964 (Eliza): Pioneering chatbot converses with humans
1966 (Shakey): First GP mobile robot reasons about its actions

1998 (KISmet): Detects and responds to people’s feelings
1999 (Aibo): Sony’s first consumer robot, an evolving pet dog
2011 (Siri): Apple’s intelligent assistant with voice interface
2014 (Eugene): Passes Turing test by fooling 1/3 of judges
2014 (Alexa): Amazon’s intelligent virtual shopping assistant
2016 (Tay): MS chatbot goes rogue, makes offensive comments

AI History from a Personal Perspective

THEME ARTICLE: EXPERT SYSTEMS: COMMERCIALIZING ARTIFICIAL INTELLIGENCE

An Artificial Intelligence Odyssey: From the Research Lab to the Real World

Peter E. Hart, Consultant, USA

In the mid-1960s, I began a journey that started in the lab working on basic research projects in artificial intelligence. As AI evolved, and particularly with the identification in the mid-1970s of expert systems as a promising technology, my attention turned to practical applications. In this narrative, I sketch a personal trajectory that mirrors the historical development of AI from its first stirrings to the beginning of the era of deep learning. My path included the first theoretical analysis of a widely used pattern classifier, the creation of the world's first robot with artificial intelligence, the development of the first expert system that provably solved an economically important problem, the founding of the first AI research lab in a commercial company, the founding of an expert systems company, and the restarting of a corporate research center. My account pays special attention to Syntelligence, the expert systems company that represented my most-focused effort to bring AI into the real world.

IEEE Annals of the History of Computing, Vol. 44, No. 1, pp. 57-72, February 2022.

The Rise and Fall of Expert Systems

THEME ARTICLE: EXPERT SYSTEMS: COMMERCIALIZING ARTIFICIAL INTELLIGENCE

The Expert Systems Business: How It Grew and Died

Paul Harmon, *Business Process Trends, Las Vegas, NV, 89130, USA*

This article describes the computer environment in the 1980s, the birth of the expert systems software business, the proliferation of companies created to sell expert systems products and services, their various products and business strategies, and finally their collapse and demise in the early 1990s. We describe the marketing strategies adopted by various vendors and the various market niches that evolved during the late 1980s and early 1990s. We offer speculation on the reasons for the growth and collapse of the expert systems market in the 1990s.

EXPERT SYSTEMS AS A NEW TECHNOLOGY

Expert systems applications were first developed at Stanford University in the 1970s. Other work with knowledge systems of various kinds

meningitis problems and prescribing appropriate drug treatments [1].

MYCIN was originally written in LISP, a computer language especially developed to support programming that involved list processing, which greatly facilitates

IEEE Annals of the History of Computing, Vol. 44, No. 1, pp. 31-43, February 2022.

How an Expert System Works

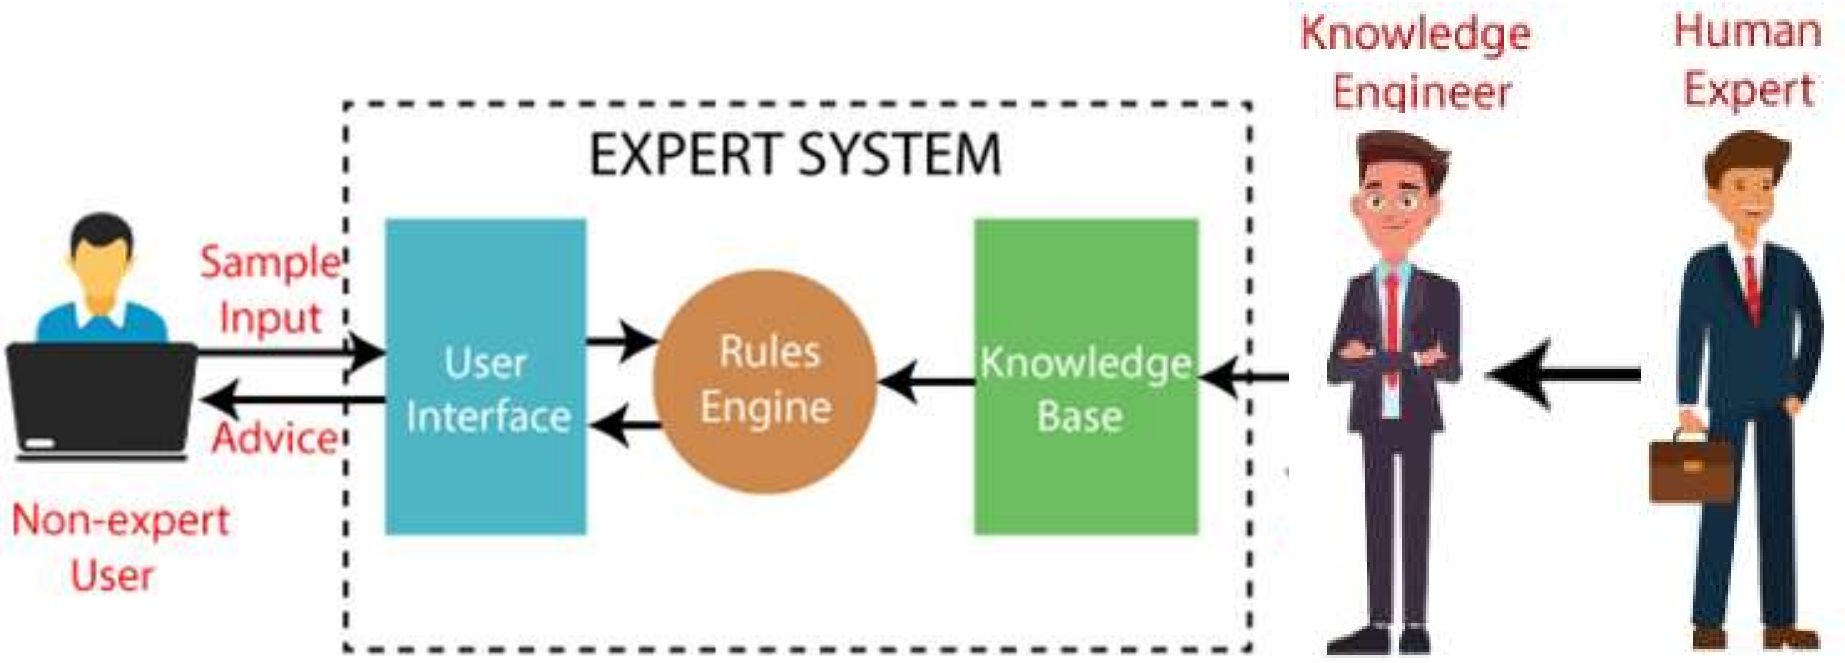

https://www.javatpoint.com/expert-systems-in-artificial-intelligence

No longer called “Expert Systems,” but ideas are still in use:

TurboTax: Asks you questions and prepares your tax return
Medical diagnosis: Deduces illness from symptoms & history

Types of Machine Learning

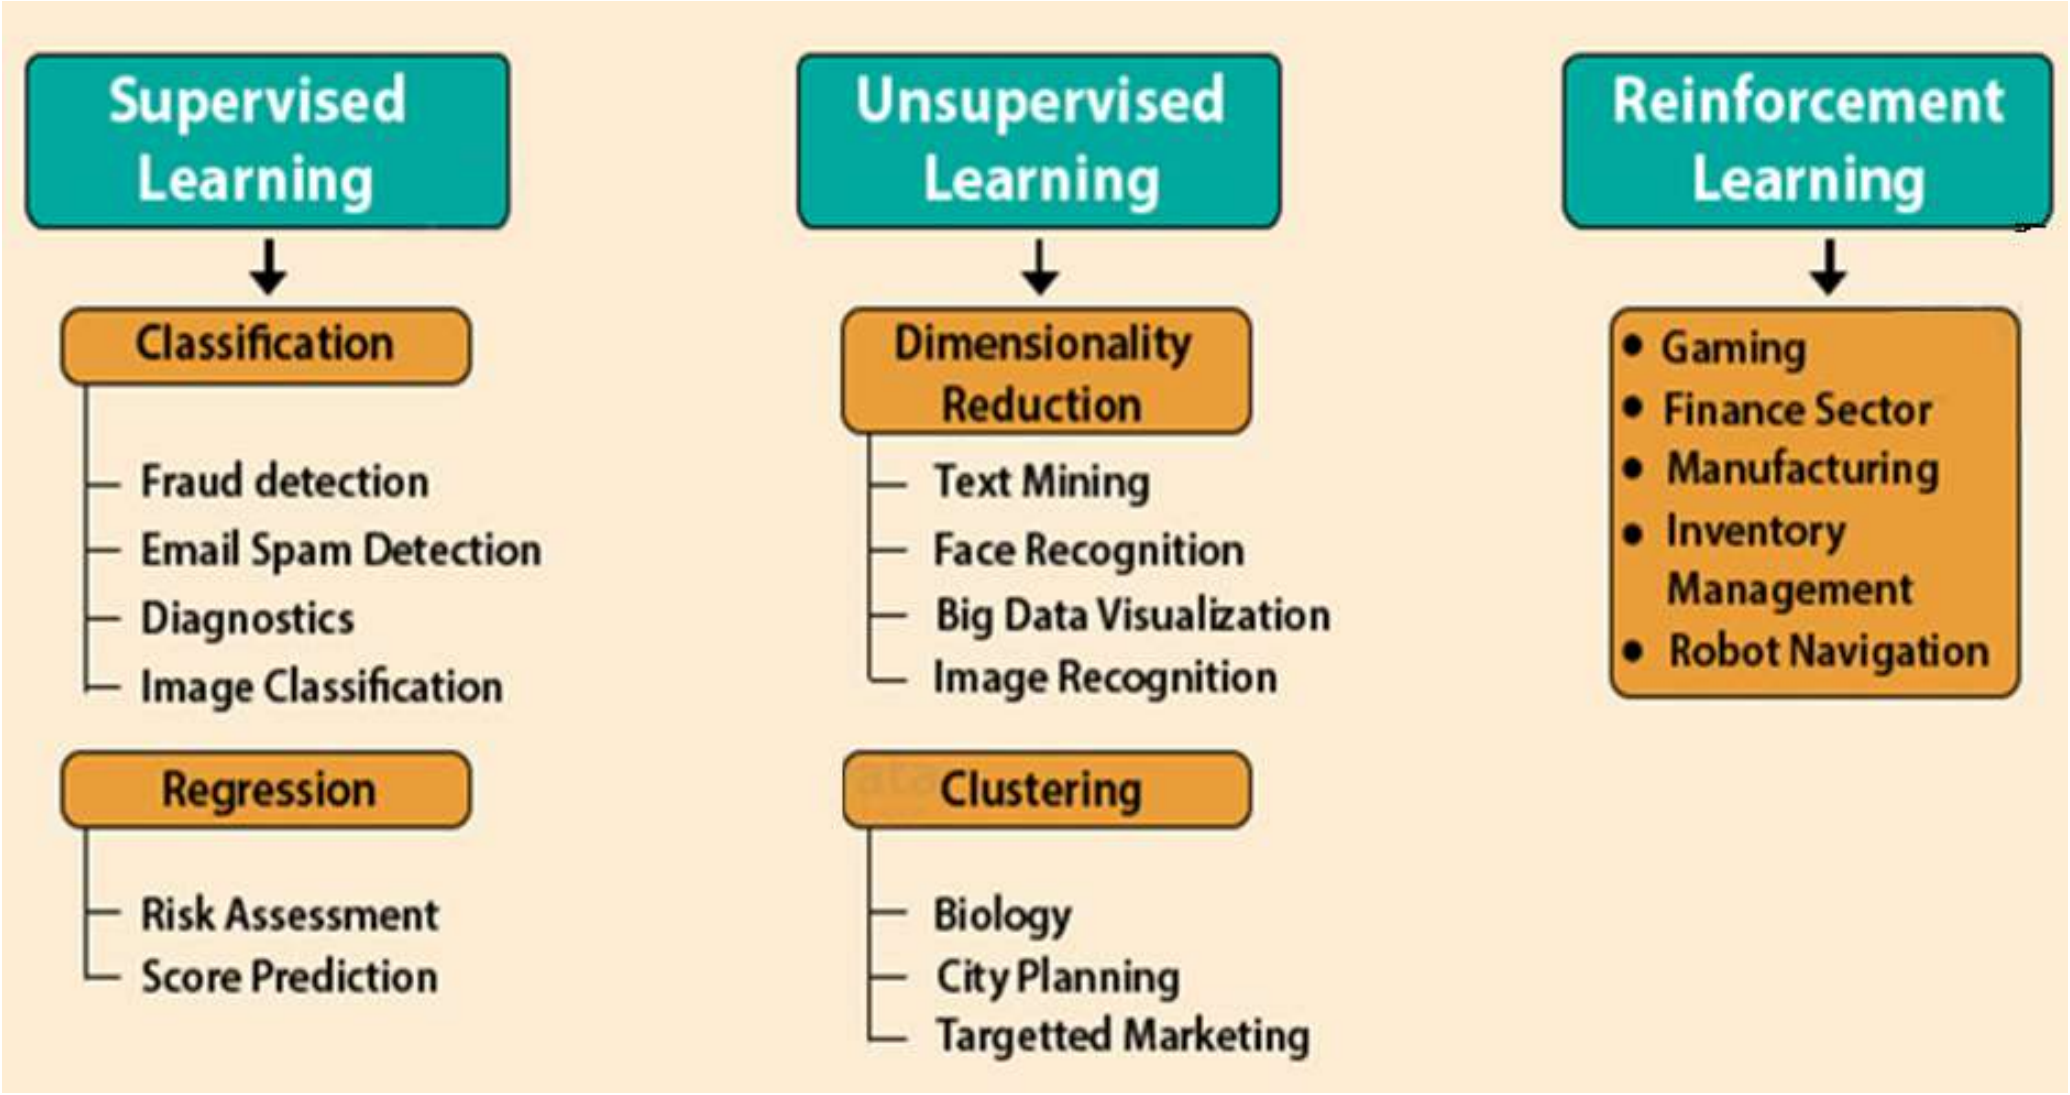

https://data-flair.training/blogs/types-of-machine-learning-algorithms/

Supervised Machine Learning

Example: Find news articles that a user might like.
Agent/Algorithm identifies candidate articles
Supervisor comments on the outcome and approach
Agent/Algorithm adapts

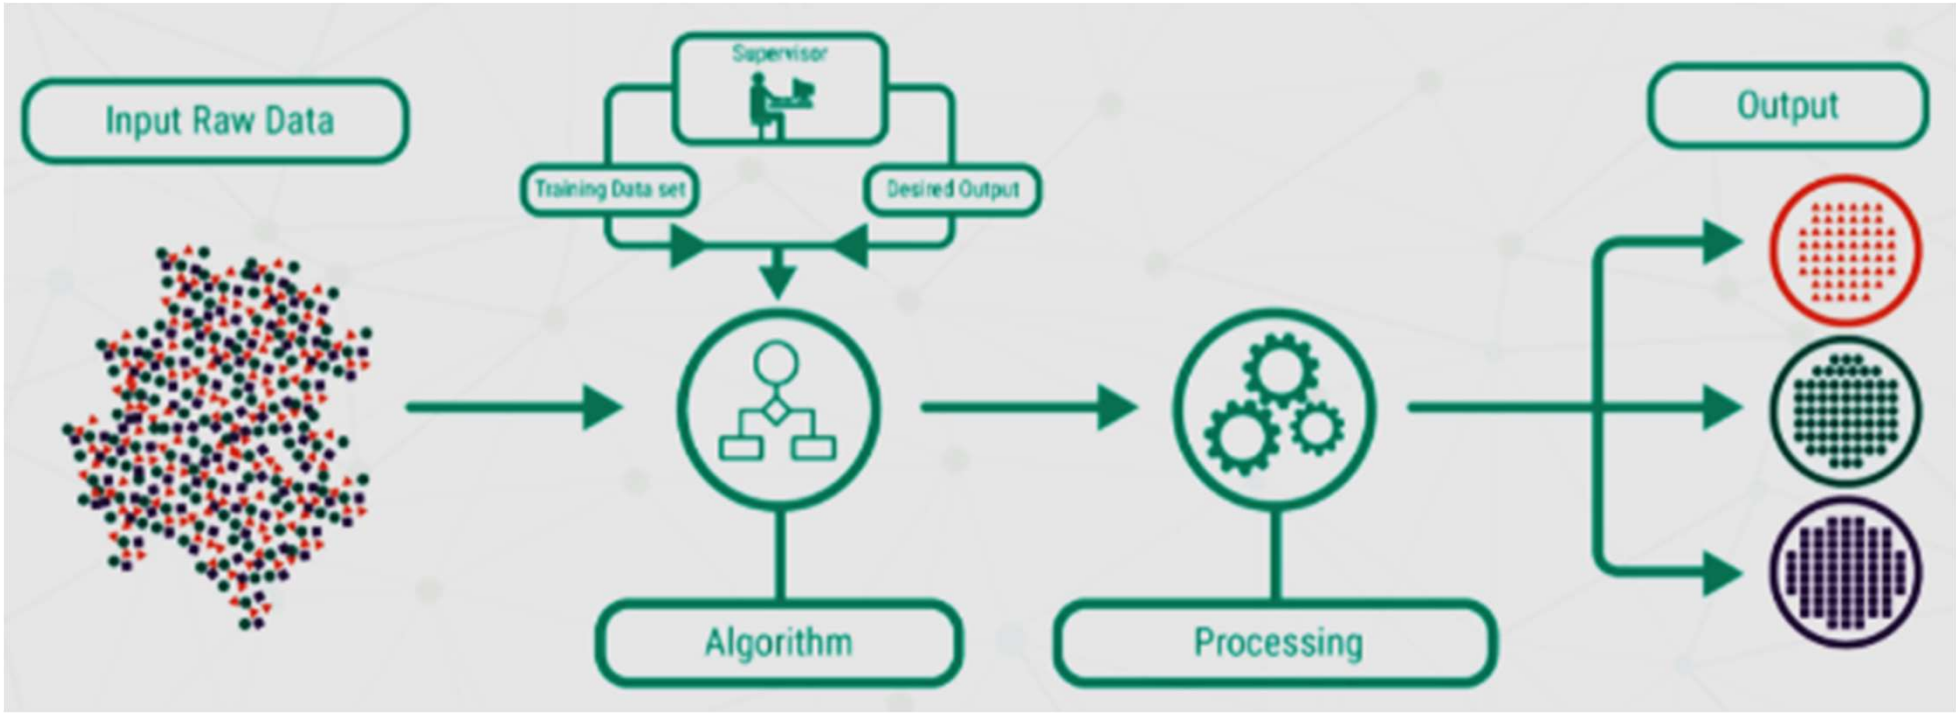

https://bigdata-madesimple.com/machine-learning-explained-understanding-supervised-unsupervised-and-reinforcement-learning/

Supervised Machine-Learning Example

Learning with a teacher

Spam filter

Starts with a basic algorithm (keywords, word patterns, typos, grammar, subject line, greeting, time pressure, …)

When you go to your spam folder and pull out an item as not being spam, you are teaching the system by example what to look for

Early on, the algorithm may act conservatively in classifying something as spam to avoid misclassifying legit mails.

As it learns more, it becomes more aggressive and more accurate

Unsupervised Machine Learning

Example: Find news articles that a user might like.
Agent/Algorithm looks at articles the user has read
Agent/Algorithm extracts key features of the articles read
Agent/Algorithm selects articles sharing same features

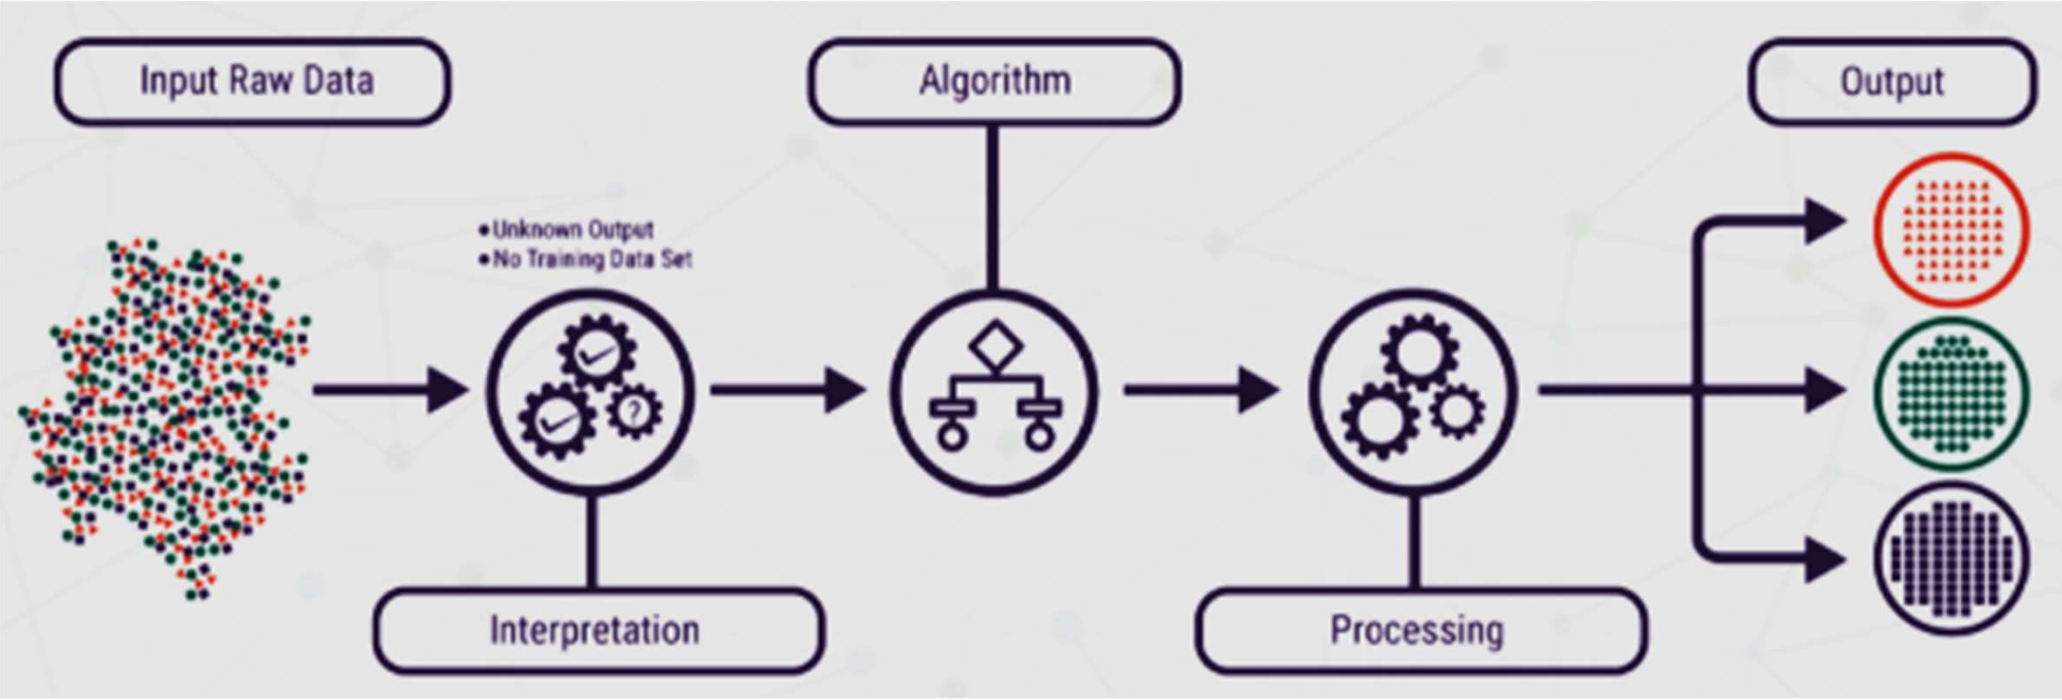

https://bigdata-madesimple.com/machine-learning-explained-understanding-supervised-unsupervised-and-reinforcement-learning/

Unsupervised ML Example

Initial estimator:

$h(x) = 12.00 + 0.20x$

$h(x) = 13.12 + 0.61x$

$h(x) = 15.54 + 0.75x$

Generally, estimation functions are much more complicated and involve many variables, and thus many coefficients (100s)

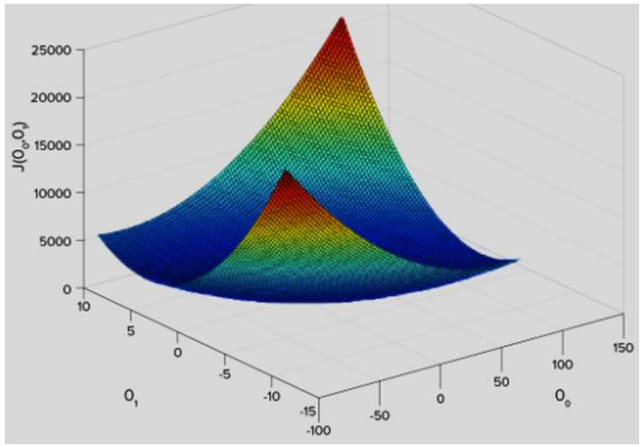

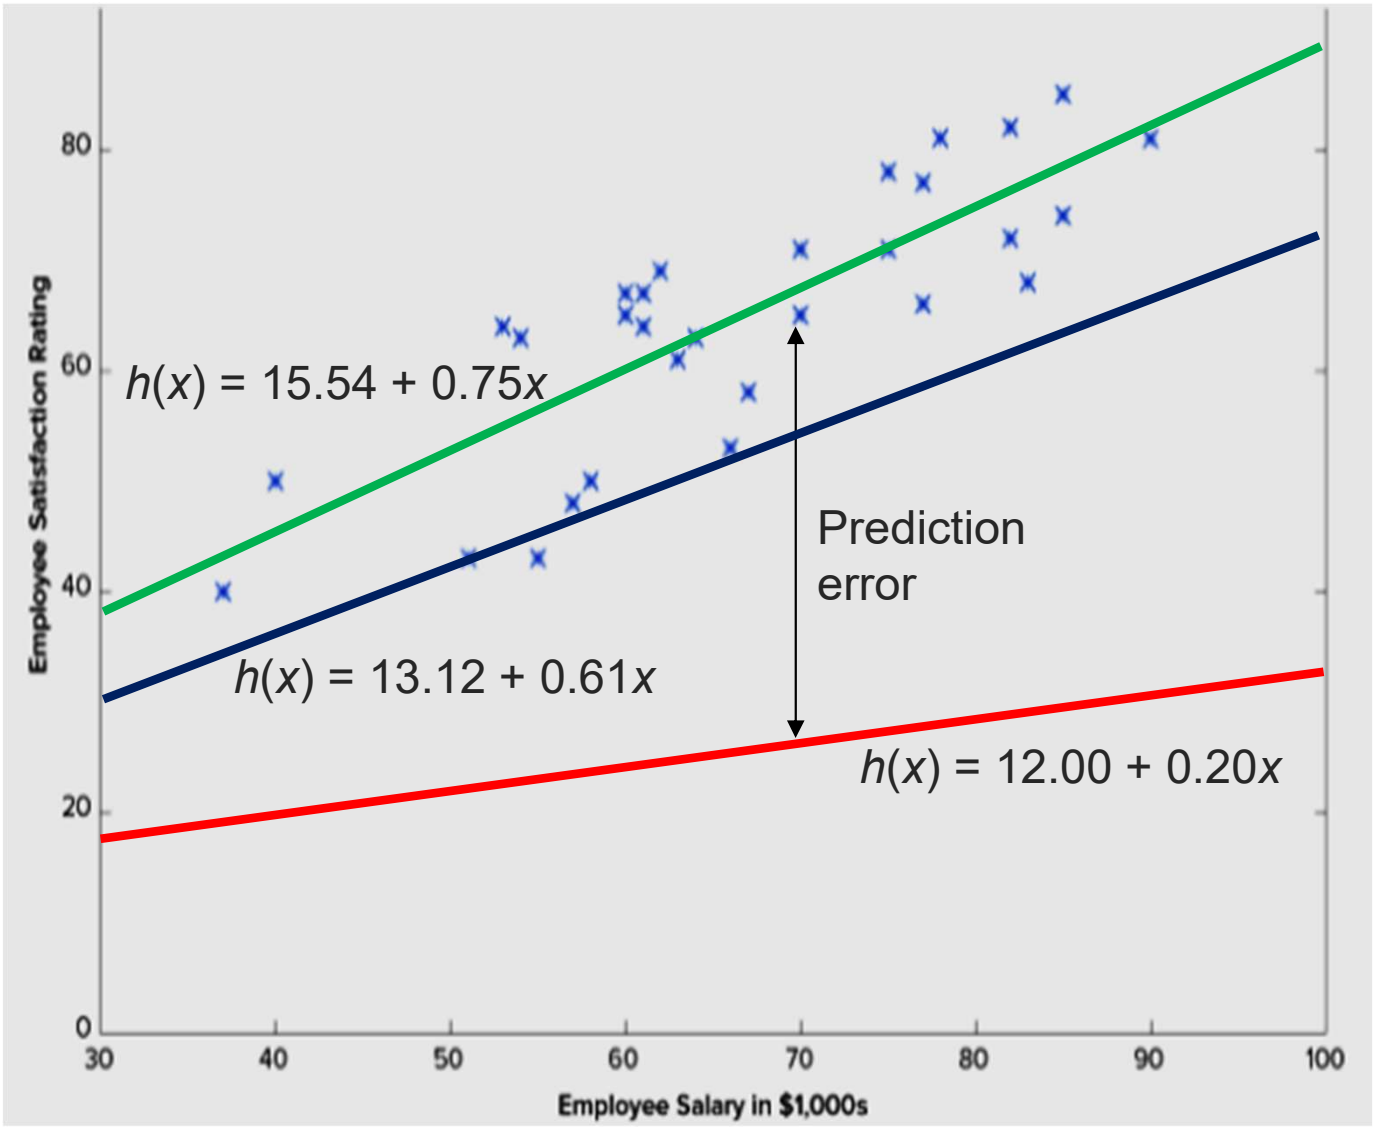

https://www.toptal.com/machine-learning/machine-learning-theory-an-introductory-primer/

Reinforcement Learning

Example: Find news articles that a user might like.
Agent/Algorithm selects based on articles the user has read
Agent/Algorithm receives feedback on its selections
Agent/Algorithm improves by adjusting selection criteria

https://bigdata-madesimple.com/machine-learning-explained-understanding-supervised-unsupervised-and-reinforcement-learning/

Reinforcement-Learning Example

News recommendation:

Reinforcement can come from user reaction or behavior

Reading the recommended story
Amount of time spent on it
Clicking on links within the story

Each story is characterized along many axes

Info in user profile:

US news	0.7
Iran news	0.5
Sports news	0.3
Opinion	0.4

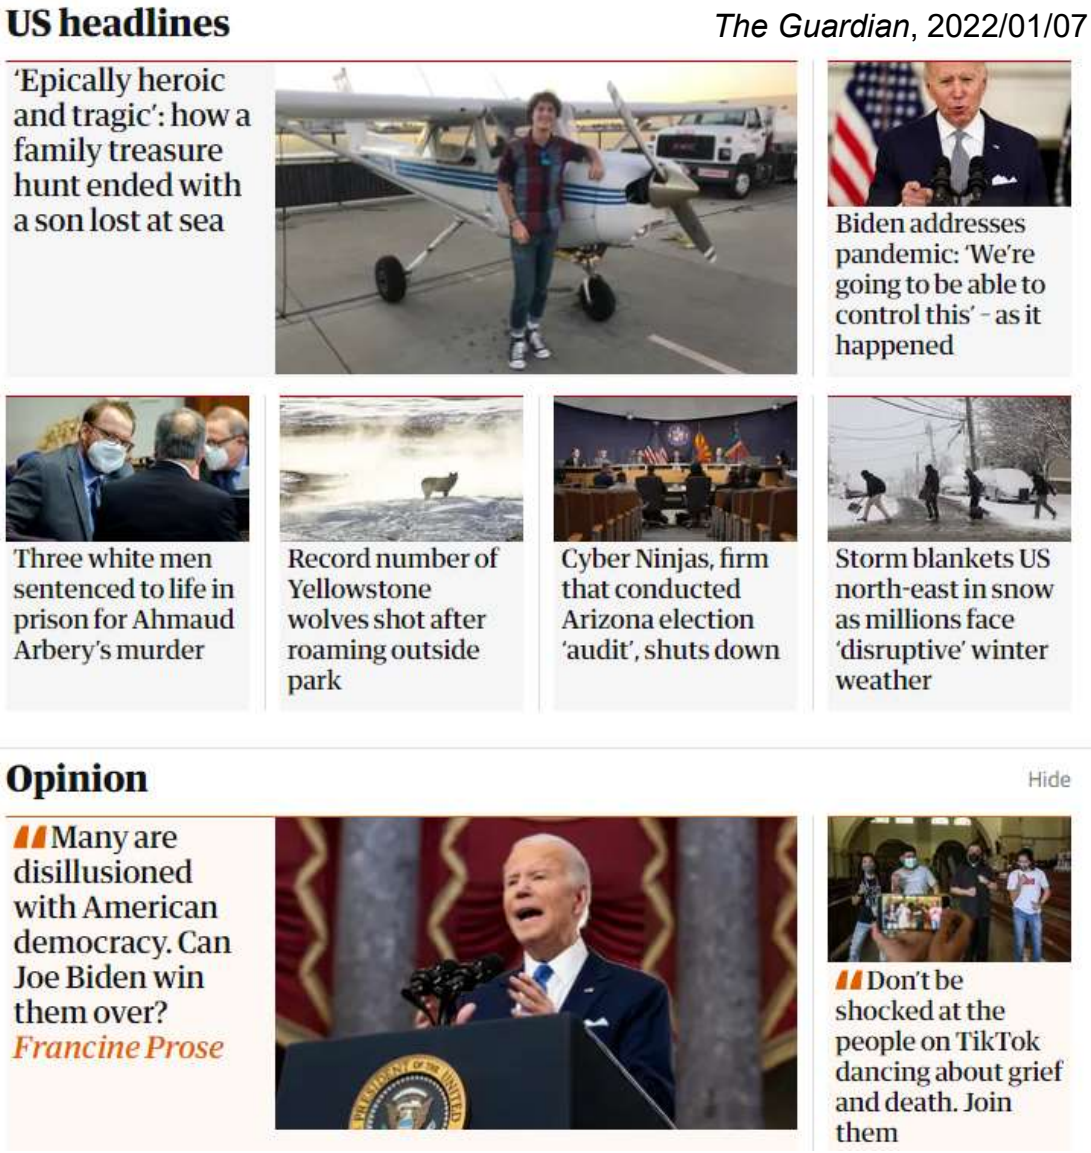

What Is Being Human?

Life 1.0: Hardware & software modified through evolution (slowly)

Life 2.0: Hardware modified slowly, software modified/augmented at an accelerated pace through learning

Life 3.0: Both hardware & software modified/augmented by humans

"… why should our simplest path to a new technology be the one that evolution came up with, constrained by requirements that it be self-assembling, self-repairing, and self-reproducing?
Evolution optimizes strongly for energy efficiency because of limited food supply, not for ease of construction or understanding by human engineers."

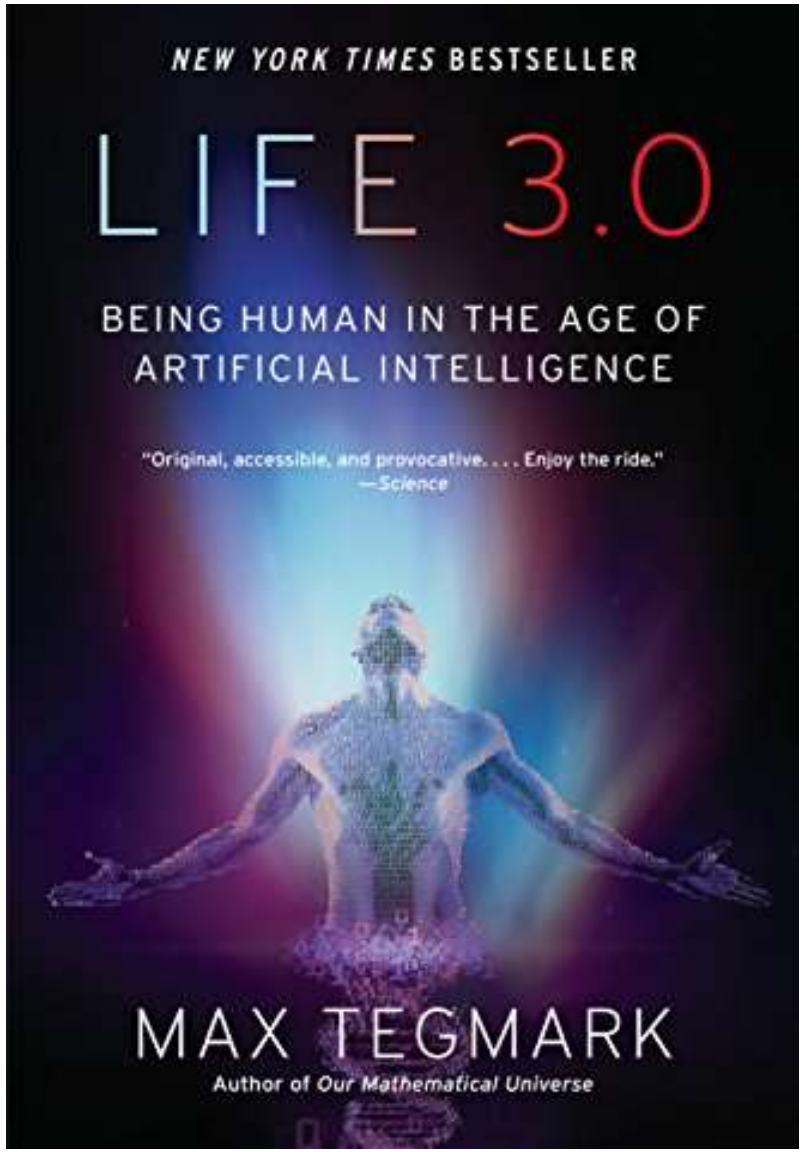

The Singularity Is Near: Is It Really?

Ray Kurzweil: Google's futurist
- AI will be smarter than humans
- Humans will transcend biology

Singularity: Point in human history, where the future of humanity will become unpredictable, because it's changing so quickly

Nanorobots (blood cells on steroids) will fix your body from the inside, but they can go haywire & kill/maim you

There are people alive today who will live forever

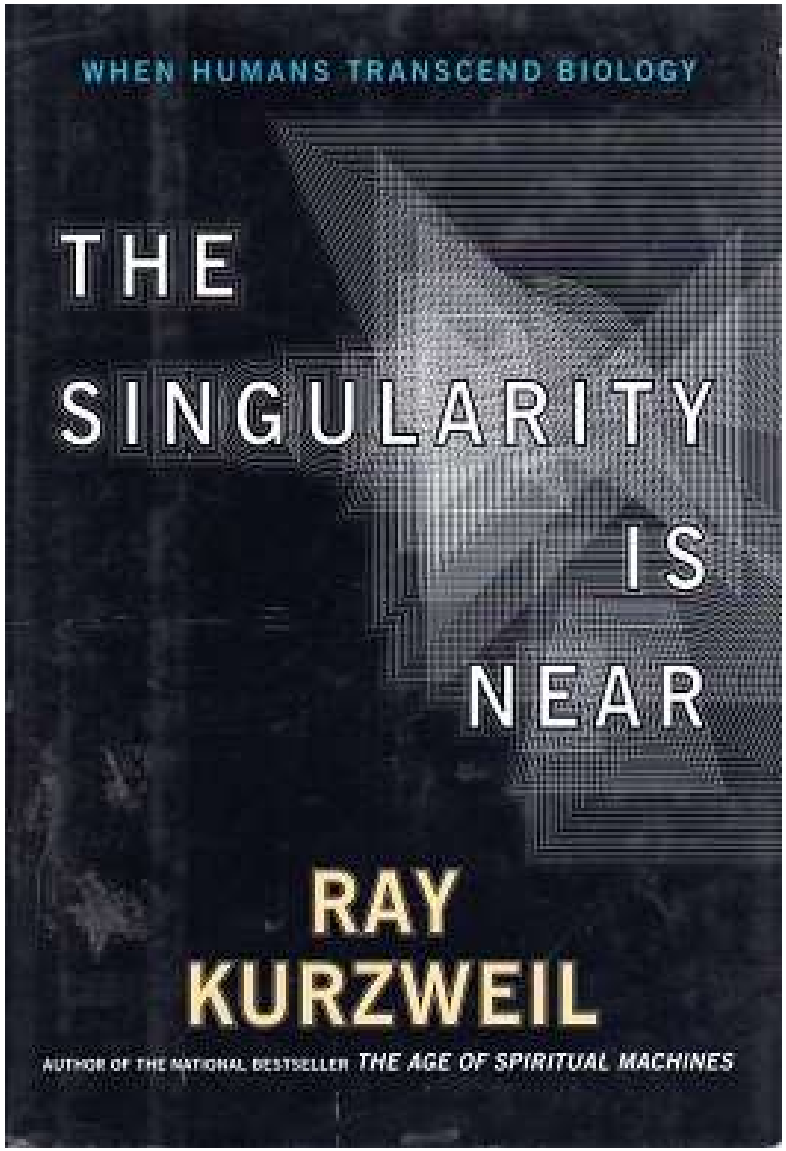

Expanding the Human Mind

Ray Kurzweil's other book:
How to Create a Mind

Our Three Brains
Reptilian brain (basal ganglia)
Limbic system (first mammals)
Neocortex
- Hierarchical layers, dealing with increasing levels of abstraction

Brain + the Cloud:
Smartphones are gateways to the Cloud … Put gateways to the Cloud in our brain (extend the neocortex)

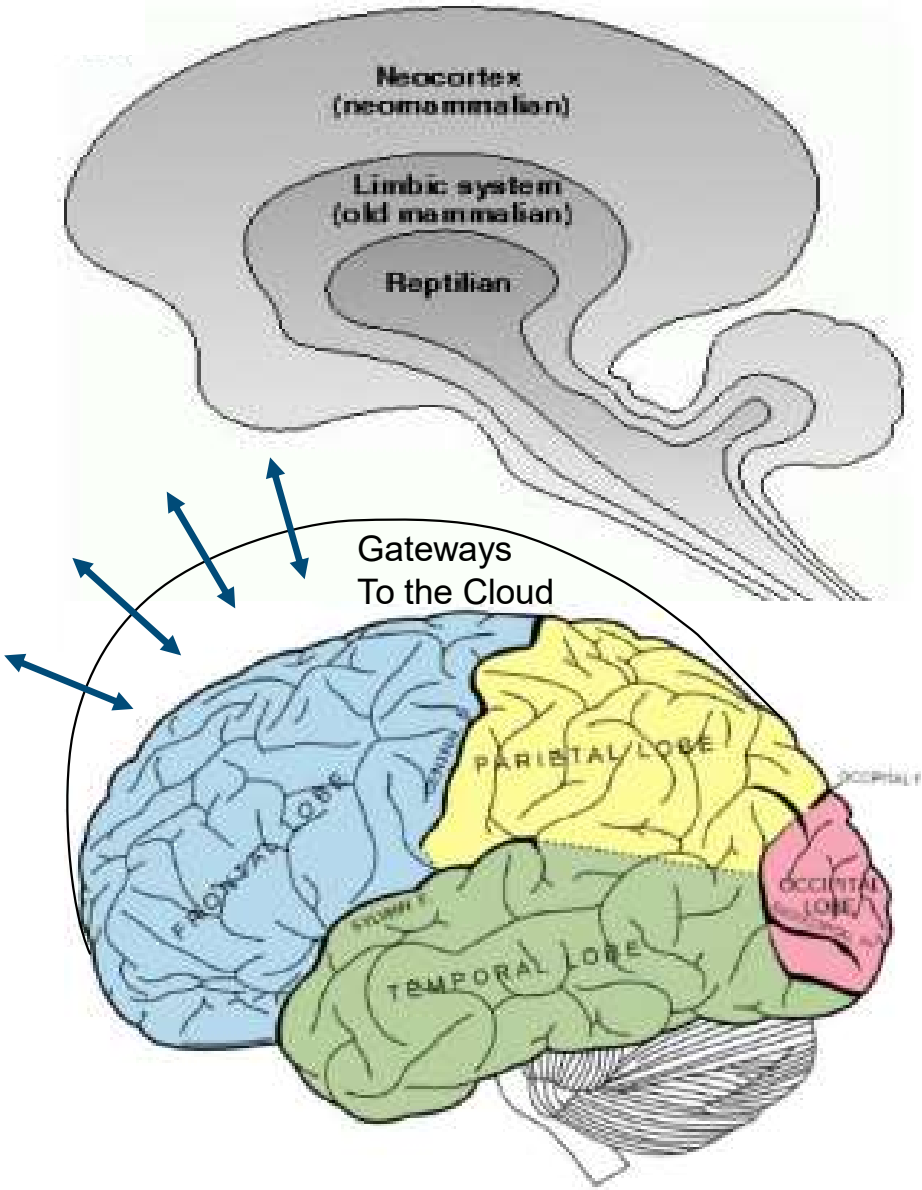

Man Controls Computers with His Mind

***New York Times*, May 15, 2022:**

Dennis DeGray, paralyzed since 2006, has regained a semblance of control over his body via a brain-computer interface (BCI).

Implanted in 2016, BCI enables DeGray to move a cursor on a computer screen by thought, via machine-learning algorithms.

DeGray has learned to control videogames, robotic limbs, and a simulated aerial drone.

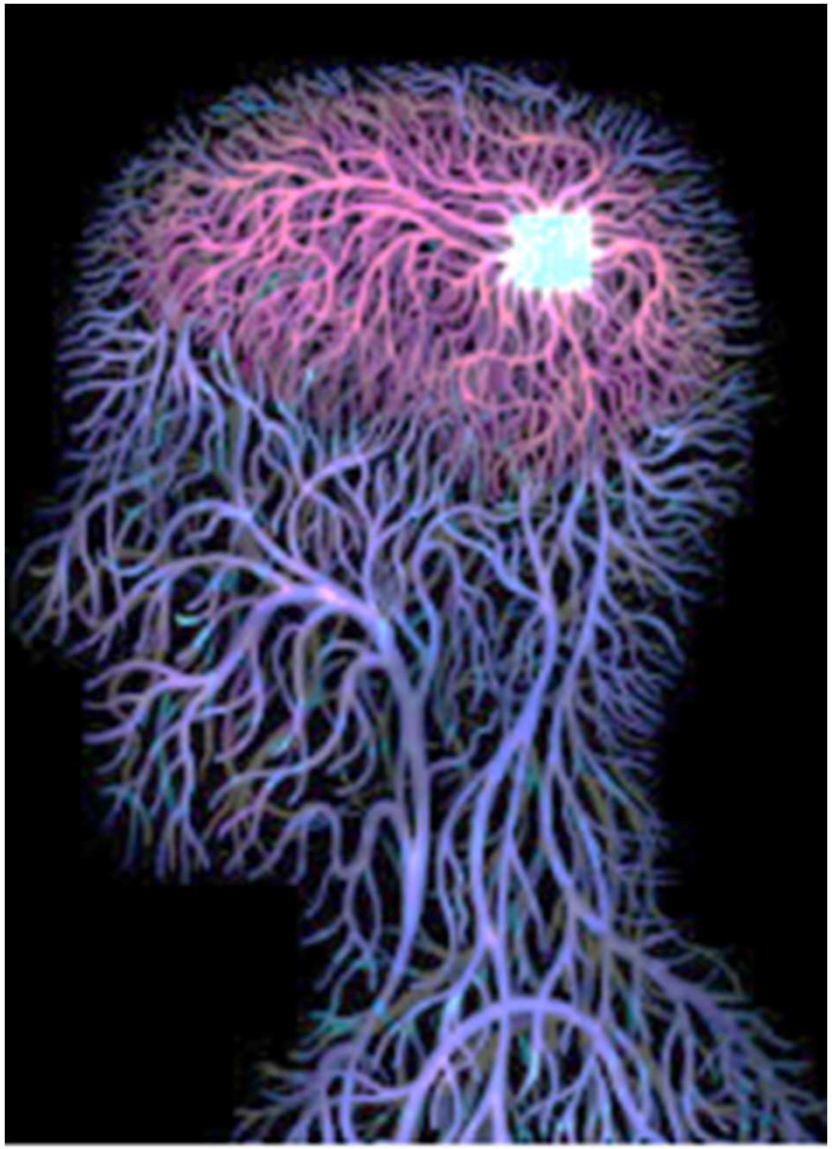

Skepticism about the Singularity

Table of Contents:

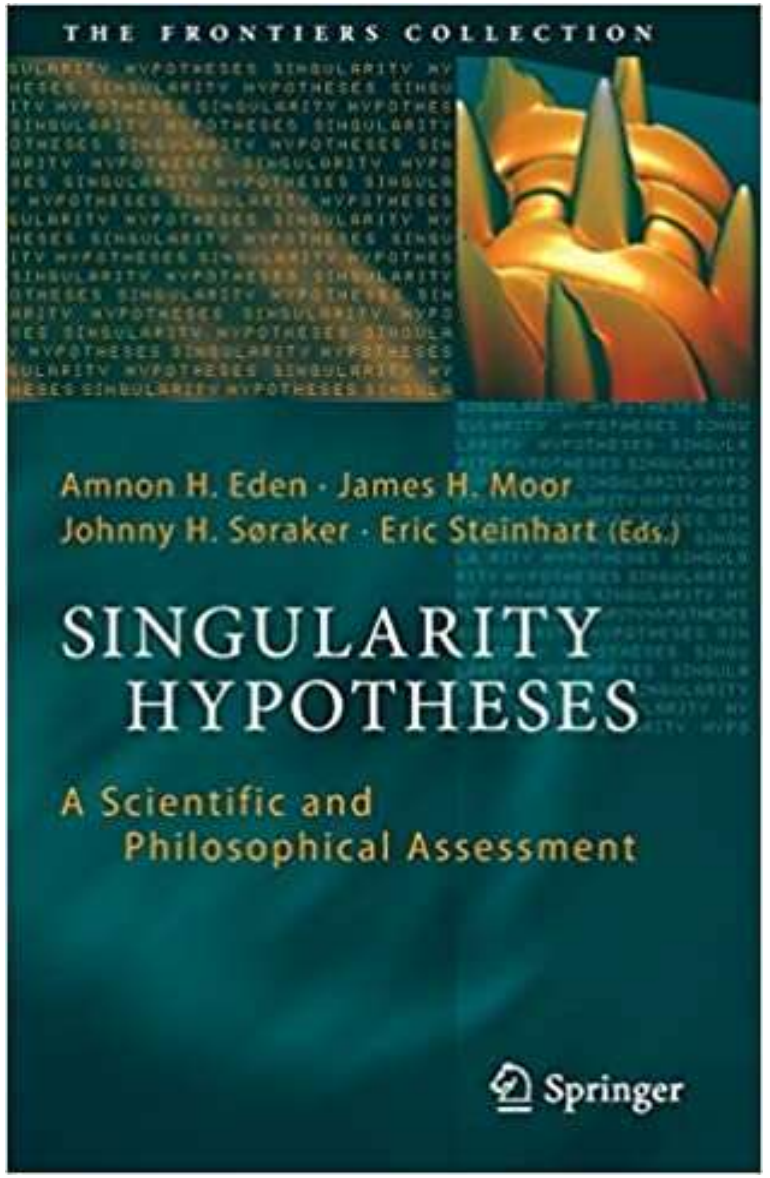

Narrow AI vs. General AI

Narrow AI deals with performing a narrow range of tasks:

Uber: Matching riders to drivers

Airport logistics: Gate assignment

Game-playing: Chess, GO

General AI handles difficult tasks with human-like flexibility:

Human-level general conversation

Expert-level dialog: Astrophysics

Machine translation: Requires Turing-test-level capabilities

Types of AI

Based on ability

- Narrow AI
- General AI
- Super AI

Based on functionality

- Reactive machines
- Limited memory
- Theory of mind
- Self-awareness

I, Robot (1950)

Asimov's laws of robots:

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.

2. A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.

3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

First introduced in the 1942 short story "Runaround," and later included in the *I Robot* collection

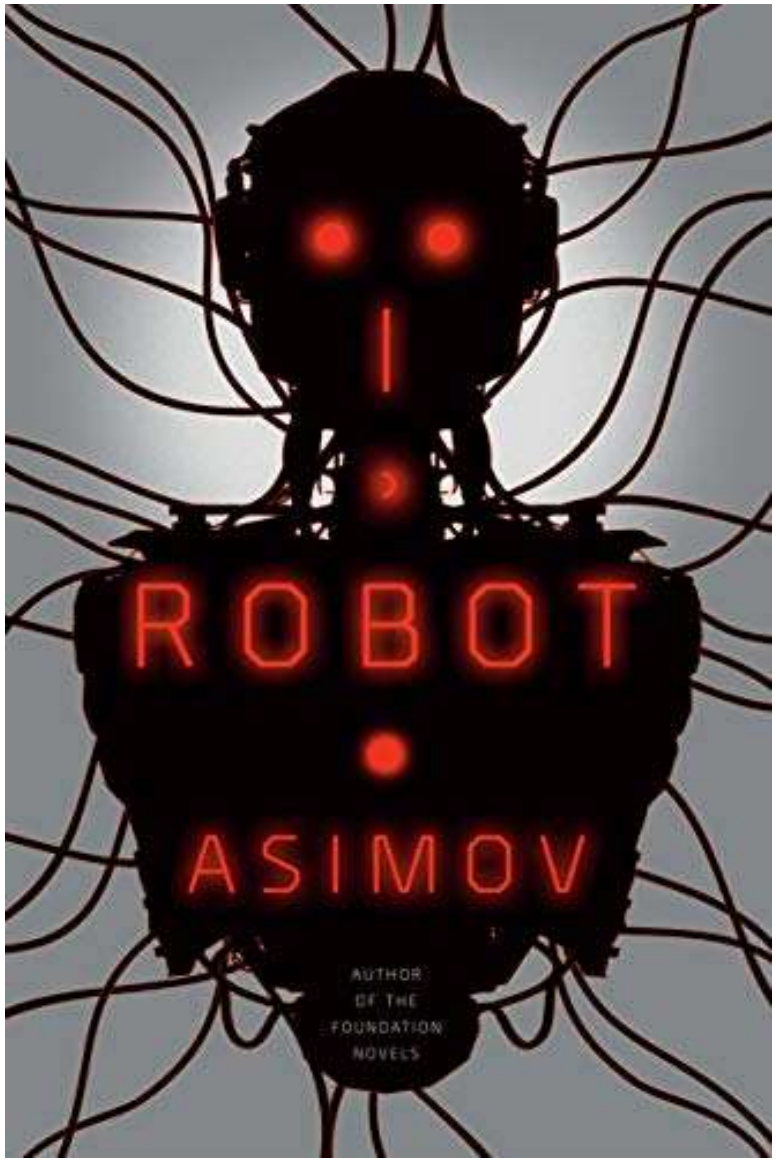

Robotic Wars

Killer robots are no longer science-fiction

Sci-fi makes us think of killer, human-like robots

Air, land, and sea drones already in use

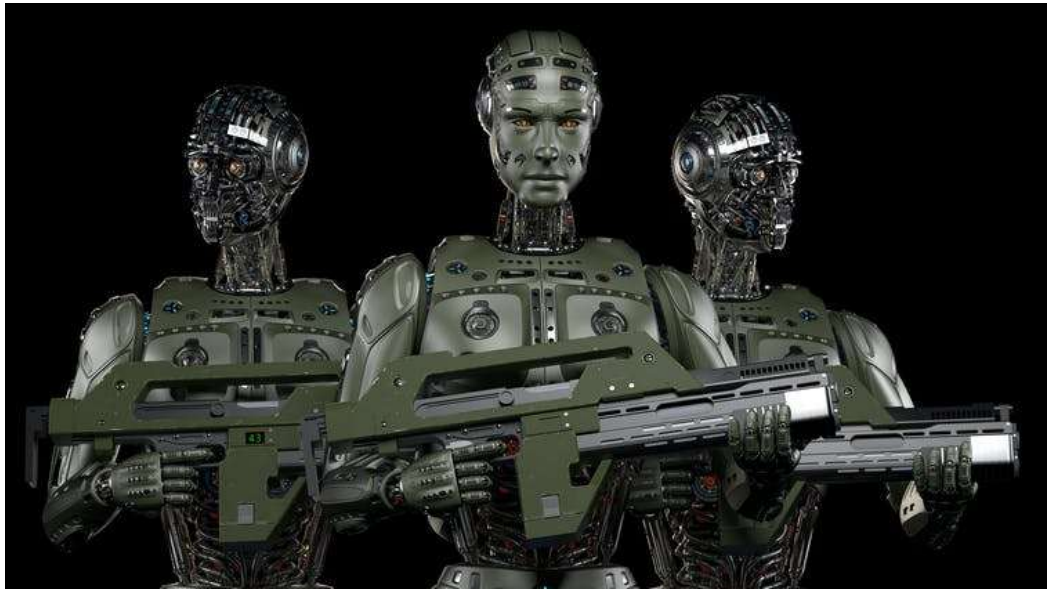

AI's Real Worst-Case Scenarios

From *IEEE Spectrum*, Issue of January 2022:

We are entering dangerous and uncharted territory with the rise of surveillance and tracking through data, and we have almost no understanding of the potential implications.

"AI doesn't have to be sentient to kill us all. There are plenty of other scenarios that will wipe us out before sentient AI becomes a problem." ~ Malcolm Murdock, machine-learning engineer and author of the 2019 novel *The Quantum Price*

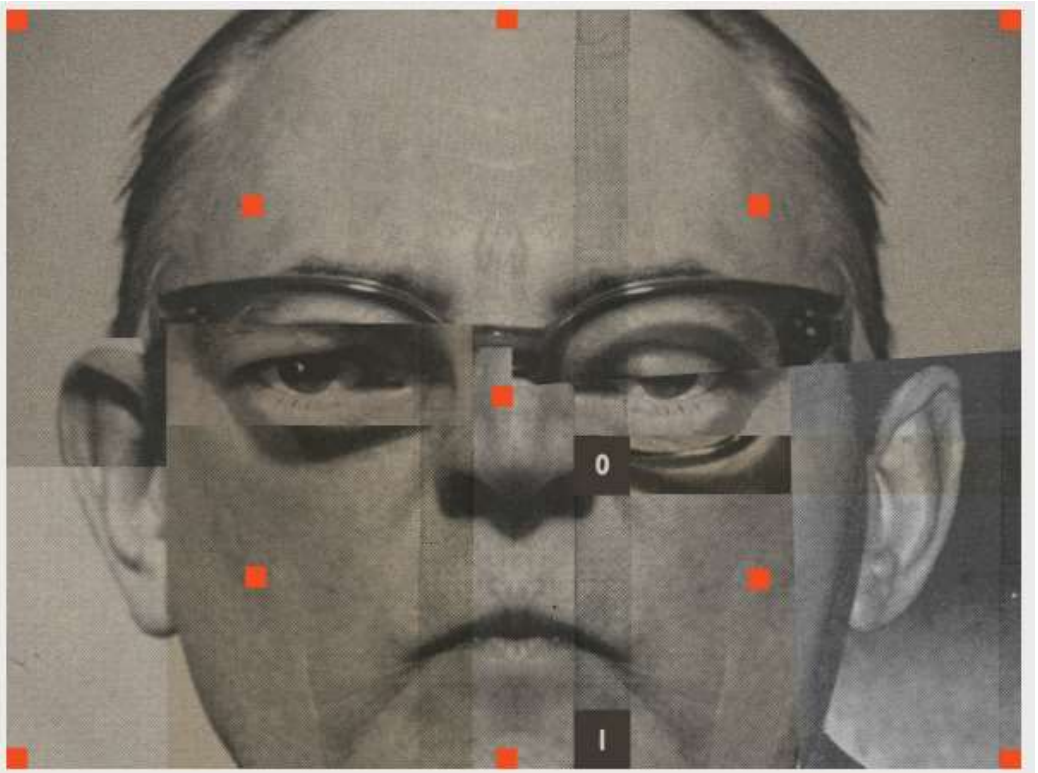

ARTIFICIAL INTELLIGENCE

AI's Real Worst-Case Scenarios › Who needs Terminators when you have precision clickbait and ultra-deepfakes?

BY NATASHA BAJEMA

Hollywood's worst-case scenario involving artificial intelligence (AI) is as familiar as any trope in blockbuster movies: Machines acquire humanlike intelligence, achieving sentience, and inevitably turn into evil overlords that attempt to destroy the human race. This narrative capitalizes on our innate fear of technology, a reflection of the profound change that often accompanies new technological developments.

However, as Malcolm Murdock, machine-learning engineer and author of the 2019 novel *The Quantum Price*, puts it, "AI doesn't have to be sentient to kill

There Are Bad News and Good News

If you're uncomfortable with how much Google knows about you, imagine if it had your brain data as well!


FIRST WIN FOR THE NEURORIGHTS CAMPAIGN

Chile plans to regulate all neurotech and ban the sale of brain data

BY ELIZA STRICKLAND

The Five Neurorights:

From *IEEE Spectrum*, Issue of January 2022

1. **The right to mental privacy**
2. **The right to personal identity**
3. **The right to free will**
4. **The right to equal access to mental augmentation**
5. **The right to protection from bias**

Ethics for Technology in General

Can you withdraw tech support when a bionic body part becomes obsolete?

Clinical ethics:

1. **Beneficence**
2. **Nonmaleficence**
3. **Autonomy**
4. **Informed consent**
5. **Truth-telling**
6. **Confidentiality**
7. **Justice**

https://www.karger.com/Article/FullText/509119

***IEEE Spectrum* magazine cover, March 2022**

Deep Learning

A most talked-about domain of machine learning (How deep is deep?)

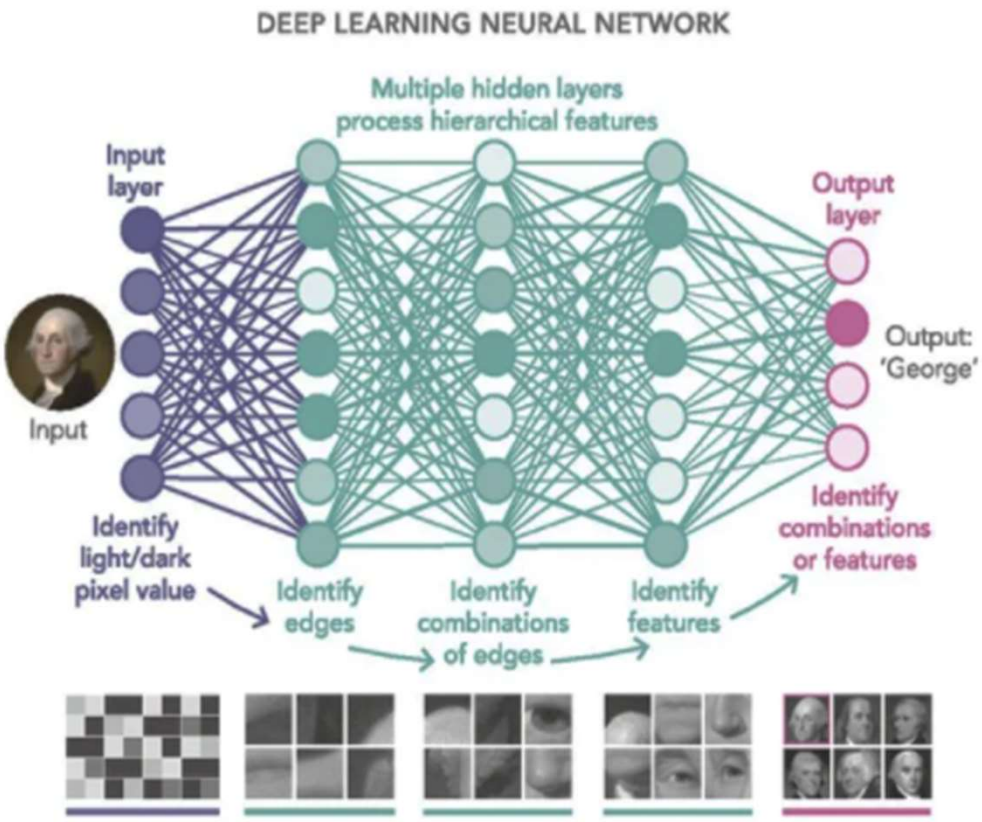

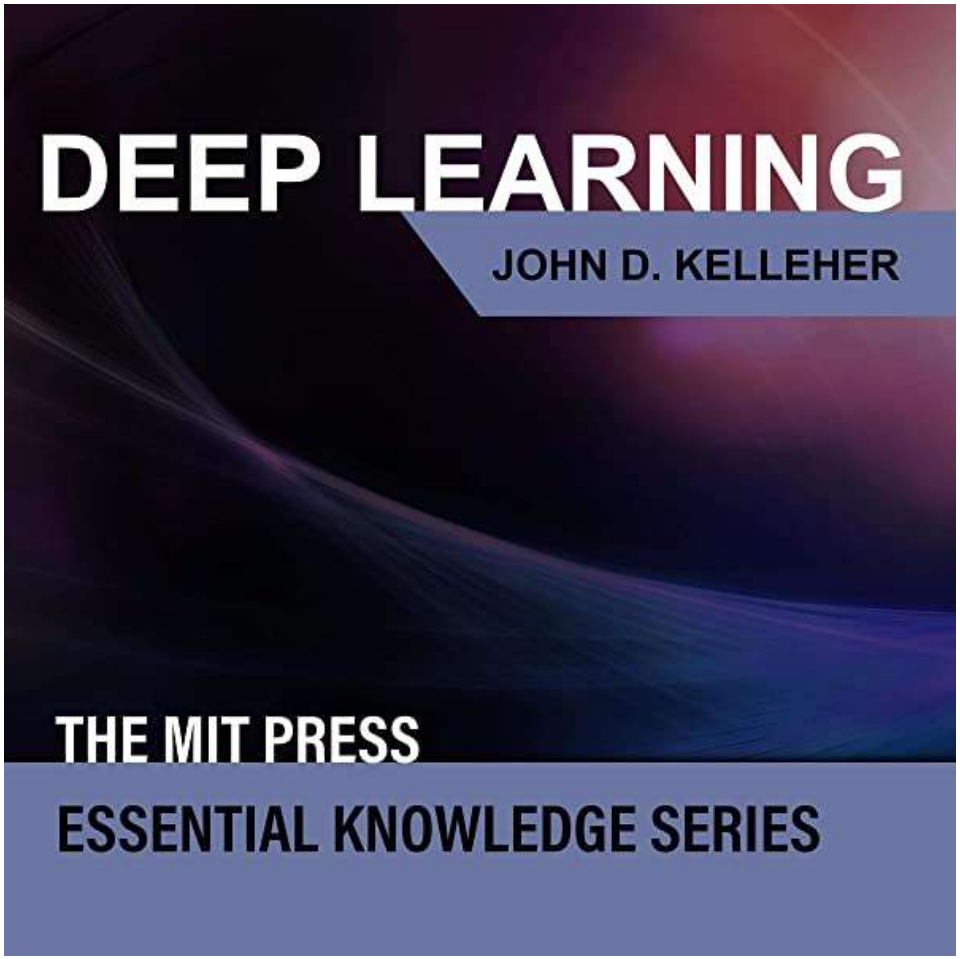

Brief intro to deep learning, by IBM Cloud Education:
https://www.ibm.com/cloud/learn/deep-learning

Simple Example of Deep Learning

Table 1.1. A dataset of loan applicants and their known credit solvency ratings

ID	Annual Income	Current Debt	Credit Solvency
1	$150	-$100	100
2	$250	-$300	-50
3	$450	-$250	400
4	$200	-$350	-300

$S = f(I, D)$

Solvency is a function of income & debt

Example:
$S = 2I + 2D$

Each node has weights for its inputs and a threshold

Fig. 1.2

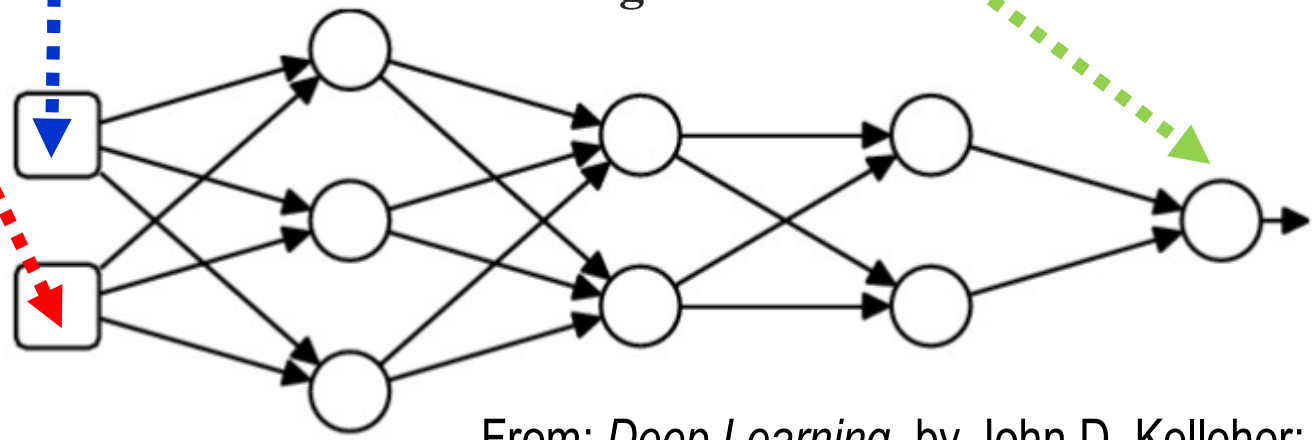

From: *Deep Learning*, by John D. Kelleher:

Resources for AI Training

According to Robert Ober (NVIDIA):

- Compute resources for AI training grow ~10x per year
- AI clusters use 100s to 1000s of synchronous GPUs

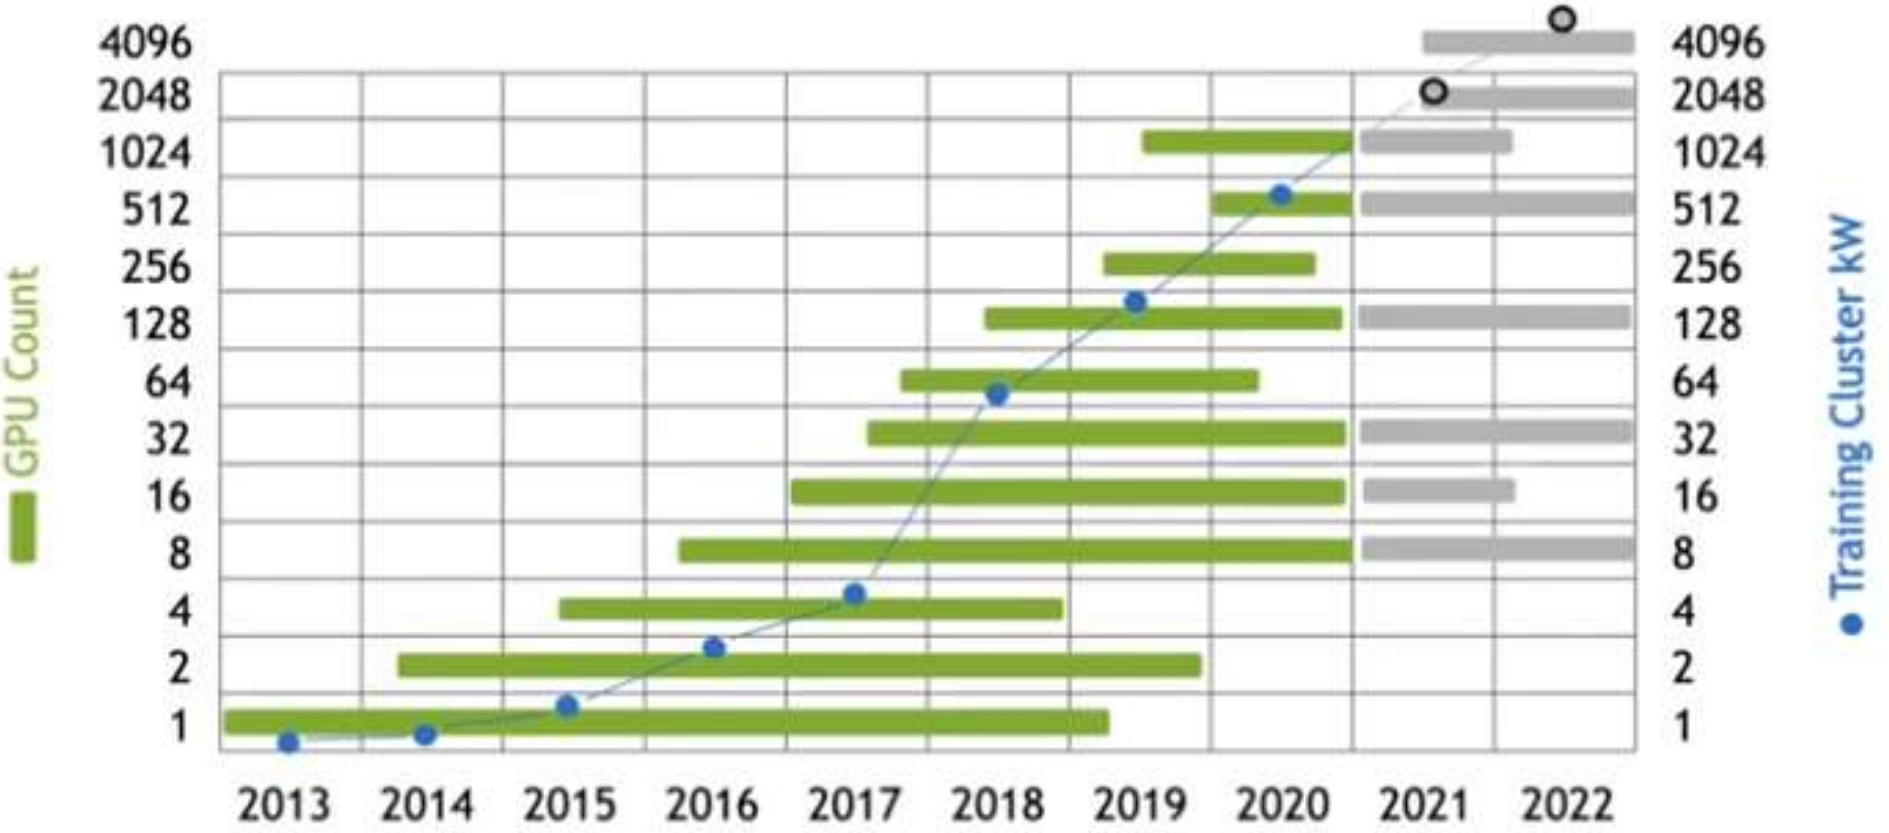

18-minute talk, “AI Cluster Trends”: https://www.youtube.com/watch?v=ulOFL_KCCtU

Trustworthy AI Systems

We need to build AI systems that we can trust

Components of trust include:

Transparency (explainability)
Fairness (absence of bias)
Reliability (availability, consistency)
Privacy (avoiding data leakage)
Safety (not doing harm to people)
Robustness (not easily thrown off)

Many researchers have shown that small changes in input data can lead to widely different decisions

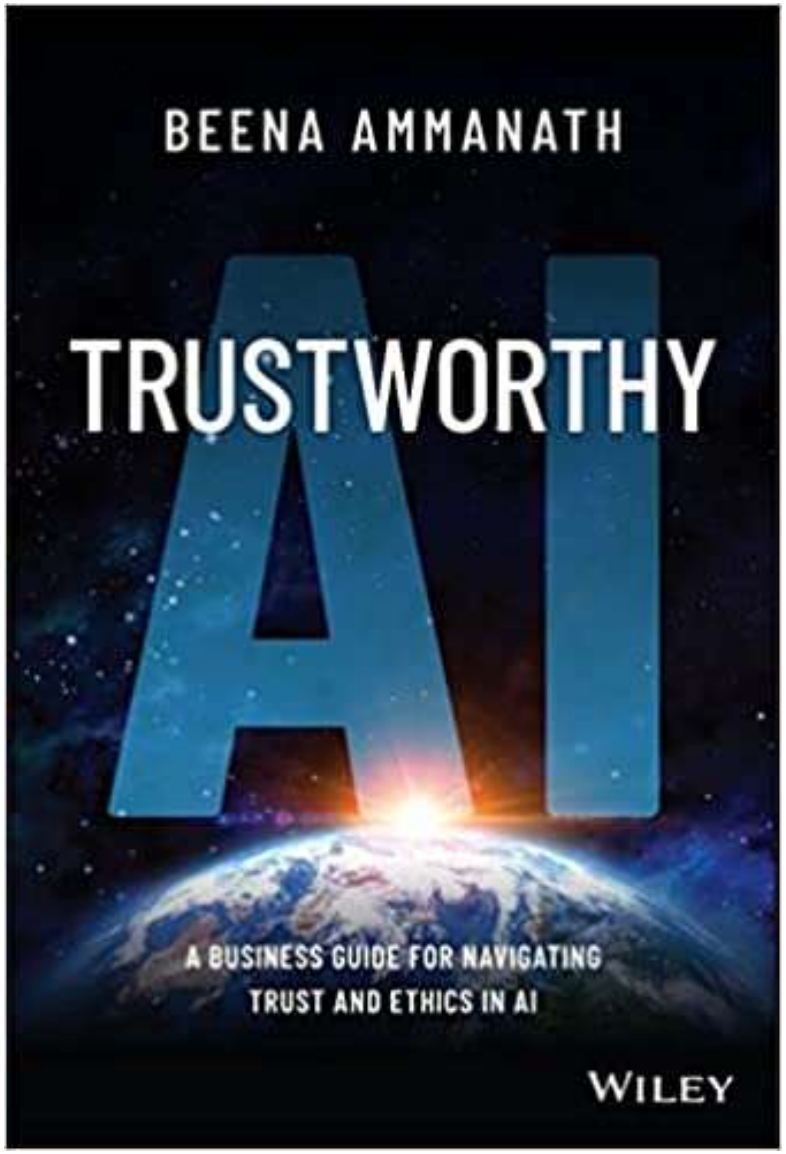

Transparent or Explainable AI

Unless humans can understand how AI systems reach their decisions, systems can go awfully wrong

ImageNet competition debacle

Millions of images to be classified into hundreds of categories
There was much cheating by trying the challenge many times and submitting the best results
(like taking an exam many times and keeping the best grade)

Using the wrong criteria, but achieving correct results

Images of forest, possibly with animals in the foreground
Determine which images have animals in them
If the trees are blurred, then there is likely an animal in front

Can AI Systems Be Creative?

Like all the previous questions we dealt with in this talk, it depends (on your definition of creativity)

AI systems can produce works of art

Paintings that are comparable in "quality" to human art
Completion of Beethoven's unfinished 10th Symphony

AI systems can prove theorems

Help with the solution of the four-color problem
Several new theorems that had eluded mathematicians

AI systems can generate new food recipes

Analyze recipes and combine ingredients in new ways

Business Angst for AI and Going Digital

New book, Nov. 2022:
Paul Leonardi (UCSB TMP)
Tsedal Neeley (Harvard Bus. Sch.)

Useful advice to businesses, but …

Small businesses may be hurt due to their inability to adapt to the new digital/automation landscape

Large businesses can afford to hire AI/ML specialists, data scientists, software/app developers, and cyber-security experts

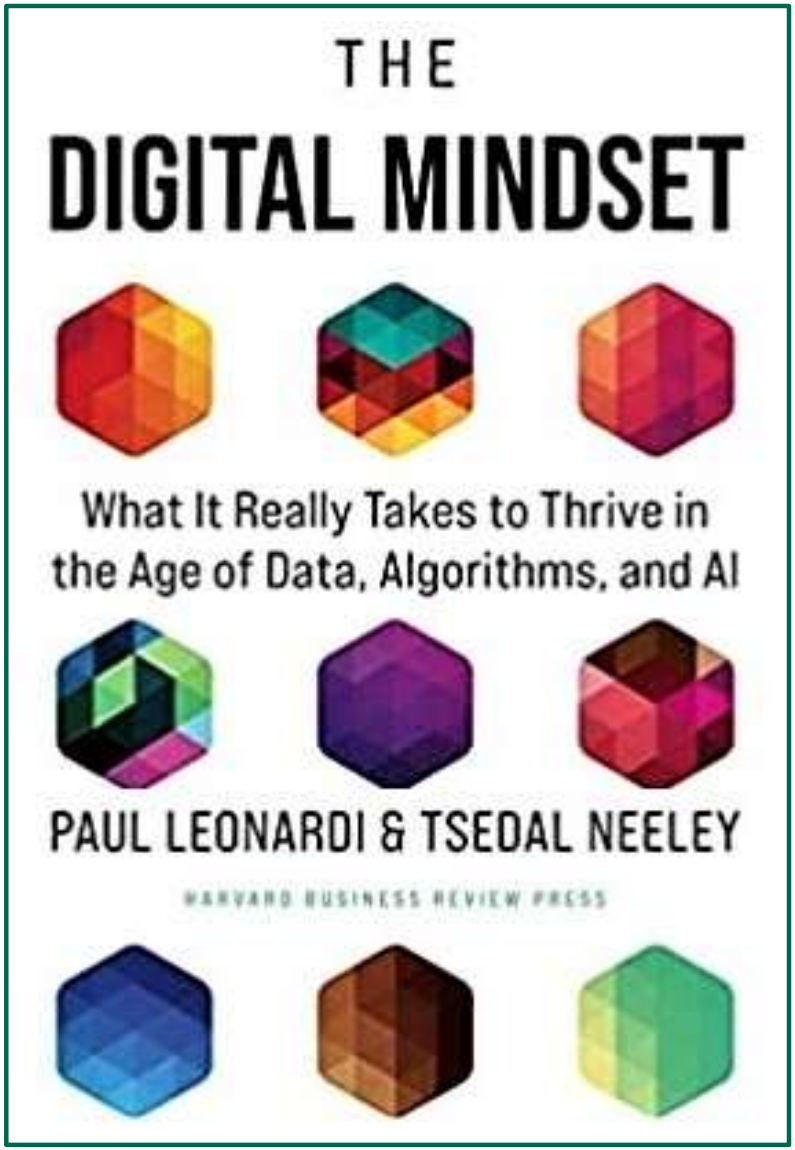

Humor from an AI / ML Specialist

What society thinks I do

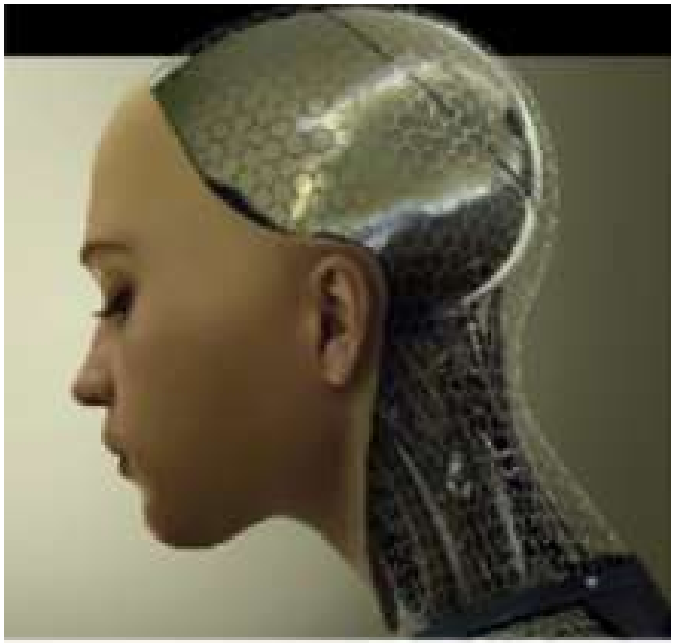

What other programmers thinks I do

$$
\begin{aligned}
&\propto \frac{\prod_{i \neq k} \Gamma(n_{m,(\cdot)}^{i,-(m,n)} + \alpha_i)}{\Gamma((\sum_{i=1}^{K} n_{m,(\cdot)}^{i,-(m,n)} + \alpha_i) + 1)} \prod_{i \neq k} \frac{\Gamma(n_{(\cdot),v}^{i,-(m,n)} + \beta_v)}{\Gamma(\sum_{r=1}^{V} n_{(\cdot),r}^{i,-(m,n)} + \beta_r)} \\
&\times \Gamma(n_{m,(\cdot)}^{k,-(m,n)} + \alpha_k + 1) \frac{\Gamma(n_{(\cdot),v}^{k,-(m,n)} + \beta_v + 1)}{\Gamma((\sum_{r=1}^{V} n_{(\cdot),r}^{k,-(m,n)} + \beta_r) + 1)} \\
&\propto \frac{\Gamma(n_{m,(\cdot)}^{k,-(m,n)} + \alpha_k + 1)}{\Gamma((\sum_{i=1}^{K} n_{m,(\cdot)}^{i,-(m,n)} + \alpha_i) + 1)} \frac{\Gamma(n_{(\cdot),v}^{k,-(m,n)} + \beta_v + 1)}{\Gamma((\sum_{r=1}^{V} n_{(\cdot),r}^{k,-(m,n)} + \beta_r) + 1)} \\
&= \frac{\Gamma(n_{m,(\cdot)}^{k,-(m,n)} + \alpha_k)(n_{m,(\cdot)}^{k,-(m,n)} + \alpha_k)}{\Gamma(\sum_{i=1}^{K} n_{m,(\cdot)}^{i,-(m,n)} + \alpha_i)(\sum_{i=1}^{K} n_{m,(\cdot)}^{i,-(m,n)} + \alpha_i)} \frac{\Gamma(n_{(\cdot),v}^{k,-}}{\Gamma(\sum_{r=1}^{V} n_{(\cdot),r}^{k,-}} \\
&\quad n_{m,(\cdot)}^{k,-(m,n)} \quad \alpha \quad n_{m,(\cdot)}^{k,-(m,n)} \quad \alpha \quad \Gamma
\end{aligned}
$$

What I really do

```
open Accord
open Accord.Statistics.Models.Regression
open Accord.Statistics.Models.Regression.Fitting


let regression = new LogisticRegression 2
```

https://lenadroid.github.io/posts/machine-learning-fsharp-accorddotnet.html

Summary and the Road Ahead

The Bottom Line:

- AI & ML have both been hyped by researchers & techies
- Separating the wheat from the chaff is quite challenging

Some Problems Being Worked on:

- Using the increasing computational power for better ML
- Developing better neural-network and other models
- Improving joint decision-making by humans + machines
- Ethics of AI & ML; ethics of technology, more generally

Going Forward:

- AI need not be human-like to be useful (factory robots)
- It's not humans vs. machines, but humans + machines
- The singularity isn't near (or if it is, no need to worry)

Further Reading

Artificial Intelligence: A Guide for Thinking Humans,
by Melanie Mitchell, Macmillan, 2019

Table of Contents (16 chs.)

Part I Background
Part II Looking and Seeing
Part III Learning to Play
Part IV Artificial Intelligence Meets Natural Language
Part V The Barrier of Meaning

Interweaves stories about the science of AI and the people behind it

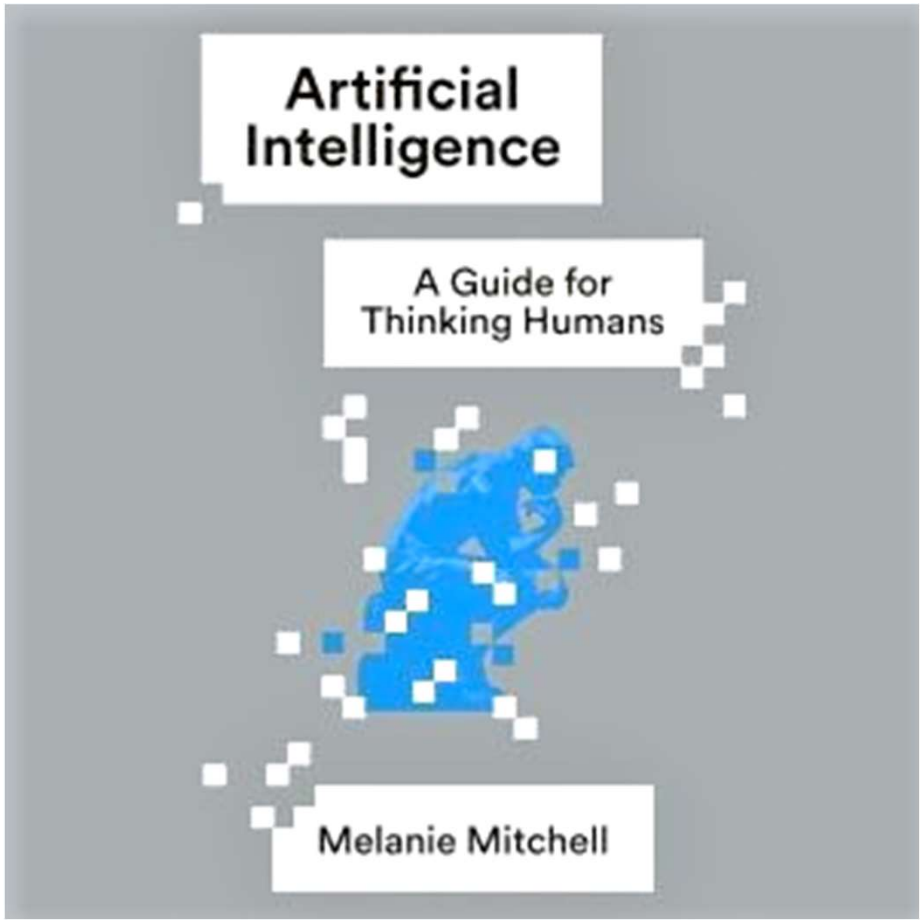

Humans in a World Dominated by AI

Life 3.0: Being Human in the Age of Artificial Intelligence **(Max Tegmark)**

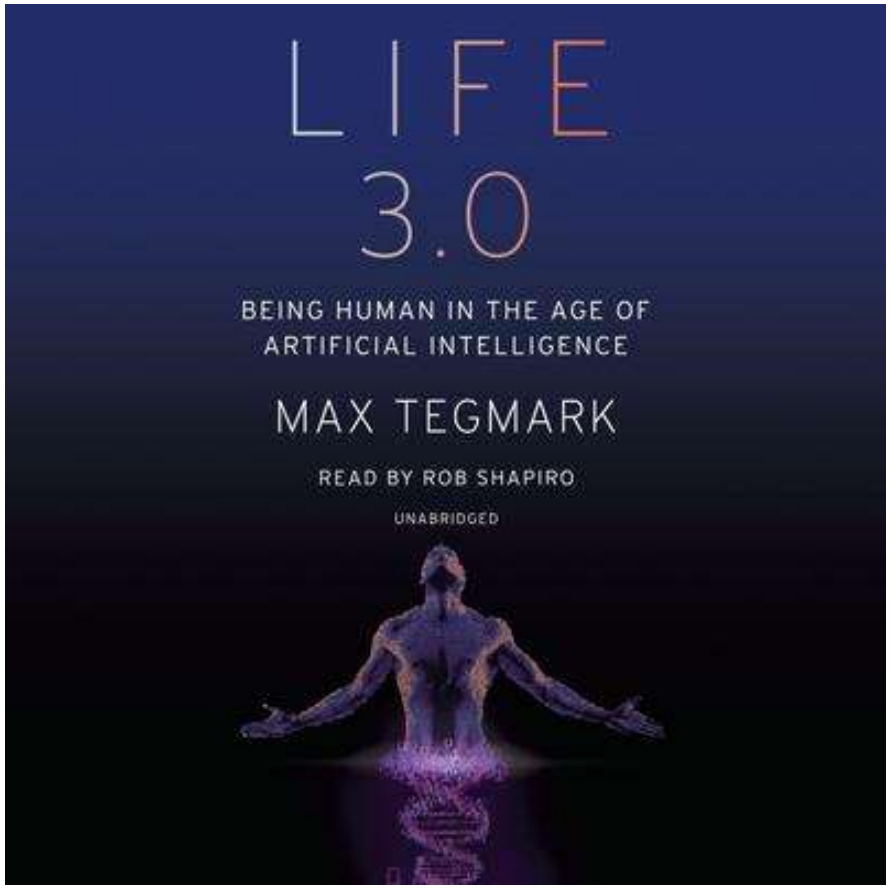

Superintelligence: Paths, Dangers, Strategies **(Nick Bostrom)**

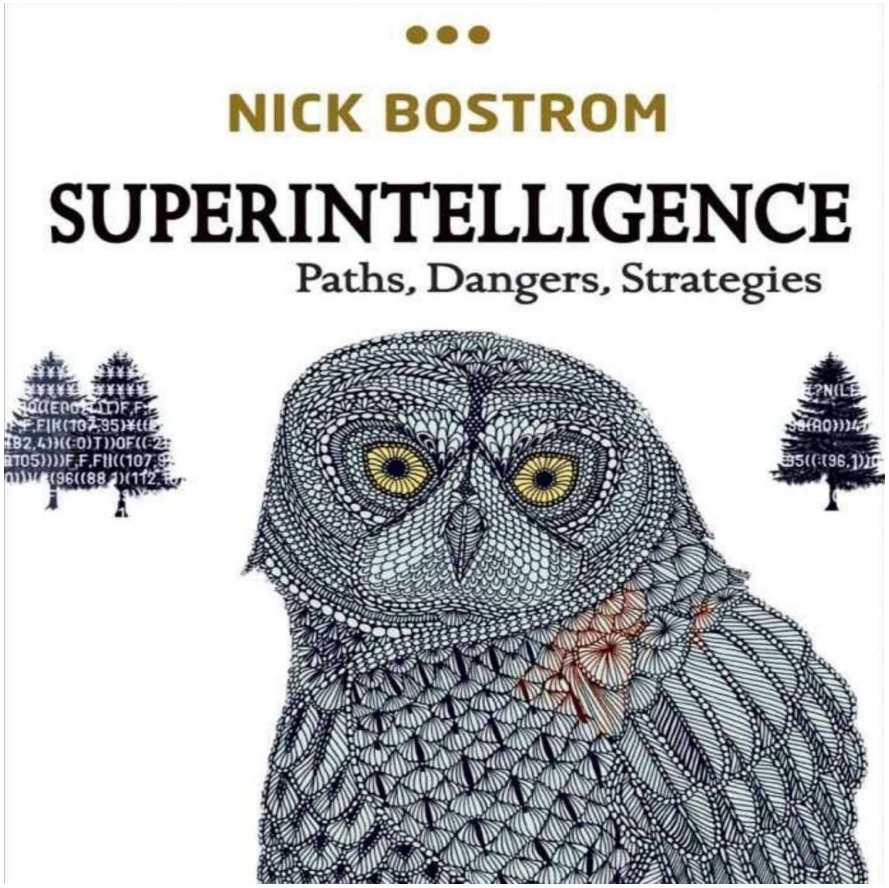

Questions or Comments?

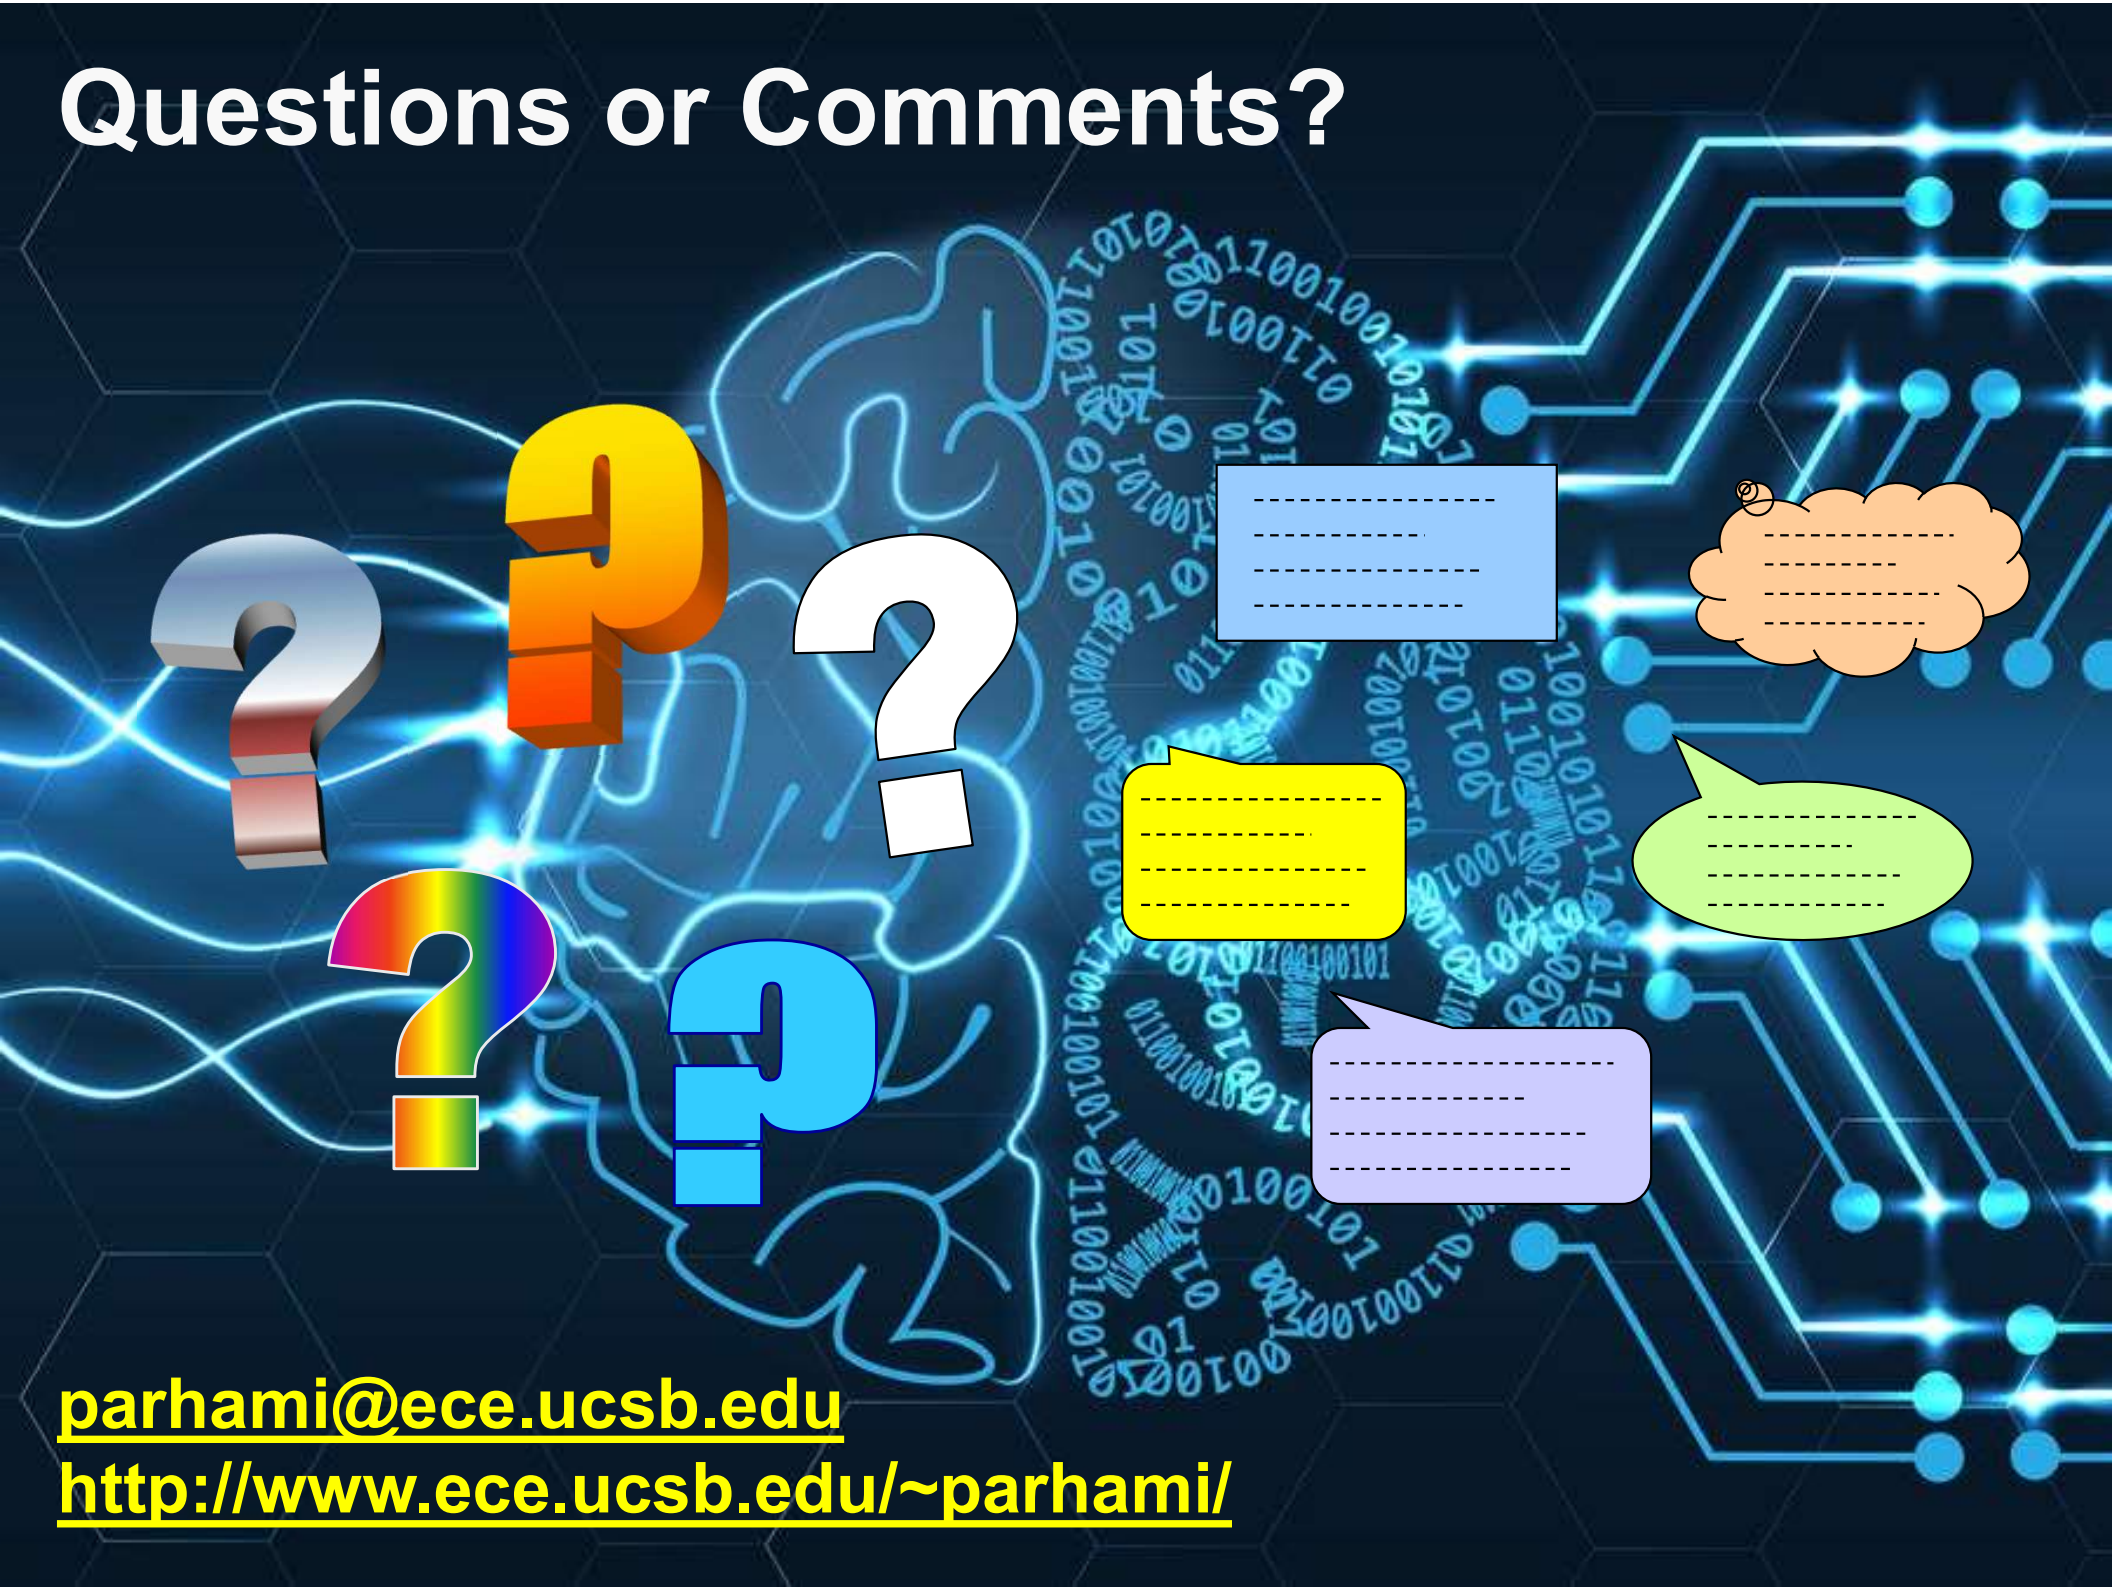

parhami@ece.ucsb.edu
http://www.ece.ucsb.edu/~parhami/

Slides from previous versions of this talk

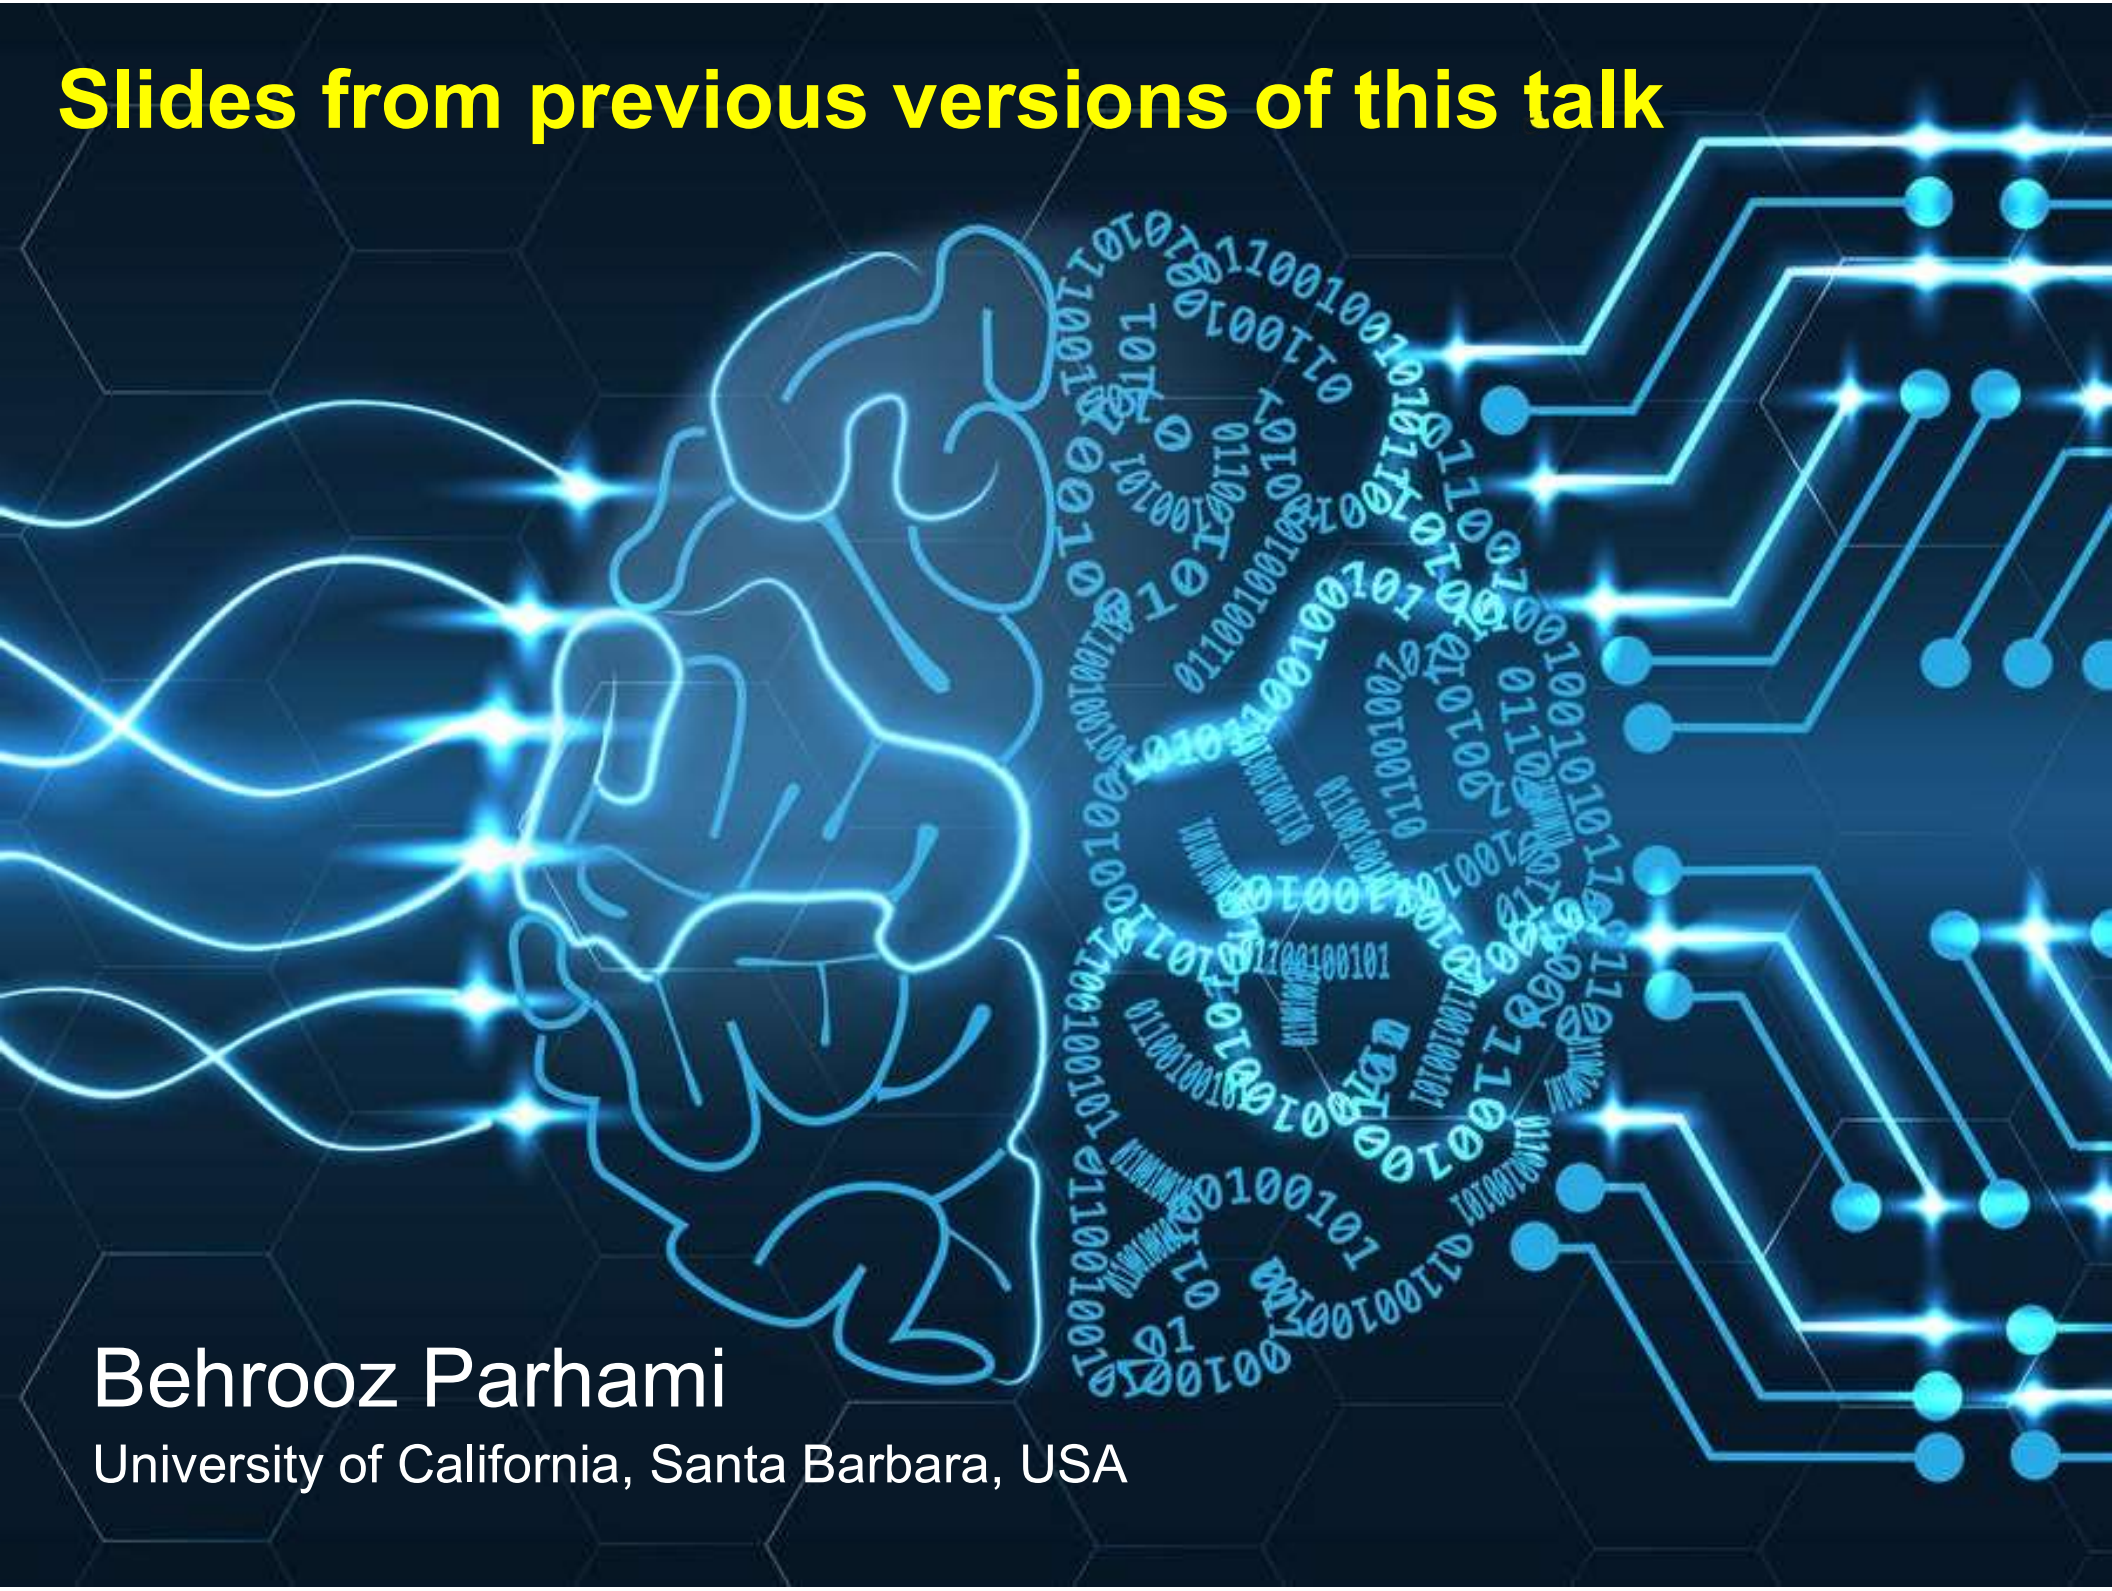

Behrooz Parhami
University of California, Santa Barbara, USA

Let's Not Call Everything Artificial Intelligence:

A Realistic Assessment of Intelligent Behavior and Machine Learning

Behrooz Parhami

University of California, Santa Barbara, USA

لطفا به همه چیز برچسب هوش مصنوعی نزنید:

تحلیلی واقع گرایانه از رفتار هوشمندانه و یادگیری ماشینی

تعریف هوش مصنوعی، و هوش به طور کلی، کار آسانی نیست. خیلی‌ها در این راه پافشاری نمی کنند و فقط می گویند که با وجود عدم توانایی در تعریف رفتار هوشمندانه، شناخت آن در عمل امکان پذیر است. تلاش در راه دستیابی به هوش مصنوعی عمومی نشیب و فرازهایی داشته است، ولی پیاده سازی رفتارهای هوشمندانه در میدانهای نسبتا محدود موفقیت آمیز بوده است. مثالهای این امر شامل مسیریابی و زمان بندی خودروها (اوبر، اسنپ)، برنامه ریزی و تدارکات (تخصیص دروازه های فرودگاه به هواپیماها)، و بازیهای فکری (شطرنج، گو) است. با این حال، تا ساختن ماشینهایی که آزمون تورینگ را با موفقیت پشت سر بگذارند و یا بتوانند ترجمه ماشینی با کیفیت بالا ارایه دهند راه درازی در پیش است. در سال‌های اخیر پژوهشگران هوش مصنوعی به این نتیجه رسیده اند که شاید مشکل تر از رسیدن به هوش مصنوعی عمومی، آگاهی رسانی و مواجهه با آثار اجتماعی داده های کلان و تصمیم گیری ماشینی باشد. در پی جمع آوری حجم عظیمی از داده ها راجع به هر یک از ما، تضمینی وجود ندارد که ماشینها یا صاحبانشان از این داده ها برای بهینه سازی سرویس های ارایه شده، و نه کنترل کردن و شکل دادن به رفتار های اقتصادی و اجتماعی ما، استفاده کنند. بنابراین، یکی از چالشهای مهم پیشاروی ما ایجاد توازن بین تسهیل پیشرفتهای علمی و فنی از یک سو و تضمین انصاف و عدالت اجتماعی از سوی دیگر است.

Let's Not Call Everything Artificial Intelligence:
A Realistic Assessment of Intelligent Behavior and Machine Learning

Defining artificial intelligence (AI), or plain intelligence for that matter, has proven more difficult than expected. Many people have thrown up their arms, taking the position that, even though we can't define AI, we'll recognize it when we see it! Despite the cycles of hype and disappointment in achieving general AI, success stories abound in making machines behave intelligently in limited domains. Examples include vehicle routing (Uber), logistics (airport gate assignments), and game-playing (Chess, GO). Meanwhile, we still have a long way to go in building machines that can pass the Turing test, as well as in domains such as machine translation, which may require the same, or even greater, capabilities. In recent years, we have come to realize that, as great as the technical challenges are in developing general AI, an even greater challenge is developing awareness and dealing with social implications of massive data repositories and automated decision-making. After collecting petabytes of data on each of us, there is no guarantee that machines, or their masters, will use the data to offer better services and optimal outcomes, rather than controlling and shaping our economic and social behaviors. A key consideration is thus ensuring a balance between facilitating technical progress and ensuring fairness and social justice.